



Topographic Bias and Variability within Non-Criterial Components of the Operant Class

Laurilyn D. Jones

OsloMet Avhandling 2019 nr 27

OSLO METROPOLITAN UNIVERSITY
STORBYUNIVERSITETET



Topographic Bias and Variability within Non-Criterial Components of the Operant Class

Laurilyn D. Jones

OSLOMET

Dissertation for the degree of philosophiae doctor (PhD)
Department of Behavioural Science
Faculty of Health Sciences
OsloMet – Oslo Metropolitan University

Autumn 2019

OsloMet Avhandling 2019 nr 27

ISSN 2535-471X (trykket)

ISSN 2535-5414 (online)

ISBN 978-82-8364-198-1 (trykket)

ISBN 978-82-8364-218-6 (online)

OsloMet – storbyuniversitetet

Universitetsbiblioteket

Skriftserien

St. Olavs plass 4,

0130 Oslo,

Telefon (47) 64 84 90 00

Postadresse:

Postboks 4, St. Olavs plass

0130 Oslo

Trykket hos Byråservice

Trykket på Scandia 2000 white, 80 gram på materiesider/200 gram på coveret

Acknowledgements

I would like to thank Dr. Francis Mechner, my mentor, advisor, leader and partner for twenty years of scientific discovery. Without him, the research discussed in this dissertation would never have been done. I'm honored to have been able to explore new areas of behavioral science with him.

Also, my warmest thanks go to Dr. Ingunn Sandaker, my advisor at OsloMet, for her consistent support and ongoing efforts on my behalf. I would not have been able even to start my degree with OsloMet without her, let alone complete it.

Finally I thank my parents, Ronald and Jean Jones, who believed firmly for over forty-five years that there is nothing I cannot do if I try.

Abstract

Topographic bias and behavioral variability are two fundamental characteristics of the operant class, and are present to some degree in all experimental analysis of behavior, whether measurable or not. They represent the overlapping effects of multiple extra-experimental contingencies, often drawn from the organism's biology and history, which interact with the programmed experimental contingencies. All operant classes have both criterial and noncriterial dimensions, with the criterial ones being those that must occur in order for the programmed consequences to take place, and the noncriterial ones being all other behaviors and aspects of behaviors associated with each operant occurrence. Bias and variability can affect both criterial and noncriterial aspects of the operant, but their effects are most often visible within the noncriterial ones. The seven experiments with human participants reported in this thesis used a novel methodology, the revealed operant, that permits the measurement of noncriterial data, thus allowing for an in-depth analysis of both topographic bias and noncriterial variability effects on common behaviors such as drawing and typing. Major findings include 1) an overall bias toward the center of the apparatus in use; 2) a bias toward movement of the hands across the horizontal plane in front of the torso, in particular a preference for moving the hand either right to left or near to far; 3) a bias toward efficiency of motion, even when the movements required were so small as to be almost indistinguishable in effort; 4) high levels of noncriterial variability even when stereotypy would have been equally reinforced, especially variability associated with performance errors; and 5) levels of variability correlated with certain measurements of bias, indicating that the two phenomena can interact under certain circumstances. These findings have implications for current and future behavioral research, as well as for the theory of the operant as a motor program, and theories of automaticity and behavioral variability.

Keywords: Topographic bias, bias, noncriterial variability, motor learning, kinesthetic bias, perceptual bias

