European Journal of Behavior Analysis

ISSN: 1502-1149 (Print) 2377-729X (Online) Journal homepage: https://www.tandfonline.com/loi/rejo20

Noncriterial behavioral variability and related topographic bias in humans

Laurilyn D. Jones & Francis Mechner

To cite this article: Laurilyn D. Jones & Francis Mechner (2020): Noncriterial behavioral variability and related topographic bias in humans, European Journal of Behavior Analysis, DOI: 10.1080/15021149.2020.1745526

To link to this article: https://doi.org/10.1080/15021149.2020.1745526

Published online: 07 Apr 2020.

EUROPEAN JOURNAL OF BEHAVIOR ANALYSIS
https://doi.org/10.1080/15021149.2020.1745526

ARTICLE

Noncriterial behavioral variability and related topographic bias in humans

Laurilyn D. Jones [a,b] and Francis Mechner [b]

[a]Department of Behavioural Sciences, Oslo Metropolitan University, Oslo, Norway; [b]The Mechner Foundation, New York, NY, USA

ABSTRACT

All operant behaviors have multiple characteristics in addition to those criterial for reinforcement, and variation occurs across all. All such characteristics can also reflect topographic bias due to historic and physiological factors. The revealed operant is constructed so that topographic aspects and variation are measurable. In two experiments humans performed a revealed operant response of 14 or more keystrokes. The first and last were mandated, while the middle 12 or more were allowed to vary. There were significant differences in variability among participants, as well as systematic effects of the experimental designs. Despite not being reinforced, variability among complete sequences was high. Test conditions in Experiment 2 resulted in a much larger increase in variability than did suspension of reinforcement in Experiment 1. There was systematic topographic bias both for and against letter keys in the center of the keyboard. There were also correlations between measures of variability and bias.

ARTICLE HISTORY
Received 22 May 2019
Accepted 18 March 2020

KEYWORDS
Behavioral variability; noncriterial variability; stereotypy; topographic bias; bias; human participants

The operant can be thought of as a basic unit in the experimental analysis of behavior: every individual occurrence or instance of an operant response is also an exemplar of a larger operant class (Skinner, 1959). In order to be considered part of any given operant class, the behaviors comprising the operant response must satisfy the criteria that define that class – for example, depressing a lever in an experimental chamber with a minimum amount of force – but all other characteristics of the operant response, such as the duration of the lever press, are free to vary and frequently do.

All operant classes, no matter how simple, can be thought of as sequences of behaviors (Mechner, 1994; Mechner et al., 1997). Some behaviors in the sequence are specified by the experimenter as criterial for reinforcement, and all the others can be termed the noncriterial dimensions of the given operant class (Herrnstein, 1966). These latter behaviors are sometimes referred to in the literature as noncontingent or non-instrumental, but for purposes of clarity, the term noncriterial will be used throughout this paper.

All dimensions of the operant response, even noncriterial ones, have the potential to be affected by reinforcement (Neuringer, 2002). For example, in what is often referred to as superstitious behavior, aspects of the operant response designated by the experimenter

CONTACT Laurilyn D. Jones ldj@mechnerfoundation.org

as noncriterial (i.e. not required for reinforcement) can nevertheless be shaped by adventitious reinforcement that controls the form of the organism's responses and thus potentially affects the reinforcement received in turn. That is why an experimental methodology that permits measurement of at least some of the noncriterial dimensions of each operant response allows an experimenter to analyze behavior in far greater depth than methodologies that only record the emission of each response as a single data point. This advantage is particularly significant when studying complex topics such as behavioral variability vs. stereotypy, and the existence of participant bias relating to a given operant response, hereafter referred to as topographic bias, to distinguish the phenomenon from the more commonly-referenced stimulus bias.

All organisms display biases, of course, in the sense that they favor certain behaviors over others. The term topographic bias refers to any consistent systematic favoring of one functionally equivalent behavior over another, which cannot be explained by the current contingencies (Jones & Mechner, 2013, 2015). It can have many different causes, for example, the anatomy or physical capabilities of the organism, which may make some behaviors easier than others; the organism's learning history; or topographical constraints on behavior due to the apparatus in use, to name only a few. These causes may or may not – usually not – be known to the experimenter.

The *revealed operant* is an experimental methodology first proposed, and discussed in detail, in the 1994 monograph of the same name (Mechner, 1994). It allows for the measurement of both criterial and certain noncriterial attributes of each individual operant occurrence of an operant class and consists of behavioral sequences that are programmed and tracked by the experimenter (Mechner, 1994, pp. 11-12; Mechner et al., 1997). Each revealed operant unit consists of a sequence of recorded actions, or sub-operants. Some of these sub-operants are mandated by the definition of the operant (criterial), and others are not (noncriterial). The first and last sub-operants in the sequence are programmed to be behaviorally distinct so as to demarcate consecutive sequences from each other and prevent fusing. It is important to note that each revealed operant sequence also functions as a single behavioral unit – a single operant response, rather than a chain of individual operant responses. As a research preparation, the revealed operant is thus ideal for studying noncriterial variability as well as topographic bias, as noncriterial sub-operants within the sequences can vary without affecting the validity, or the reinforcement, of the operant response as a whole. The experiments presented in this paper used a revealed operant class consisting of a sequence of keystrokes on a modified computer keyboard to measure both variability and topographic bias with respect to typing behaviors, and to discover if there were quantifiable interactions between the two.

Variability, both criterial and noncriterial, has been studied extensively (see Neuringer, 2002 for one review). Variability in many different types of *noncriterial* operant dimensions has been shown to increase in extinction in both animal and human studies (Antonitis, 1951; Iversen, 2002; Kinloch et al., 2009; Margulies, 1961; Morgan & Lee, 1996; Neuringer et al., 2001; Notterman, 1959). It is also possible to generate higher levels of noncriterial variability by changing the response requirement (Mechner et al., 1997; Tatham et al., 1993). In general, intermittent reinforcement often results in greater noncriterial variability than continuous reinforcement (Antonitis, 1951; Boren et al., 1978; Eckerman & Lanson, 1969), though the literature also contains examples of contrary results (Gates & Fixen, 1968; Herrnstein, 1961a; see R. Lee et al., 2007, for a review).

In contrast, a search of the behavior analytic literature shows relatively limited investigation of systematic participant bias as a phenomenon in its own right. Researchers generally attempt to reduce potential bias for or against the responses they use in their experimental designs, just as they attempt to reduce the effects of uncontrolled variables (Sidman, 1960). The study of bias does, however, occur in the experimental analysis of matching. Since Herrnstein's (1961b) formulation of the matching law, the frequently-recorded deviations from perfect matching are commonly referred to as either undermatching or overmatching. Baum (1974) called these effects bias, and included a variable of that name as part of the equation describing how closely the relative rate of different responses match their relative reinforcement rates. Mechner (1994, pp. 40-41) explained that such bias effects can be due to noncriterial topographic variations in response duration, where longer durations necessarily create undermatching and shorter durations overmatching. Very few matching studies have focused on response topography as an experimentally-controlled or recorded variable. One of the few that did so, Sumpter et al. (1995), showed topographic bias in an experiment comparing an operant consisting of key pecks to one that required hens to push a "door" (two rods which the birds could put their heads between). Their experiments found consistent and predictable bias for the key peck over the door push.

Within the variability literature, researchers occasionally also examine noncriterial topographic bias as a byproduct of their procedures. Antonitis (1951) found a general bias toward the center of the horizontal response slot for his rats during conditioning; he also found that individual rats showed strong idiosyncratic biases for particular preferred positions along the slot. Eckerman and Lanson (1969) also found a center bias in pigeons, but as in Antonitis' result, this was probably due to the location of the feeder. In that study there was also, however, a very slight systematic bias toward the right side of the apparatus, which persisted throughout the experiment.

Noncriterial variability and topographic bias are linked phenomena that often occur together. There is very little quantitative data in the behavioral literature on bias, and no results exploring whether noncriterial variability and topographic bias can have interactive effects on behavior. In order to generate the relevant quantitative data, the experiments presented in the present paper use the revealed operant methodology.

In both experiments, human participants learned to execute an operant sequence in which almost all of the behaviors were noncriterial and allowed to vary, thus providing a large amount of data that could be analyzed in multiple ways. Other than the criterial operant sequence itself, which remained exactly the same under all conditions, the contingencies put in place during Experiment 2 were designed to be very different from those used during Experiment 1, allowing a further analysis of the effect of different experimental variables on noncriterial variability, topographic bias and combinations of the two.

Experiment 1

Method

Participants

12 adults over the age of 18 were recruited to participate in the study by means of flyers posted on local college campuses. In the initial telephone screening interview, participants were informed that they could earn up to a total of 300 USD by completing ten experimental sessions, each approximately an hour in length and each taking place at the same time of day on ten consecutive days. Compensation was set at this relatively high level to prevent the loss of participants that had occurred in previous multi-day studies.

Participants received a flat fee of 10 USD per session for their participation, which they were given in cash at the end of each session. The remainder of their compensation consisted of the monetary reinforcers earned during the experiment, which totaled 207.06 USD over the ten days. However, they did *not* receive each day's earnings on that day but instead were paid the total as a lump sum by check only after the completion of the final session. Again, this procedure was put in place to discourage dropouts.

All participants were informed when they signed up that they were free to withdraw from the experiment at any time and keep the already-earned 10 USD per session participation fees, but would receive the final payment only if they completed all ten sessions. Of the 12 participants initially recruited, 4 still withdrew during the course of the experiment, leaving a total of 8 who completed the entire study and whose results are presented here.

Participants were asked to maintain relative consistency in sleeping, eating and caffeine consumption during the ten days of the study. This behavior was tracked by having them fill out a form detailing these variables every day. This was done in order to equalize the establishing operations preceding each session of the experiment as much as possible, and the particular variables used were chosen due to their well-documented effect on human performance. Due to the nature of laboratory experiments with human participants, it is difficult to control these types of variables in any more systematic way, but requiring documentation of their behavior between sessions may help increase participants' awareness of the need to keep such behavior constant. At the beginning of the first session, all participants received a written explanation of all requirements of the study as well as their right to withdraw at any time, and signed a document giving their informed consent to participate in the study under these requirements. After they had completed the last session of the experiment, they filled out a questionnaire on their experience, and the experimenter then debriefed them fully.

Apparatus and setting

Four Dell 486 desktop computers were used as the apparatus in both experiments, placed at four computer workstations separated by screens. Each had a 14-inch CRT monitor placed at slightly below eye level, and a keyboard on a keyboard tray at a comfortable height for typing. The keys in use during this experiment were limited by means of a custom-made particleboard mask, shown in Figure 1, which covered the entire keyboard with the exception of 12 letter keys (T, Y, U, I, G, H, J, K, V, B, N and M), the space bar, enter key, number keypad and a few function keys.

Figure 1. Picture of the keyboard used in both experiments, showing the subset of 12 letter keys in use during the experiments.

The 12 letter keys used were chosen to fall between the normal positions of the left and right hands when touch typing. The keyboard mask thus made it harder for the participants to use their normal hand positions. All the experimental software was written using the Euphoria programming language. It controlled the experimental input and output, provided the needed visual and auditory stimuli, and tracked every keystroke performed by the participant.

Experimental design

In both Experiments 1 and 2, the operant class performed was a sequence of 14 or more keystrokes as follows: a press on the space bar to start, followed by *at least* 12 letter keypresses from those available on the masked keyboard, followed by the enter key at the end. For most participants, after some practice this sequence took in the range of two to three seconds to execute. A feedback stimulus was programmed to help participants keep track of the responses they emitted: the entire screen turned blue when the space bar was pressed to begin each operant response, returned to black with the press of the enter key, and remained black until the participant began to perform the next operant response. While the screen was blue, acceptable keystrokes produced a subtle click feedback noise 100 milliseconds in duration. At no point did the monitor display the characters typed by the subjects.

The specific letters making up the response sequence (other than the space bar and enter keys) were not mandated and were allowed to vary freely except for one restriction: a modified within-response-only version of a Lag 2 requirement, under which participants were not allowed to type any letter identical to either of the two previous letter keypresses *within that sequence*. For example, typing “gg” as part of a given operant sequence was never acceptable at any time during either experiment; neither was “ghg”. However, “ghjg” was allowed under this restriction, representing the minimum acceptable level of within-response variability. If these restrictions were violated, the repeated letters were not counted toward the total of 12 required nor did they generate the soft click noise programmed as feedback for valid keystrokes. Comparison of the current letter key to the two previous ones only occurred within each individual sequence. At no time was this modified Lag 2 requirement applied to letters from more than one operant sequence. This feature of the operant class design was chosen for two reasons: it generates enough moment-to-moment variability to allow measurement of potential bias for or against specific letter keys, and it serves as an analog of a procedure used in certain studies of variability that require switching between keys during the emission of each sequence (Machado, 1997). Thus, in both experiments, a certain minimum level of variability *within* each operant sequence was a criterial dimension of the operant class, but the level of variability among complete responses was completely noncriterial and did not affect reinforcement in any way.

Responses were classified as invalid if terminated by pressing the enter key before the participant had emitted at least 12 keystrokes fulfilling the response requirement, and such responses were not reinforced, regardless of how many total keystrokes they contained. However, it was technically possible for a participant who realized their mistake in mid-response to simply type additional letters fulfilling the variability restriction before pressing the enter key, which was always required to end a sequence. There was no upper limit on the number of keypresses per operant response, so long as at least 12 acceptable keypresses were registered before the enter key was pressed. Each response could contain as many key presses as the participant chose. This ensured that each response would end only with the behaviorally-distinct terminal sub-operant (the enter key press), rather than resetting after a certain number of keystrokes. It is important for each revealed operant sequence to be distinct, rather than fused with the next one.

Experiment 1 was designed to determine if the behavior would settle at some level of stability (i.e, with each participant’s overall level of variability remaining within the same general range from session to session) after many sessions of intermittent reinforcement, before testing the stability of the performance by instituting an extinction protocol – in other words, terminating reinforcement. Unpredictable reinforcement, especially in experimental situations which allow for variability, can potentially increase variability levels. On the other hand, it can also lead to the development of idiosyncratic “superstitious” stereotyped response patterns. A variable ratio schedule was selected for the baseline condition as it encourages higher rates of responding than do interval schedules. For the test condition, the withholding of reinforcement was chosen in order to see whether this operation would replicate the well-documented phenomenon of extinction-induced increase in variability with this type of operant class.

Baseline condition with procedure

The first nine sessions of Experiment 1 were programmed for 440 responses each. Because only valid operant responses – those meeting the criteria explained above – were counted toward this total, the actual length of each session varied slightly, depending on how many mistakes each participant made that day. Because participants worked at different speeds, the duration of each session also varied. Valid responses were reinforced throughout the first nine sessions on a VR4 schedule, programmed in an unpredictable pattern that repeated every 80 responses while still keeping the number of reinforcers earned relatively consistent throughout the session. The reinforcer used was money, programmed to vary among 11, 19 or 26 cents on an unpredictable schedule. The 110 reinforcers programmed for each of the first nine sessions added up to 21.34 USD per session, making the average reinforcer amount 19.4 cents.

Reinforcement was signaled by a 440 Hz beep sound emitted for .125 seconds, accompanied by a message appearing in the middle of the computer screen, which read "You just earned X cents. Ring it up." A response, analogous to consumatory responses for reinforcers used in many animal experiments, was required after each reinforcer presentation, consisting of typing the amount earned on the number keypad, followed by a press of the enter key. Whenever participants performed this post-reinforcer response, the amount earned was added to a running total, which was displayed at all times throughout the session in the top right corner of the screen. Participants were unable to start another operant response by pressing the space bar until they had correctly performed the post-reinforcer response, as the experimental software would not respond to any but the required input. They were thus always aware of each reinforcement, and of the total amount they had earned at any point during a given session.

In addition to this reinforcement schedule, for the first 240 operant responses of the first session only, the computer screen also displayed a large (3" x 3") green square for .5 seconds after *every* valid response as feedback, regardless of whether or not the response received reinforcement. This was done in an attempt to ensure that the intermittent nature of the programmed reinforcement contingency did not impede the participants from learning how to produce consistently valid operant responses.

When each participant arrived for the first session they were first asked to read and sign the written agreement specifying the requirements of the study and the rights of the participant. They then took a seat at a computer workstation to participate in a demo of the software. The experimenter instructed the participants on how to perform each operant response by saying: "Always press the space bar first, then type 12 or more letters, then press the enter key to finish. Press the space bar again to start the next sequence, and so on. The 12 or more letters can be anything you want, except that you can't repeat either the last letter you just typed, or the one before it during each pattern. When the green square appears on the screen you know you did it correctly. Then you can keep going, and do it over and over again."

When the participant triggered the first reinforcement in the demo, the instructor said, "Sometimes when you type sequences correctly you will earn money. Every time that happens, the computer will show you a message, like this one, and say 'Ring it up.' Just type the amount you earned on the number keypad and press enter. The computer will add it to your total, which you can see here in the corner of the screen. Then you can keep going and type another sequence." When each participant could reliably generate valid responding and perform the post-reinforcer response the demo ended, and the participant logged in to begin the first session.

On days two through nine of the study participants signed in with the experimenter when arriving on the premises and then started their sessions themselves. If they talked at any time during any of the sessions, or used their phones, they were asked to be quiet and shut off all electronics while working.

Test session with procedure

Session 10 consisted of 500 revealed operant responses, again with only valid ones being counted toward the 500. In this session no reinforcers were given, nor was the green square presented, but all other elements of the experiment were identical to those in the previous sessions. No instructions were given regarding the new contingency, and the procedure followed by experimenter and participants was exactly the same as in the previous sessions. If participants had questions or complaints during Session 10 – which many of them did – the experimenter simply said "Please continue working."

Experiment 1 results

Noncriterial variability

An operant class consisting of a lengthy sequence of discrete, highly variable behaviors that are precisely tracked, such as that used in these experiments, offers a preparation for studying noncriterial variability in detail through the use of multiple measurements. One statistic that helps reveal differing levels of variability for this type of operant class is the proportion of each participant's responses in each session that are unique: i.e., the fraction of the total operant responses emitted during the session that remains if one removes all repetitions of responses previously emitted during the same session. Only the first – or only – occurrence of each unique operant response is counted toward this value; if it is repeated within the same session by the same participants, the repetitions are not considered. This ratio of unique sequences to total sequences emitted within each session might be termed the Unique Value, and it serves as a useful analog for the commonly used U-value measure of variability, as it too is a single value falling somewhere between 0 (perfectly predictable) and 1 (perfectly unpredictable). There are certainly limitations to the information provided by this statistic (just as there are for U-value – see Barba, 2012; Kong et al., 2017), but for the current experiment it provides an excellent indicator of each participant's overall level of variability in each session.

Figure 2 presents this Unique Value statistic for each participant in Experiment 1. Interestingly, this graph shows an almost perfectly bifurcated data set: half of the participants (shown with open data points and dashed lines) remain highly variable throughout the experiment, with extremely high Unique Values approaching 1, while the other half (represented with closed data points and solid lines) start out more variable in Session 1, when learning the task, then drop very quickly to relatively low levels of variability which they maintain for the remainder of the study. The split between the two groups occurs after the first or second session and remains consistent thereafter. The double vertical line between Sessions 9 and 10 indicates the change in experimental contingencies. There is no large increase in variability in Session 10 when the condition change was imposed.

In spite of the significant differences among participants, some fairly consistent effects of elapsed time, repetition and reinforcement on noncriterial variability levels can be

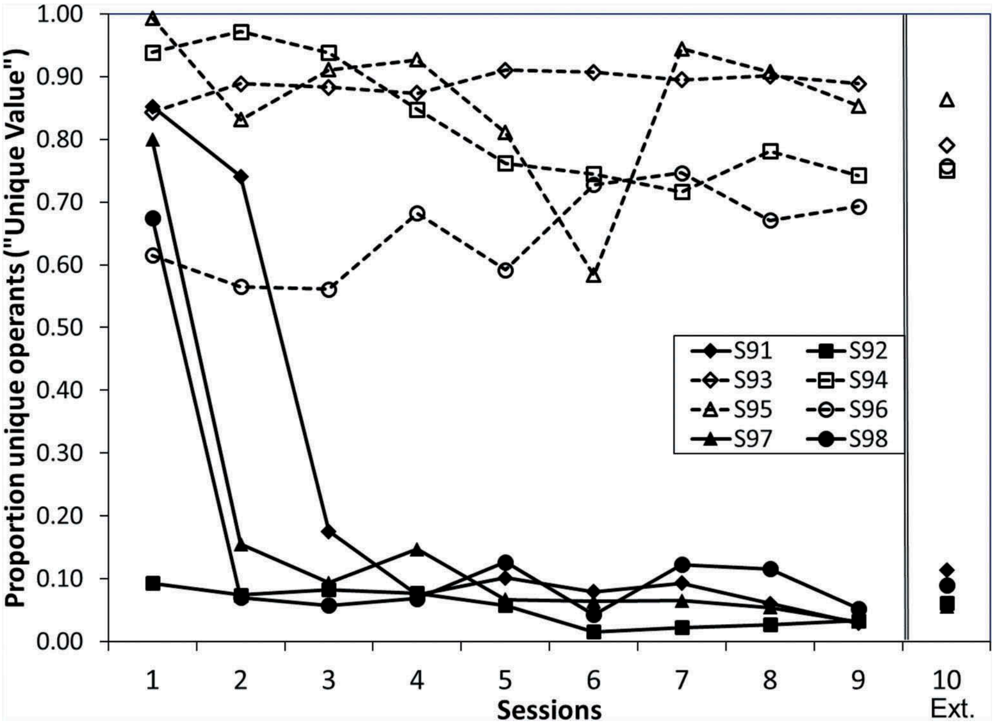

Figure 2. Experiment 1: Proportion of the total responses emitted by each participant in each session that are unique (Unique Value). The ten sessions of the Experiment are on the X axis. The high-variability group's data is shown with dotted lines and open markers, while the low-variability group's data is graphed with solid lines and closed markers. The double vertical line between sessions 9 and 10 shows the condition change.

seen. For the four participants with low Unique Values (hereafter referred to as the low-variability group), noncriterial operant variability tends to start higher in the first session, decline sharply in the second and third sessions before stabilizing, then rise *very* slightly in the final session. Notably, these session to session trends are almost completely absent in the four high-Unique-Value participants (hereafter termed the high-variability group), who are more variable across sessions, across a much wider range of Unique Values than the more tightly-grouped low variability participants, and do not show a systematic increase in noncriterial variability during the final session.

In addition to looking at Unique Value as a measurement of noncriterial operant variability, we can also measure the variability in *length* of the operant responses emitted, because participants were allowed at all times to type as many letters as they wanted within the body of each response. Figure 3 shows the proportion of total operant responses in each session of Experiment 1 that were longer than the minimum 12 letters required (thus representing variation from the norm), for all participants. Note that the three participants who emit the longest responses throughout the study are three of the four participants in the high-variability group (open data points, dashed lines). In fact, the difference in operant length between the high-variability and low-variability groups is statistically significant using a one-way ANOVA (*F*(1,

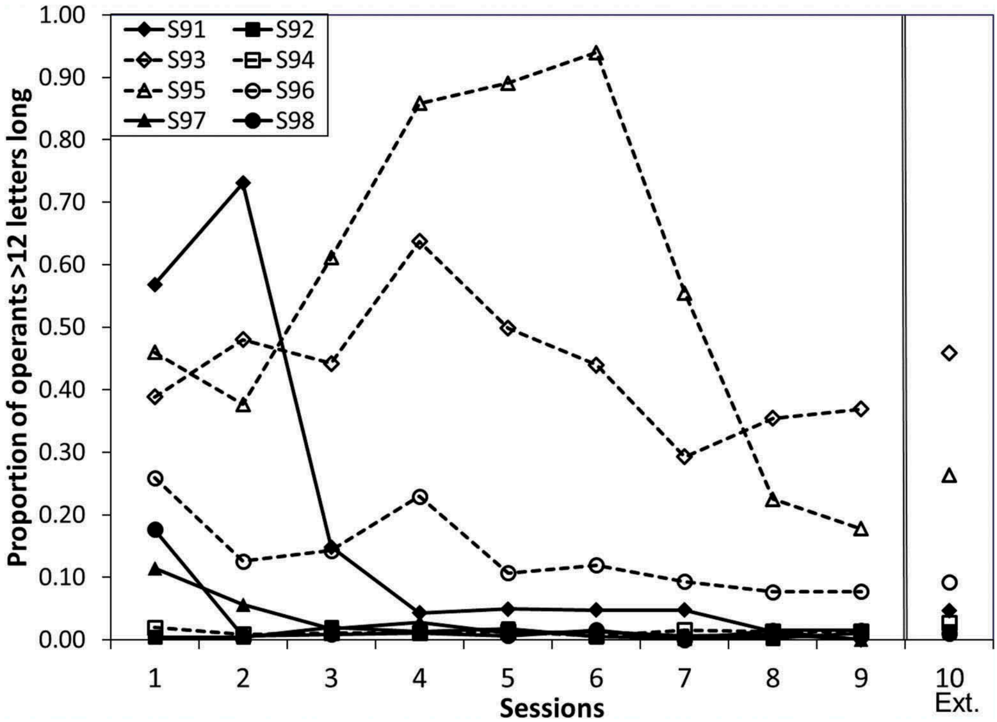

Figure 3. Experiment 1: Proportion of the total number of responses emitted by each participant in each session that are more than 12 letters long (the minimum criterial length).

78) = 21.98, p = .00001). The session-to-session trends are very similar to those seen in the Unique Value data: in the low-variability group, operant length tends to start higher in the first session, decline quickly, then stabilize, while participants in the high-variability group are more variable in operant length from session to session and cover a wider range of values.

Finally, although the accuracy with which this type of operant class was performed is not a direct measure of variability, it is quite closely related to it, as variability in general often occurs in conjunction with mistakes or disruptions in the behavior stream (Mechner & Jones, 2011, 2015). In all sessions of Experiment 1, the proportion of invalid operant responses (those not conforming to the criterial requirements) that were unique within the session in which they were emitted was much higher than the proportion of unique responses in the session as a whole. The average difference in Unique Value between total and invalid operant responses is .13 for the high-variability group of participants and .69 for the low-variability group. This effect is robust, being found in all participants.

In addition, the total number of errors that each participant made during each session of the study is also correlated with that participant's level of variability. Figure 4 shows the percentage of each participant's operant responses that were invalid during each of the 10 sessions. Although only two of the four high-variability participants' values (open data points, dashed lines) are visibly much higher than the others, the difference in

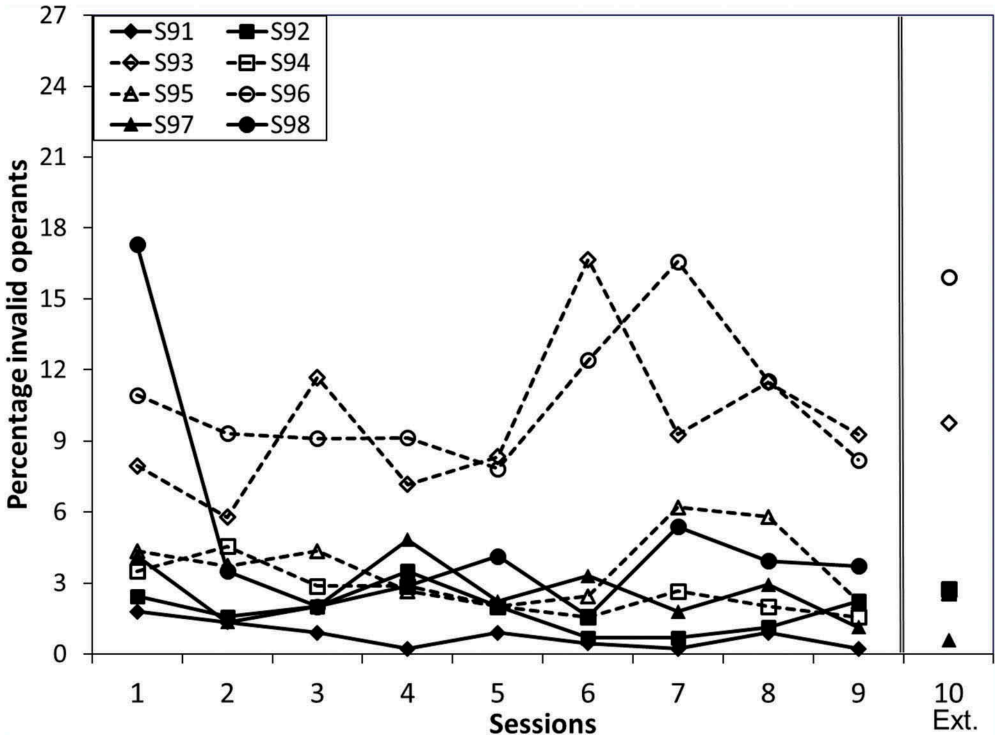

Figure 4. Experiment 1: Percentage of responses emitted by each participant in each session of that are invalid (not counted toward the variable ratio required for reinforcement). Please note that due to the low error rates overall the Y axis is expanded rather than scaling from 0 to 100, in order to show the differences between participants.

accuracy between the two groups of participants is still statistically significant ($F(1, 78) = 27.64$, p = .000001).

Interestingly, the consistent differences between the two distinct groups of participants in Experiment 1 seen in the data are also reflected in their answers on the post-experimental questionnaire. The four participants in the high-variability group, when asked "How did you select the keystrokes required?", gave the following answers:

- "All I could really see was that the more spread out and different the patterns were from each other the more money you would make."
- "I tried to use different letters … . I feel I made more money when I made a mistake."
- "At first it was random but then I selected them by process of elimination – starting on different letters and then moving down from there."
- "I guess the first few sessions I'd key in all the letters in various patterns . . . . I'd key in whatever patterns on the keyboard just to keep it from getting monotonous."

In contrast, the four participants in the low-variability group provided the following answers when asked the same question:

- “I just pressed the keys one by one, and I continued to do the same during the whole experiment.”
- “I just stuck with the same pattern until I figured it was near the end of a session then I would change it.”
- “I started repeating a system beginning from the top-left of the keyboard and finishing at the bottom-right of it. Sometimes I changed the system to see if any changes in the amount of money occurred.”
- “Of course it was more efficient to press the same buttons … . I don’t press spontaneously, instead I have 6 different systems connected to each other.”

Bias

The nature of this experiment, in which participants were allowed to type whichever letters they liked from among the set of 12 provided, also lends itself quite well to an analysis of human topographic bias as it affects typing behavior. In order to measure bias levels, we first calculated which of the twelve letter keys available to the participants were chosen more or less often, i.e., were favored over others. There are significant preference patterns for certain of the twelve keys. These do not appear to be based on individual letter preference but instead are strongly linked to the spatial position of the letter keys on the keyboard. In addition, they vary based on the ordinal position of the particular keystroke within the operant sequence. The data are thus presented organized in that manner (to see the keyboard again, please refer back to Figure 1).

Throughout Experiment 1, the participants demonstrated biases with reference to their use of the letter keys in the *middle or center* of the section of the keyboard available for typing, with corresponding biases affecting choice of the keys on the edges of that space. In particular, we looked at bias for or against letter keys from the middle of the three rows (letters g, h, j and k), vs. those on the top and bottom rows (t, y, u, i, v, b, n, and m). If the participants had typed letters at random, i.e. with no biases for or against specific letters, we would expect them to have selected approximately one third (.33) of their keystrokes from the middle row, as those letters represent 4 out of the 12 available ones. We also measured bias for or against those letter keys in the center vertical section of the keyboard (letters y, u, h, j, b, and n) vs. those on the left or right sides (t, g, v, i, k, and m). Again, if participants had chosen letter keys randomly, we would expect approximately half (.50) of the keystrokes emitted to be from the center (6 of 12 letters).

Within the sequence of keystrokes that make up the operant class, the first and last letters of each sequence represent anchor points which are behaviorally distinct from the others. Therefore, additional analyses were performed in order to measure preference for the middle row and center section separately for both the first and last letter keystroke of each operant sequence.

The three panels on the left side of Figure 5 show the proportion of keystrokes chosen from the middle row by each participant in each session of Experiment 1, while the three panels on the right show the proportion of keystrokes chosen from the center section. The top two panels of Figure 5 show the proportion of the *total* keystrokes from each of these categories, while the two middle panels present the same values as measured when looking only at the *first* keystroke of each operant sequence, and the two bottom panels show the same preferences for the *last* keystroke in each sequence. The gray horizonal

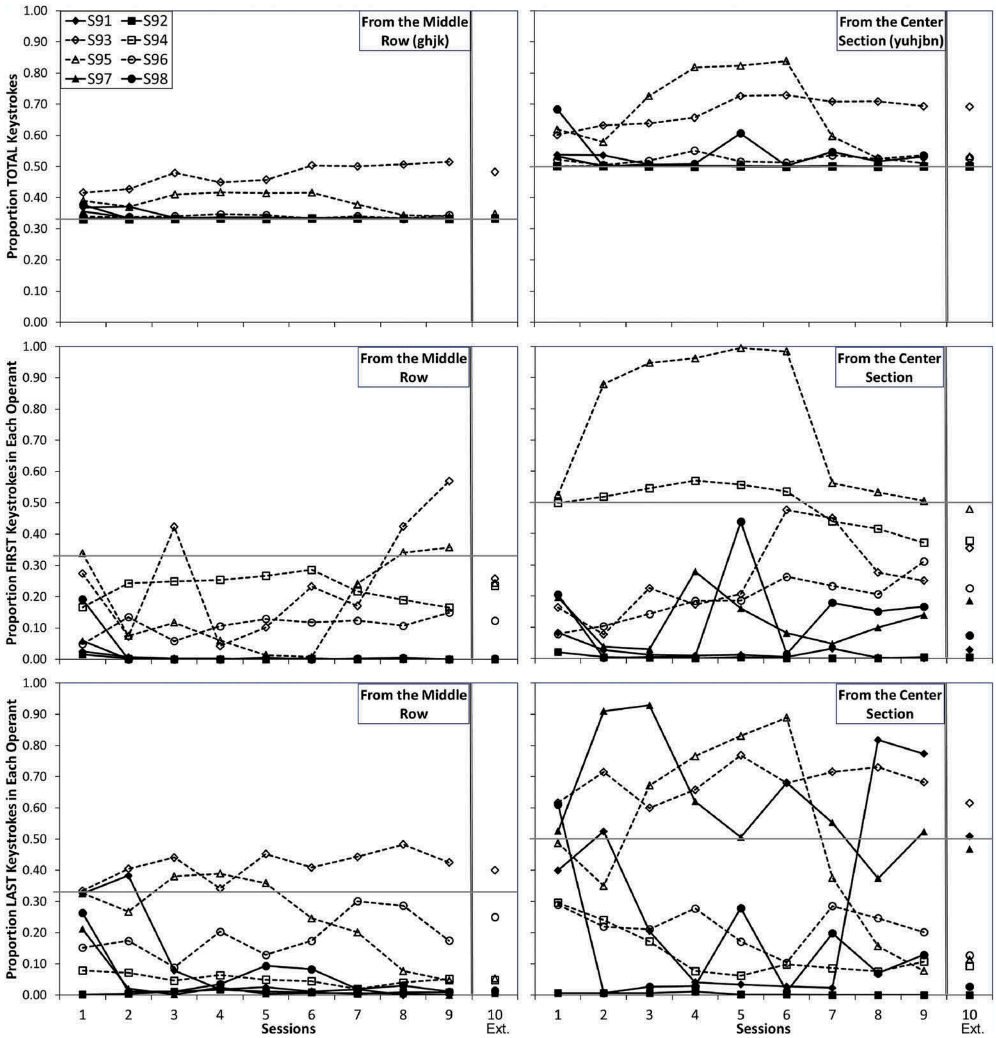

Figure 5. Experiment 1: Proportion of individual keystrokes emitted by each participant in each session from A. the middle row of three available – the three panels on the left, and B. the center section of the keyboard – the three panels on the right. The top two panels show the proportion of total keystrokes from these areas, the middle two panels show the proportion of the first keystrokes of each sequence, and the bottom two panels show the proportion of last keystrokes. The X axis of each panel is the 10 sessions of the Experiment.

line across each panel shows the indifference point of .33 for letters from the middle row (left side) and .50 for letters from the center section (right side), i.e., where the data points *would* cluster if the participants had demonstrated no significant bias regarding letter key selection. As in previous figures, the double vertical line between the last two data points represents the condition change from VR4 to extinction.

For most of the participants, the proportion of *total* keystrokes emitted in each session falls very close to those indifference points, both for keystrokes from the middle row (top left panel of Figure 5) and for those from the center section (top right panel), and is remarkably consistent from session to session. What bias toward letters from the middle and/or center does exist is seen almost exclusively in participants from the high-variability group (open data points, dashed lines). In fact, due to the extreme uniformity of the data for the low-variability participants, the difference in preference between the high-variability and low-variability groups is statistically significant for both the middle row ($F(1, 78) = 20.91$, p = .00002) and the center section ($F(1, 78) = 19.71$, p = .00003), despite the fact that not all high-variability participants are biased toward the middle/center keys. In addition, it is interesting that no participant, even those in the low-variability group, ever drops *below* the indifference point, either for letters from the middle row or center section, so there is no bias against the middle/center of the keyboard in any participant during any session – at least, when looking at total keystrokes emitted.

By contrast, the two middle panels of Figure 5 show that almost all participants generally have a strong bias *against* letters from the middle row and center section when choosing the first keystroke of each operant sequence, with this bias being strongest and most consistent among the low-variability participants, whose proportion values for these categories of letter keys are frequently zero. *All* high-variability participants are far more likely to choose letters from both the middle and center for the first keystroke of each sequence than are low-variability participants. Again, these differences between groups are statistically significant ($F(1, 78) = 84.82$, $p = 4.23\text{x}10^{-14}$ for the middle row, $F(1, 78) = 66.65$, $p = 4.53\text{x}10^{-12}$ for the center section).

The bottom left panel of Figure 5 shows that the same bias against letter keys from the middle row is also present when looking only at the last letter chosen in each operant sequence – again, with high-variability participants being far more likely to choose letters from the middle row than low-variability ones ($F(1, 78) = 40.82$, $p = 1.11\text{x}10^{-8}$). However, the proportion of last keystrokes in each sequence chosen from the center section of the keyboard – seen in the bottom right panel of Figure 5 – was much more variable across participants and across sessions, with no clear effects. In general, bias both for or against keys from the middle row was more consistent than bias involving keys from the center section of the keyboard, which varied more from session to session, as well as across participants.

In addition to measuring bias for or against letter keys, we also analyzed the transitions between keystrokes, which are part of the physical motion required to complete the operant response just as much as the keypresses themselves. Once the participant has chosen the first letter of each sequence, different movements of the hand are required depending on which of the other 11 letter keys he or she chooses next, and so on and so forth throughout the entire sequence. For purposes of analysis, the decision was made to focus on the first and most behaviorally distinctive transition, that between the first and second letters in each sequence. We measured how close those two letter keys were to each other – in other words, after choosing the initial letter, did the participant then choose a second one immediately adjacent to it on the keyboard, or farther away?

Figure 6 presents the total proportion of second letters adjacent to the first in each operant sequence, for all participants in all sessions of Experiment 1. The gray horizontal line again represents indifference, or where the proportions would be if the second letter was chosen by random chance (.33, as there are twice as many non-adjacent letters than adjacent ones available for each of the twelve keys). Four of the eight participants chose

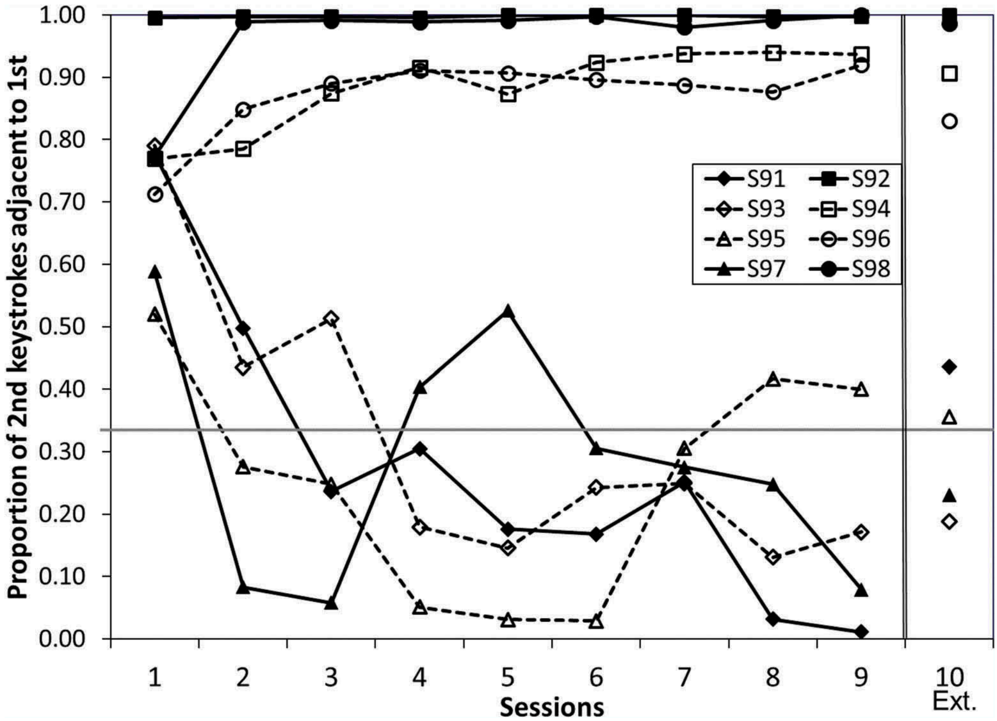

Figure 6. Experiment 1: Proportion of the total second letter keystrokes in each operant sequence that are adjacent to the first letter keystrokes emitted by each participant in each session.

adjacent letters for the second keystroke almost exclusively throughout the study, while the other four started out choosing adjacent letters more frequently than chance, then dropped approximately to or below chance levels and stayed that way for the rest of the experiment, albeit with considerable variation among participants. Interestingly, this split in the Experiment 1 participants does *not* correlate with the division between high- and low-variability groups that exists in all other results presented, both variability and bias.

Experiment 1 discussion

The contingencies put in place in Experiment 1 generated a large amount of both noncriterial variability and bias, representing experimental effects that are often not visible to the experimenter. Furthermore, the results clearly show that variability and topographic bias are linked phenomena, as participants responded to the experimental contingencies in either one of two ways: 1) with very high overall variability and a tendency to use keystrokes from the middle/center section of the keyboard more often than from the edges, or 2) with very low overall variability, coupled with a strong bias toward choosing the first and last letter keys in each operant sequence from the edges of the space available for typing. This is an interactive effect, in which certain spatial biases are seen in high-variability participants and different ones in low-variability participants. The participants' answers on post-experimental questionnaires match their differing response strategies, suggesting the possibility that self-generated "rules" for responding may have been exerting some control over their behavior.

Overall, it is interesting that operant-to-operant variability levels were so high despite being completely noncriterial, and thus never required for reinforcement. Participants could have chosen to type the exact same 12-letter sequence for every operant response during the entire experiment, but none did so. Even the low-variability group still emit anywhere between 1.5% and 12% completely unique sequences in each session, a relatively robust level of operant-to-operant variability. In addition, in this experiment we did not see a large increase in noncriterial variability when monetary presentations were discontinued, as one might have expected given the results in the literature on extinction-induced variability. Of course, it is entirely possible that this result is an artifact of the experimental design, as we have no independent measure of the reinforcing effects of the money presentations.

After analyzing the data from Experiment 1, the decision was made to run another experiment using the same operant class but with a completely different set of experimental contingencies, in order to determine how much of the results observed in Experiment 1 were due to the specific experimental design used, particularly the unpredictable nature of the reinforcement, which has been shown to affect the development of superstitious or idiosyncratic response patterns in many experimental settings. The contingencies put in place for Experiment 2 were designed to create a laboratory analog of performance learning.

Experiment 2

Method

Participants

16 new adult participants were recruited in the same manner as that described for Experiment 1, except that these were explicitly told that the first 9 sessions of the study would be practice sessions, and the final one would be the test session during which they would earn the bulk of their compensation. Of the 16 participants initially recruited for Experiment 2, 7 either failed to appear on the first day or withdrew during the course of the experiment, leaving a total of 9 who completed the entire study.

As in Experiment 1, participants received a participation fee of 10 USD per session in cash, which they kept even if they later withdrew. In Experiment 2, however, the remainder of the participants' compensation consisted of the reinforcers earned during Session 10 of the experiment, as monetary reinforcers were not given in Sessions 1 through 9. This amount was not fixed but varied depending on their performance during the test session. Participants in Experiment 2 followed the same experimental management procedures described above for Experiment 1.

Apparatus and setting

The setting and apparatus for Experiment 2 were identical to those used in Experiment 1 and described above.

Experimental design

The operant class learned by participants in Experiment 2 was identical to that used in Experiment 1. The overall experimental design, however, was deliberately made as different as possible. Instead of nine sessions of intermittent reinforcement followed by an extinction protocol, Experiment 2 was designed with nine sessions of continuous feedback followed by a tenth session featuring a linked group of related contingencies designed to stress the performance (i.e., cause the participants to make mistakes). This type of research design can be thought of as a translational one mimicking the real-world behavior of *skill learning*, or any behavior that includes a physical component such as motor learning, which is complex enough that it cannot be emitted in its correct or criterial form without prior practice. The sessions of continuous feedback are analogous to practice sessions, while the final test session is analogous to being required to perform the skill that has been learned, with consequences based on the performance. For example, a musician must learn to play a particular piece of music, practicing the required motor program many times in order to perfect it before performing the piece before an audience.

Baseline condition with procedure

The first nine sessions of Experiment 2 consisted of 400 operant responses each. As in Experiment 1, only valid responses were counted toward this total. However, in Experiment 2, *all* valid responses during Sessions 1 through 9 received continuous feedback in the form of the 3" x 3" green square described under Experiment 1 above. No monetary reinforcers were given.

The procedure for instructing participants at the beginning of the first session was identical to that used in Experiment 1 and described above, except that it did not include a demonstration of monetary reinforcement or the post-reinforcer response, as those were not used in this session. Experimental procedures for checking in, starting and ending Sessions 2 through 9 were also identical to those previously described for Experiment 1.

Test session with procedure

In Session 10 of Experiment 2, referred to as the test session, participants were subject to a group of interrelated contingencies collectively designed to disrupt the stability of the operant behavior that had been established over the last nine sessions. While it is true that this type of test session requires changes in multiple experimental variables simultaneously, thus unfortunately making it impossible to determine the contribution of each variable, these contingencies must be applied together in order to mimic the experience of skill performance in the real world. First of all, in Session 10 the green square did *not* appear to provide constant and reliable feedback. Instead, valid operant responses were reinforced continuously with a relatively large reinforcer of 80 cents each. The procedure for signaling reinforcement and the post-reinforcer response required were identical to those described above in the Method section for Experiment 1. Session 10 lasted for a total of 440 responses, whether these were valid or not. All operant attempts were counted toward this total.

In addition to positive monetary reinforcement of valid operant responses, invalid responses were punished by a loss of 40 cents each. Furthermore, a time limit was imposed during Session 10: if participants paused for too long between keystrokes at any time during the performance of each response, or between responses, they were also penalized with the loss of 40 cents. These time limits were personalized for each participant, due to their differences in working speed, and were set at five times the participant's average time between keystrokes during the previous session. Participants were thus forced to work without pausing throughout Session 10 to avoid losses. All 40-cent losses were signaled by a 110 Hz low-pitched tone emitted for .125 seconds and a message in the middle of the computer screen reading "You just lost 40 cents." That amount was then automatically deducted from the participant's earnings total in the top right corner of the screen.

Finally, an additional variability requirement was added for Session 10 of Experiment 2, this one affecting the operant response as a whole. In order to be considered valid, the 12-or-more letter sequence making up each response could not be identical to the immediately previous one (i.e., containing the exact same letter keystrokes in the same order). This Lag 1 schedule, which the participants were explicitly informed of, was imposed in order to prevent participants from reacting to the stressors of the test session with complete stereotypy, i.e. by emitting the exact same keystrokes for every operant response in the session. Similar stereotypy as a reaction to test conditions had been observed in previous, related studies (Mechner & Jones, 2011, 2015). The within-response variability requirement from Sessions 1 through 9 also remained in effect during Session 10.

At the beginning of Session 10, participants were informed of the requirements of the test session – reinforcement, punishment, post-reinforcer response, the time limits on pausing, and the Lag 1 schedule – by the experimenter. Participants then took part in a demo of the reinforcement features of the session that was identical to that undergone by Experiment 1 participants at the beginning of Session 1 (described above), before they began the final session. The total amount earned by participants during Session 10 varied widely, from a low of 220.00 USD to a high of 331.60 USD out of a total possible amount of 368.00. USD The reinforcer amount for Session 10 had been set deliberately high as the final session was designed to be stressful and produce many mistakes. However, earnings during the final session were higher than had been anticipated. The average amount earned in the final session of Experiment 2 was 299.12 USD.

Experiment 2 results

Variability

Figure 7 shows the Unique Value statistic described in the Results section of Experiment 1 above as a measure of overall variability for each participant in each session of Experiment 2. This figure shows a very wide range of variance among the nine participants, and each participant also varies from session to session. The two distinct groups of participants seen in Experiment 1 (high-variability and low-variability) are not present in these results. Overall, however, there was a decrease in Unique Value across sessions, except for one anomalous participant (S104). There is also a noticeable increase in the proportion of unique operant responses in Session 10 for 6 of the 9 participants, far larger than that observed in Session 10 of Experiment 1. Some increase in variability in Session 10 of Experiment 2 would, of course, be expected due to the imposition of the Lag 1

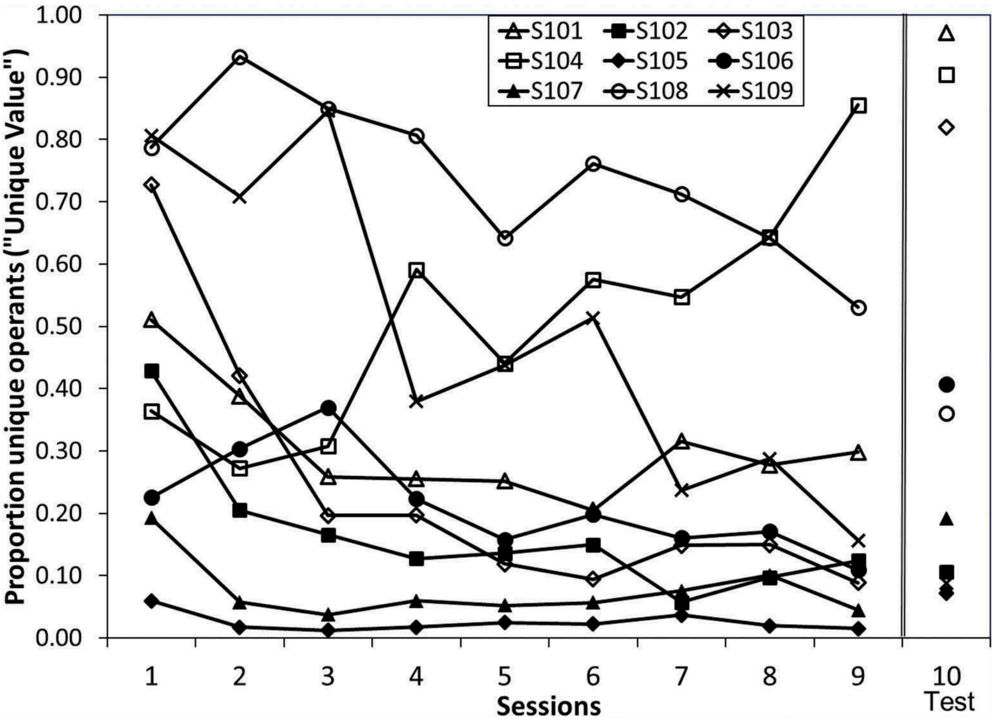

Figure 7. Experiment 2: Proportion of unique responses emitted by each participant in each session, referred to as the Unique Value.

requirement. However, the increase recorded is far beyond what would have been necessary to fulfil that minimal requirement.

Figure 8 shows an additional measurement on noncriterial variability: the proportion of responses in each session of Experiment 2 that were longer than the minimum 12 letters required, for all participants. Again, there is a large amount of variation among the participants, more so than among Experiment 1 participants when looking at the same statistic. In general, the average length of operant responses emitted in the last three sessions of Experiment 2 is greater than in the corresponding sessions of Experiment 1.

Looking at accuracy, in Experiment 2 the overall proportion of *invalid* responses (those not conforming to the criterial requirements) that were unique was always much higher than the proportion of unique operant responses in the session as a whole, just as it was in Experiment 1. The average difference in Unique Value between total and invalid operant responses in Experiment 2 is .50; this effect is consistent, occurring in all participants. Figure 9 shows the percentage of each participant's responses that were invalid during each of the 10 sessions. Again, there is no clear trend in this data, except for a large increase in errors during Session 10 for all participants, an effect not seen in Session 10 of Experiment 1. On average, Experiment 2 participants made more errors throughout the study than Experiment 1 participants.

Finally, it is worth noting that on the debriefing questionnaire for Experiment 2, the participants' answers to the question "How did you select the keystrokes required?" were

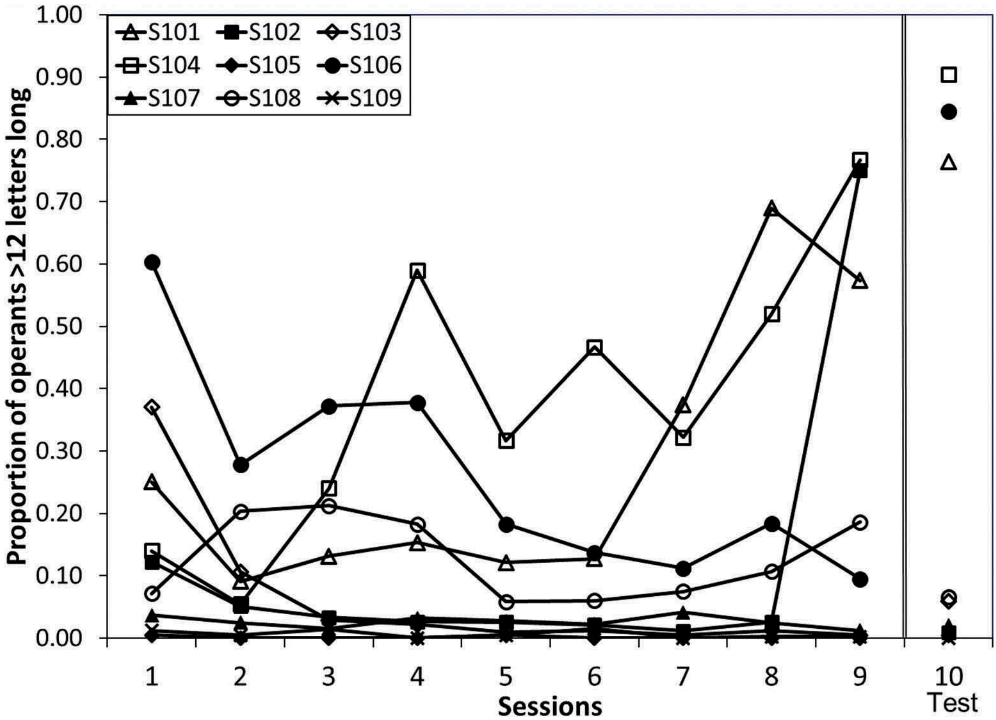

Figure 8. Experiment 2: Proportion of responses more than 12 letters long emitted by each participant in each session.

far more varied in their content than the answers given by Experiment 1 participants, and they did not necessarily accurately describe their actual behavior during the study.

Bias

As in Experiment 1, the bias measured for or against the 12 letter keys available was found to be linked to the position of the keys on the keyboard, so the results will be presented organized into the same spatial categories (middle row vs. top and bottom rows, and center section vs. left and right edges). Figure 10 shows this data for each participant in each session of Experiment 2: again, the three panels on the left side show the proportion of keystrokes chosen from the middle row, and the three panels on the right show the proportion of keystrokes chosen from the center section. As in Figure 5, presented earlier, the top two panels of Figure 10 show the proportion of total keystrokes from these categories, the two middle panels show the proportion of *first* keystrokes, the two bottom panels show the same values for the *last* keystroke in each operant sequence, and the gray horizonal line across each panel shows the indifference point.

In general, the bias data for Experiment 2 shows the same overall trends as the variability data: we see more variance than in Experiment 1, both among participants and from session to session within each participant's results. The proportion of *total* keystrokes emitted in each session of Experiment 2 still falls reasonably close to the indifference points for some participants, for keystrokes both from the middle row (top left panel of Figure 10) and those from the center section (top right panel), but the data

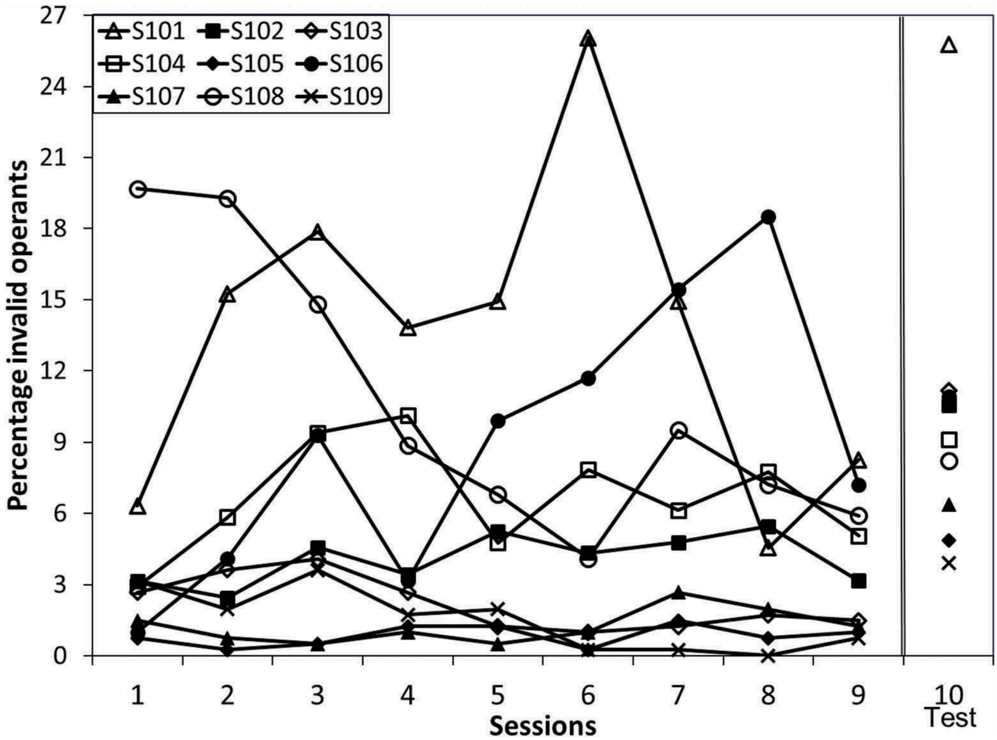

Figure 9. Experiment 2: Percentage of invalid responses emitted by each participant in each session. Please note that due to the low error rates overall the Y axis is expanded rather than scaling from 0 to 100, in order to show the differences between participants.

points spread across a much wider range of values than found in the Experiment 1 results, with clear bias toward middle/center keys for four of the nine participants. In addition, there are two anomalous participants (S101 and S106) who show a completely different bias pattern for keys from the middle row. These two start at the indifference level in Session 1, then drop nearly to zero, using almost no keystrokes from the middle row at all during sessions 3 to 9 before showing a sharp rise in use of those letter keys during the final test session. This usage pattern was not seen at all in Experiment 1, where *no* participant ever dropped below the indifference point for letters from the middle row when looking at total keystrokes.

The left middle panel of Figure 10 shows that, as in Experiment 1, almost all participants demonstrated a strong bias *against* letters from the middle row when choosing the first keystroke of each operant sequence, albeit with a few anomalous (very high) individual session values from otherwise consistent participants. Interestingly, there is a sharp increase in first keystrokes from the middle row in Session 10. In the bottom left panel we see another divided data set, with an increasing use of keys from the middle row for the *last* keystroke among approximately half of the participants, while the remainder drop almost to zero in their use of these keys for the last keystroke after the first or second session, only for this category to, again, increase sharply in Session 10.

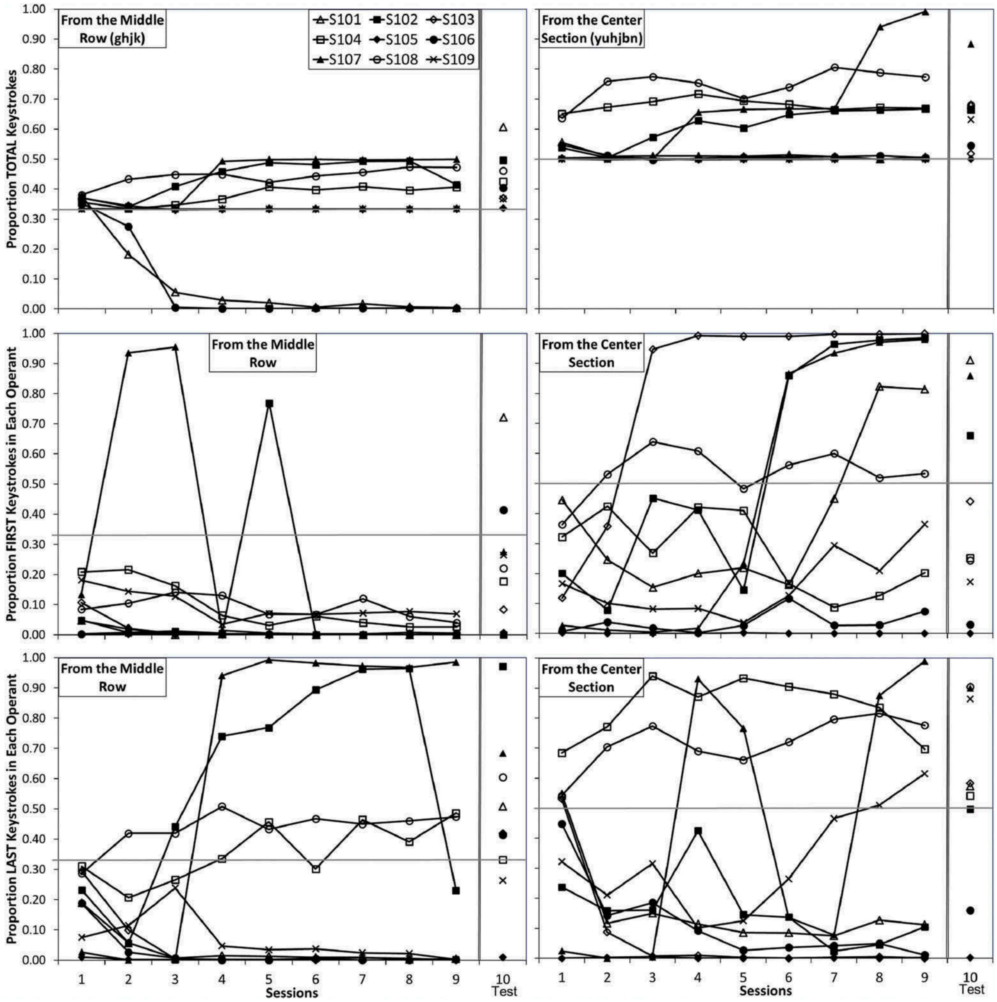

Figure 10. Experiment 2: Proportion of individual keystrokes emitted by each participant in each session from A. the middle row of three available – the three panels on the left, and B. the center section of the keyboard – the three panels on the right. The top two panels show the proportion of total keystrokes from these areas, the middle two panels show the proportion of the first keystrokes of each sequence, and the bottom two panels show the proportion of last keystrokes of each sequence.

The data on usage of keys in the center vertical section of the keyboard, when looking at the either the first and last keystrokes of each sequence (right middle and bottom panels) shows no clear trend or preference, with almost the greatest possible differences between participants, and wide swings in the values of individual participants from session to session. However, one effect that is noticeable, not only in these two panels but in the entire figure, is that imposition of the test session contingencies in Session 10 very frequently causes a sharp change in preference for these spatial groupings of letter keys. Whether the effect is to increase usage of a given category of keys or decrease it depends on the individual participant, as does the amount of the change, but overall in

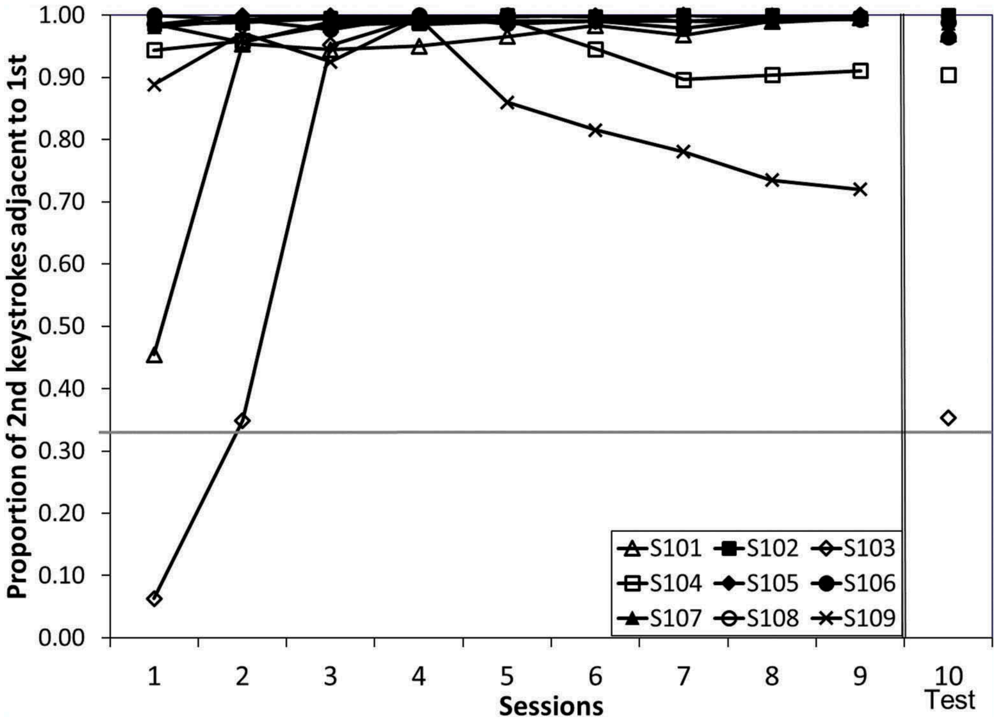

Figure 11. Experiment 2: Proportion of total second letter keystrokes in each operant sequence that are adjacent to the first letter keystrokes emitted by each participant in each session.

Experiment 2 the difference between bias measured in Session 9 and that seen in Session 10 is considerably larger than in Experiment 1.

As was done for Experiment 1, we also analyzed the transition between the first and second keystrokes for all operant sequences in Experiment 2. Figure 11 presents the total proportion of second letters adjacent to the first in each operant sequence, for all participants in all sessions of Experiment 2. The gray horizontal line again represents indifference. This statistic does not show the large differences between participants seen in the rest of the Experiment 2 data. Instead, *all* participants choose second letter keystrokes adjacent to the first almost all of the time, in sharp contrast to the split data set seen when measuring transition bias in Experiment 1.

Experiment 2 discussion

Overall, the experimental contingencies of Experiment 2 generated large amounts of noncriterial variability and significant topographic bias, as in Experiment 1, but the *specific* variability and bias effects – and the combination of the two – seen in this second study were different from those in the first. The split between high-variability and low-variability groups, which was consistent throughout almost all the Experiment 1 results, did not occur in Experiment 2, which suggests that those extreme patterns of responding may be more likely to develop under conditions of intermittent reinforcement. Throughout Experiment 2 there was more variance among participants' behavior than in Experiment 1, as well as more variation from session to session in the behavior of individual participants. Thus continuous feedback coupled with the possibility of future reinforcement seems to exert relatively weak control over individual behavior, at least compared to a more traditional experimental design.

There are, however, some noticeable trends in the results of Experiment 2 despite the large differences between participants. Once again, as in Experiment 1, we see that overall noncriterial operant-to-operant variability is high, even in those participants with the lowest Unique Values and lowest number of operant sequences longer than the minimum length. Given the general finding that continuous reinforcement results in less behavioral variability than intermittent, one might expect variability under these conditions to be lower; but in the present type of experiment, it appears that continuous feedback does *not* function similarly to continuous reinforcement, at least in this aspect. Furthermore, there was a significant change in behavior in almost all participants when the condition change was implemented between sessions 9 and 10. In Experiment 2, the imposition of the package of stressors in the final test session caused a significant increase in errors, and a correlated disruption in established behavior patterns with respect to both noncriterial variability and topographic bias. This was not the case in Experiment 1.

Finally, data regarding consecutive keystrokes shows that participants in Experiment 2 overwhelmingly chose second letter keys adjacent to the first one typed, throughout the experiment. What might be thought of as response effort – it is easier and faster to type keys that are closer together – was minimized in all participants under these conditions of continuous feedback, in a way that was not the case for the Experiment 1 participants under conditions of intermittent reinforcement on a VR4 schedule, who instead were more likely to type the letter keys making up their operant sequences based on the keys' absolute rather than relative position on the keyboard.

General discussion

Some of the experimental findings presented in this paper apply to both studies, in spite of their different experimental designs, and may thus tell us something about noncriterial variability and topographic bias in general, and about the way these two linked phenomena interact in shaping complex human behavior. Other findings from the two experiments demonstrate different results due to their different experimental designs, with possible implications for the effects of intermittent reinforcement and non-reinforcement on this type of behavior, vs. a more interrelated set of experimental contingencies mirroring real-world conditions outside the laboratory.

The most notable of the general conclusions one can draw from the two experiments is that requiring a minimal level of criterial variability *within* each operant sequence leads to surprisingly high levels of noncriterial variability *among* operant responses – far higher than one might expect from such a highly-practiced, automatized task. The performance of half of the Experiment 1 participants, and most of those from Experiment 2, did not become stereotyped at all but continued to include a very large proportion of unique sequences throughout the study. Many participants also showed considerable variation in the length of the sequences they emitted, a type of variability requiring even more response effort than the production of unique sequences of the minimum length. Even performances with the lowest noncriterial variability levels overall never became completely stereotyped but continued to contain unique operant responses in every session. Overall these findings are in line with those of Machado (1997), who found that requiring one within-sequence type of variability in pigeons (a changeover between keys within each response sequence) led to high levels of noncriterial variability among complete sequences, even though stereotypy would have been equally adaptive. His results also showed differences among individual participants similar to those seen in the present data.

Interestingly, Schwartz (1982) found the opposite – his human participants rapidly developed highly stereotyped response sequences which they repeated over and over even when noncriterial variability was allowed, and subsequently developed higher-order stereotypy when a Lag 1 criterial variability requirement was implemented. Higher-order stereotypy is the term used to refer to a repeating pattern of n + 1 responses, which often emerges after participants have been subjected to a Lag n schedule. However, his procedure was *very* different from ours, as it used a series of discriminative stimuli to signal every sub-response within the operant sequence, making it, in effect, a chain of separate operant responses. With the revealed operant methodology, in contrast, when a Lag 1 schedule was implemented in Session 10 of Experiment 2, no participant, even the least variable ones, developed higher-order stereotypy.

It is worth noting that the revealed operant, unlike other types of behavior sequences often used as operant classes in studies of variability, always has a behaviorally distinct beginning and ending sub-operant (Mechner et al., 1997): in the case of the current revealed operant, the space bar and enter key were always required as the first and last sub-operant in each 14-keystroke sequence. This feature can help slow or disrupt the automatization or fusion of long strings of consecutive operant responses into a single sequence. Automatization, mentioned earlier, is a pervasive effect seen in the development of motor skills. Through sheer repetition, separate elements of a behavioral sequence become fused into single longer routines that are usually faster and more fluent than the sum of their component parts (Glencross, 1973; Newell & Rosenbloom, 1981).

For a majority of the participants, variability levels tended to be higher during the first and possibly also second session, as participants were initially learning the task. This finding matches results from the literature showing that noncriterial variability is almost always higher at the start of acquisition – possibly due to the phenomenon of induction, or response generalization across a range of similar responses – and then decreases as the criterial response is automatized under differential reinforcement (Jarmolowicz et al., 2015). In addition, in both experiments, we found that invalid operant responses were consistently more likely to be variable ones, as their average Unique values were significantly higher than those measured for the session as a whole. In other words, in all sessions and for all participants, a much greater proportion of the invalid responses were unique than the valid ones. Noncriterial variability is thus correlated with disruptions in criterial behavior.

The second generally applicable finding, seen in both studies, is that there are large, consistent differences among individual participants in how they approach this type of task. These differences show up in the bias data as well as in measures of variability. High levels of variation among individuals should certainly not be too surprising in experiments such as these, in which participants have a virtually unlimited number of valid options for each operant response they emit. Other variability researchers have also found differences among their human participants (Eckerman & Vreeland, 1973; Peleg et al., 2017). V. L. Lee's (1996) results, in particular, show a bimodal data set similar to that found in Experiment 1 – half of her participants generated consistently high levels of noncriterial variability throughout the study, while the other half displayed more variability at the beginning of the experiment, then dropped to much lower levels and stayed low. Dracobly et al. (2017), when studying variability in children, also found a bimodal distribution of participants reflecting either low or high variability in one experiment.

What are we to make of these individual differences? It is possible that the participants in these experiments were, at least to some degree, typing different letter keys within their operant response pursuant to differing response strategies. In Machado's (1992, 1993) experiments on criterial variability, pigeons responded to a frequency-based variability requirement either with a higher-order stereotypical sequence alternating between responses, or unpredictably, even though those two strategies were approximately equivalent in terms of number of reinforcers obtained. The various self-generated rules regarding letter choice that the participants in both experiments wrote on their debriefing questionnaires are also strikingly similar to those found by Hunziker et al. (2002) when studying variability. Their participants all stated in post-experimental questionnaires a desire to "discover the rule of the game." They had all developed individual, idiosyncratic ideas about what that rule was, and none were able to fully discern the contingencies involved. The researchers hypothesized that there are factors controlling variable human behavior other than the currently prevailing contingencies, such as learning histories that resulted in the search for rules and strategies, often verbally formulated ones (Hunziker et al., 2002). These factors are not under an experimenter's control and differ from participant to participant, with resulting individual differences.

Pre-existing bias, of course, is one of the factors that are not under the experimenter's control, and bias and variability are linked phenomena, as the results of these experiments clearly show. Any attempt by human participants to vary their responses stochastically can easily be affected by pre-existing biases due to the physical characteristics of the organism, or its learning history. A study of bias thus enhances our study of behavioral variability, helping us to understand the limits placed on it by the organism's physical features and pre-existing behavior patterns. In general, humans are not very good at spontaneously emitting unpredictable behavior, although with training they can learn to increase their unpredictability – see Neuringer (1986) for one such example. Without feedback that allows them to compare their performance with a truly stochastic model, human participants tend to consistently avoid repetition and symmetry and emit too many alternating responses (Lopes & Oden, 1987). This type of responding can be adaptive under a Lag schedule even though it is not truly stochastic. Noncriterial variability in human operant responding, therefore, is *not* necessarily related to unpredictability or stochasticity, but is under the control of contingencies. However, these are not necessarily the same contingencies that control the criterial dimensions of the behavior, and may not be known to the experimenter.

The majority of participants in both experiments demonstrated a slight but consistent preference for choosing more of their *total* letter keystrokes from the center horizontal section of the keyboard. This result corresponds with those reported in the authors' two earlier articles on bias, which used very different types of operant classes. In an experiment in which human participants were required to draw lines on a computer graphics tablet, the participants showed a strong bias for the center of the three starting points from which those lines could be drawn, regardless of the programmed independent variables of the study (Jones & Mechner, 2013). This type of pre-existing bias in humans may be related to perceptual bias, in addition to effects of neural and anatomical coordinative structures. An instance of such bias is the perceptual phenomenon referred to as center bias: the general fixation of the gaze at the center of the visual field (Tseng et al., 2009). It is possible that the participants in the current studies, when looking at the keyboard in front of them, tended to focus on the letter keys in the center of their visual field. They may also have chosen more of their total letter keystrokes from the center of the keyboard because they were easier to type and required less movement of the hand, or because they avoided edge interference from the mask that covered the other keys.

Some participants in both experiments also showed a slight bias toward letters from the middle row of the keyboard. This bias was significant during test conditions in Session 10 of Experiment 2, being present in 8 of the 9 participants. Jones and Mechner (2015) found analogous results when human participants spent multiple sessions typing several different *criterial* sequences of letters (i.e., ones where the specific letters had been chosen by the experimenters rather than the participants), with only the green square for feedback. In the final session of those studies, under test conditions almost identical to those in Experiment 2, participants were biased toward those sequences that contained more letters from the middle row of the keyboard. It is possible that under conditions that constrained their performance, the participants in Experiment 2 were more likely to fall back on a learning history of touch-typing, in which typists are taught to hover their hands over the middle row of the keyboard, referred to as the home row. The fact that typing is fastest and more error-free on the home row was confirmed many years ago in industrial research on typing behavior (Gilbreth & Gilbreth, 1920).

The majority of the participants in both of the present studies also demonstrated a strong bias *against* the center/middle keys when typing the first letter of each sequence, preferring to start each operant response with a letter key somewhere on the edges of the space that was left uncovered on the keyboard. There was also a weaker and far less consistent bias toward ending the operant sequence with a keystroke chosen from the edge of the keyboard space. Both of these biases echo findings from the authors' two previous experiments, in which humans, whether drawing lines on a graphics tablet or typing sequences of letters, generally preferred to move their hands either from left to right, or from closer to farther away from themselves (Jones & Mechner, 2013, 2015). In either case both the starting and ending points of each sequence were on the edges of the space available for movement, indicating that participants were perhaps responding to the experimental constraints placed on their behavior by moving their hands back and forth across the space available to them.

The third general finding of our two studies, therefore, is that significant topographic bias regarding letter keys typed exists among adult human participants, and can be measured if one is willing to use a methodology that allows it. In particular, these biases show evidence of being kinesthetic in origin, relating in part to the spatial position of the keys on the keyboard, rather than to any verbal component of the letters themselves. They are more complex than the biases based on simple laterality (right vs. left) that have sometimes been observed in previous studies. Because the operant class used here involves 12 keys grouped roughly in a rectangle, rather than merely two or three keys arranged in a line from left to right, these experiments are some of the first reported ones that allow the measurement of more complex spatial preferences. It is entirely possible that such preferences occur in addition to, or in place of, simple lateral bias in any studies providing more than two or three keys as response options.

Finally, in half of the Experiment 1 participants and in all the Experiment 2 participants, we found that second keystrokes adjacent to the first in the sequence were strongly favored. This type of bias would seem to be a kinesthetic preference based on minimizing the motion required to complete the operant response – it's simply easier to type letters that are closer together, and the same bias was also found in our previous experiment using a criterial operant sequence: participants favored sequences containing more adjacent keys (Jones & Mechner, 2015). It is interesting that the experimental design of Experiment 2 generated this effect much more consistently than that of Experiment 1. The VR reinforcement schedule in Experiment 1 may well have caused superstitious conditioning of specific sequences that were then maintained throughout the rest of the experiment, contributing to individual differences. This would explain why only half the participants in Experiment 1 show this effect.

These experiments also allow us to actually measure and analyze the *interaction* between pre-existing topographic bias and noncriterial behavioral variability, two of the behavioral dimensions that shape the larger operant class. The evidence for this interaction is perhaps our most general finding. For both experiments the effects of bias and variability were interrelated; however, the specific nature of the combined effect was different for each. In Experiment 1 the overall preference for keys from the center of the keyboard was stronger in those participants who showed more variability. We can speculate that these were using the center of the keyboard as a sort of home base, or anchorage point from which many different letters were more easily accessible. This was true for first and last keystrokes, measured separately, as well as for total keystrokes. In Experiment 1, the accuracy of each participant's performance is also strongly correlated with their overall level of variability, perhaps due to the effect of intermittent reinforcement and possible superstitious conditioning.

In Experiment 2, this type of correlation did not occur, but there was a sharp rise in errors for all participants during the test conditions imposed in Session 10. This effect was not seen during the cessation of reinforcement in Session 10 of Experiment 1, and a corresponding shift in variability level and/or bias was recorded for a majority of the participants. The imposition of intended stressors such as high-value reinforcers, punishment contingencies, and strict time limits in Session 10 of Experiment 2 thus caused a change (whether an increase or decrease) in variability, bias and accuracy for most of the participants, whereas unsignaled suspension of reinforcement, in Experiment 1, did not. Given the well-documented phenomenon of extinction-induced variability, it is surprising that the suspension of reinforcement did not have the effect of extinction. This may also be the reason why the few studies done on extinction-induced noncriterial variability with human participants have had somewhat mixed results (Kinloch et al., 2009; Maes, 2003; Morgan & Lee, 1996). Kinloch et al. (2009), when they suspended reinforcement for their participants, were able to elicit greater variability in one noncriterial dimension of their operant class (interresponse time) but not in a different one (the size of the rectangle drawn). It is important to note, however, that the current results on reinforcement discontinuation vs. test session variability come from only two sessions of data from two studies, and may well be an artifact of the specific design used.

In all experimental studies of complex human behavior, there are always many more factors at work than the experimenter can identify, much less control. The present studies represent an attempt to delve into two of those factors, bias and variability, in more depth than is usually possible in operant behavior research, including analysis of interactive effects. The findings here have limitations, based as they are on only two experiments with a particular type of revealed operant methodology, but they suggest potential implications for future research involving choice among alternative behaviors, in particular studies of noncriterial behavioral variability. The revealed operant methodology, as used in the present experiments, allows the measurement of topographic bias and variability separately. The authors hope to see further studies of these interrelated phenomena, using this methodology or similar ones. The present studies also support the broader argument that much can be learned by looking inside individual occurrences of operant responses. The revealed operant preparation may provide the behavioral microscope that enables this type of exploration.

Acknowledgments

The authors would like to thank Dr. Ingunn Sandaker of Oslo Metropolitan University for her support.

Disclosure statement

No potential conflict of interest was reported by the authors.

Ethical standards

The research protocols described in this paper were conducted in full accordance with the 1964 Helsinki Declaration and its later amendments, and were performed with the informed consent of all participants.

ORCID

Laurilyn D. Jones http://orcid.org/0000-0001-5918-5546
Francis Mechner http://orcid.org/0000-0002-1206-7347

References

Antonitis, J. J. (1951). Response variability in the white rat during conditioning, extinction, and reconditioning. *Journal of Experimental Psychology*, *42*(4), 273–281. https://doi.org/10.1037/h0060407

Barba, L. S. (2012). Operant variability: A conceptual analysis. *The Behavior Analyst*, *35*(2), 213–227. https://doi.org/10.1007/BF03392280

Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior*, *22*(1), 231–242. https://doi.org/10.1901/jeab.1974.22-231

Boren, J. J., Moerschbaecher, J. M., & Whyte, A. A. (1978). Variability of response location on fixed-ratio and fixed-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, *30*(1), 63–67. https://doi.org/10.1901/jeab.1978.30-63

Dracobly, J. D., Dozier, C. L., Briggs, A. M., & Juanico, J. F. (2017). An analysis of procedures that affect response variability. *Journal of Applied Behavior Analysis*, *50*(3), 600–621. https://doi.org/10.1002/jaba.392

Eckerman, D. A., & Lanson, R. N. (1969). Variability of response location for pigeons responding under continuous reinforcement, intermittent reinforcement, and extinction. *Journal of the Experimental Analysis of Behavior*, *12*(1), 73–80. https://doi.org/10.1901/jeab.1969.12-73

Eckerman, D. A., & Vreeland, R. (1973). Response variability for humans receiving continuous, intermittent, or no positive experimenter feedback. *Bulletin of the Psychonomic Society*, *2*(5), 297–299. https://doi.org/10.3758/BF03329279

Gates, J. J., & Fixen, D. L. (1968). Response variability in retarded children. *Journal of Experimental Child Psychology*, *6*(2), 306–320. https://doi.org/10.1016/0022-0965(68)90094-5

Gilbreth, F. B., & Gilbreth, L. M. (1920). *Motion study for the handicapped*. Routledge.

Glencross, D. J. (1973). Temporal organization in a repetitive speed skill. *Ergonomics*, *16*(6), 765–776. https://doi.org/10.1080/00140137308924567

Herrnstein, R. J. (1961a). Stereotypy and intermittent reinforcement. *Science*, *133*(3470), 2067–2069. https://doi.org/10.1126/science.133.3470.2067-a

Herrnstein, R. J. (1961b). Relative and absolute strength of responses as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, *4*(3), 267–272. https://doi.org/10.1901/jeab.1961.4-267

Herrnstein, R. J. (1966). Superstition: A corollary of the principles of operant conditioning. In W. K. Honig (Ed.), *Operant behavior* (pp. 33–51). Appleton-Century-Crofts.

Hunziker, M. H. L., Lee, V. P. Q., Ferreira, C. C., da Silva, A. P., & Caramori, F. C. (2002). Behavioral variability in humans: Effects of rules and contingencies. *Psicologia: Teoria E Pesquisa*, *18*(2), 139–147. https://doi.org/10.1590/S0102-37722002000200004

Iversen, I. H. (2002). Response-initiated imaging of operant behavior using a digital camera. *Journal of the Experimental Analysis of Behavior*, *77*(3), 283–300. https://doi.org/10.1901/jeab.2002.77-283

Jarmolowicz, D. P., Hudnall, J. L., Darden, A. C., Lemley, S. M., & Sofis, M. J. (2015). On the time course of rapid and repeatable response induction. *Behavioral Processes*, *120*, 116–119. https://doi.org/10.1016/j.beproc.2015.08.010

Jones, L. D., & Mechner, F. (2013). Systematic operant bias observed in human participants during research on choice. *European Journal of Behavior Analysis*, *14*(2), 295–311. https://doi.org/10.1080/15021149.2013.11434462

Jones, L. D., & Mechner, F. (2015). Kinesthetic operant bias in keyboard-based choice behavior. *European Journal of Behavior Analysis*, *16*(2), 202–228. https://doi.org/10.1080/15021149.2015.1093796

Kinloch, J. M., Foster, T. M., & McEwan, J. S. A. (2009). Extinction-induced variability in human behavior. *The Psychological Record*, *59*(3), 347–370. https://doi.org/10.1007/BF03395669

Kong, X., McEwan, J. S., Bizo, L. A., & Foster, T. M. (2017). An analysis of U-value as a measure of variability. *The Psychological Record*, *67*(4), 581–586. https://doi.org/10.1007/s40732-017-0219-2

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*(3), 429–455. https://doi.org/10.1007/BF03395586

Lee, V. L. (1996). Superstitious location changes by human beings. *The Psychological Record*, *46*(1), 71–86. https://doi.org/10.1007/BF03395164

Lopes, L. L., & Oden, G. C. (1987). Distinguishing between random and nonrandom events. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*(3), 392–400. https://doi.org/10.1037/0278-7393.13.3.392

Machado, A. (1992). Behavioral variability and frequency-dependent selection. *Journal of the Experimental Analysis of Behavior*, *58*(2), 241–263. https://doi.org/10.1901/jeab.1992.58-241

Machado, A. (1993). Learning variable and stereotypical sequences of responses: Some data and a new model. *Behavioral Processes*, *30*(2), 103–130. https://doi.org/10.1016/0376-6357(93)90002-9

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), 1–25. https://doi.org/10.1901/jeab.1997.68-1

Maes, J. H. R. (2003). Response stability and variability induced in humans by different feedback contingencies. *Animal Learning & Behavior*, *31*(4), 332–348. https://doi.org/10.3758/BF03195995

Margulies, S. (1961). Response duration in operant level, regular reinforcement, and extinction. *Journal of the Experimental Analysis of Behavior*, *4*(4), 41–48. https://doi.org/10.1901/jeab.1961.4-317

Mechner, F. (1994). The revealed operant: A way to study the characteristics of individual occurrences of operant responses. In S. Glenn (Ed.), *Cambridge center for behavioral studies monograph series: progress in behavioral studies, monograph #3,* (pp. 11–12). Cambridge Center for Behavioral Studies.

Mechner, F., Hyten, C., Field, D. P., & Madden, G. J. (1997). Using revealed operants to study the structure and properties of human operant behavior. *The Psychological Record*, *47*(1), 45–68. https://doi.org/10.1007/BF03395212

Mechner, F., & Jones, L. D. (2011). Effects of sequential aspects of learning history. *Mexican Journal of Behavior Analysis*, *37*(1), 109–138.

Mechner, F., & Jones, L. D. (2015). Effects of repetition frequency on operant strength and resurgence of non-criterial features of operants. *Mexican Journal of Behavior Analysis*, *41*(2), 63–83. https://doi.org/10.5514/rmac.v41.i2.63729

Morgan, D. L., & Lee, K. (1996). Extinction-induced response variability in humans. *The Psychological Record*, *46*(1), 145–159. https://doi.org/10.1007/BF03395168

Neuringer, A. (1986). Can people behave “randomly”? The role of feedback. *Journal of Experimental Psychology: General*, *115*(1), 62–75. https://doi.org/10.1037/0096-3445.115.1.62

Neuringer, A. (2002). Operant variability: Evidence, functions, and theory. *Psychonomic Bulletin & Review*, *9*(4), 672–705. https://doi.org/10.3758/BF03196324

Neuringer, A., Kornell, N., & Olufs, M. (2001). Stability and variability in extinction. *Journal of Experimental Psychology: Animal Behavior Processes*, *27*(1), 79–94. https://doi.org/10.1037/0097-7403.27.1.79

Newell, A., & Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition,* (p.38). Lawrence Erlbaum Associates.

Notterman, J. M. (1959). Force emission during bar pressing. *Journal of Experimental Psychology*, *58*(5), 341–347. https://doi.org/10.1037/h0042801

Peleg, R., Martin, N., & Holth, P. (2017). Variability and resistance to change of different response positions in response sequences emitted under continuous and lag reinforcement schedules. *European Journal of Behavior Analysis*, *18*(2), 264–275. https://doi.org/10.1080/15021149.2017.1310597

Schwartz, B. (1982). Reinforcement-induced behavioral stereotypy: How not to teach people to discover rules. *Journal of Experimental Psychology: General*, *111*(1), 23–59. https://doi.org/10.1037/0096-3445.111.1.23

Sidman, M. (1960). *Tactics of scientific research*. Authors Cooperative.

Skinner, B. F. (1959). The generic nature of the concepts of stimulus and response. *The Journal of General Psychology*, *12*(1),40–65. 1935. Reprinted in Skinner, B. F., *Cumulative record* (pp. 347-366). New York: Appleton-Century-Crofts. https://doi.10.1037/11324-024. https://doi.org/10.1080/00221309.1935.9920087

Sumpter, C. E., Foster, T. M., & Temple, W. (1995). Predicting and scaling hens’ preferences for topographically different responses. *Journal of the Experimental Analysis of Behavior*, *63*(2), 151–163. https://doi.org/10.1901/jeab.1995.63-151

Tatham, T. A., Wanchisen, B. A., & Hineline, P. N. (1993). Effects of fixed and variable ratios on human behavioral variability. *Journal of the Experimental Analysis of Behavior*, *59*(2), 349–359. https://doi.org/10.1901/jeab.1993.59-349

Tseng, P., Carmi, R., Cameron, I. M., Munoz, D. P., & Itti, L. (2009). Quantifying center bias of observers in free viewing of dynamic natural scenes. *Journal of Vision*, *9*(7), 1–16. https://doi.org/10.1167/9.7.4