Sammendrag

Topografiske tilbøyeligheter (bias) og adferdsvariabilitet er to grunnleggende kjennetegn ved den operante klassen, og de finnes til en viss grad i all eksperimentell adferdsanalyse, enten de er målbare eller ikke målbare. De representerer effekten av flere kontingenser som ikke kontrolleres eksperimentelt, ofte fra organismens biologi og historie, som interagerer med de programmerte eksperimentelle kontingensene. Alle operante klasser har kriteriebaserte dimensjoner og dimensjoner som ikke er kriteriedefinert. De kriteriebaserte dimensjonene er nødvendige for at de programmerte konsekvensene skal forekomme, mens de ikke kriteriedefinerte er all annen adferd og aspekter ved den som er assosiert med hver forekomst av operanten. Bias og variabilitet kan påvirke både kriteriebaserte og andre aspekter ved operanten, men effekten deres er oftest synlig ved de aspektene som ikke er kriteriedefinert. De syv eksperimentene med mennesker som beskrives i denne avhandlingen bruker en ny metode, den avslørte operanten, som tillater måling av ikke kriteriedefinerte data. Det muliggjør en dybdeanalyse av både topografisk bias og effekter av ikke kriteriedefinert variabilitet ved vanlige adferdsformer, som tegning og skrivning på tastatur. Hovedfunnene omfatter 1) en generell tilbøyelighet i retning av sentrum på det apparatet som er i bruk, 2) en tilbøyelighet til håndbevegelser over det horisontale plan foran kroppen, spesielt en preferanse for å bevege hånden fra venstre mot høyre eller fra nært til fjernt, 3) en tilbøyelighet til effektive bevegelser, også når bevegelsene var så små at den innsatsen som krevdes var omtrent lik, 4) høye nivåer av ikke kriteriedefinert variabilitet også når stereotypi ville blitt tilsvarende forsterket; spesielt variabilitet assosiert med feil utførelse, og 5) nivåer av variabilitet korrelert med visse målinger av bias, hvilket tyder på at de to fenomenene kan interagere under bestemte betingelser. Disse funnene har betydningsfulle implikasjoner for fremtidig adferdsforskning, for teorien om operanten som et motorisk program, og for teorier om automatisitet og adferdsmessig variabilitet.

Nøkkelord: Topografisk bias, bias, ikke kriteriedefinert variabilitet, motorisk læring, kinestetisk bias, perseptuell bias

Table of Contents

List of Tables and Figures.....	3
Introduction: Characteristics of the Operant Class as a Whole	5
Variability.....	6
Topographic bias.....	8
Theoretical Framework for the Experimental Work in the Dissertation	11
All operants are sequences of behavior, and all operants are choices	11
All operants occur in biological and historical context	16
General Methodology of the Studies Conducted for the Dissertation.....	25
Purpose of the dissertation	25
The revealed operant	25
The study of learning history variables in the operant laboratory	27
Multi-phase experimental design	29
Ethical considerations	32
Summaries of Papers Included in the Dissertation	35
Paper I	35
Paper II	36
Paper III	38
General Discussion	40
The results in summary	40
Theoretical findings of the present work	45
Implications for current and future research	49
Concluding remarks	51
References	53

Papers

I. Systematic Operant Bias Observed in Human Participants During Research on Choice

Jones, L. D. & Mechner, F. (2013) *European Journal of Behavior Analysis* 14, 295-311.

DOI: <https://dx.doi.org/10.1080/15021149.2013.11434462>

II. Kinesthetic Operant Bias in Keyboard-Based Choice Behavior

Jones, L. D. & Mechner, F. (2015) *European Journal of Behavior Analysis* 16, 202-228.

DOI: <https://dx.doi.org/10.1080/15021149.2015.1093796>

III. Noncriterial Behavioral Variability and Related Operant Bias in Humans

Jones, L. D. & Mechner, F. (2020) *European Journal of Behaviour Analysis*.

DOI: <https://dx.doi.org/10.1080/15021149.2020.1745526>

List of Tables and Figures

Paper I	Figure 1.1	Screengrab from Experiment 1
	Figure 1.2	Results of Experiment 1, by number of prior repetitions
	Figure 1.3	All nine operant classes learned in Experiment 1
	Figure 1.4	Results of Experiment 1, by starting position and angle of line
	Figure 1.5	Individual data from Experiment 1
	Figure 1.6	Screengrab from Experiment 2
	Figure 1.7	Results of Experiment 2, by number of prior repetitions
	Figure 1.8	Initial lines of all nine operant classes learned in Experiment 2
	Figure 1.9	Results of Experiment 2, by angle and direction of line
	Figure 1.10	Individual data on clockwise operants from Experiment 2
	Figure 1.11	Individual data on counter-clockwise operants from Experiment 2
	Figure 1.12	Results of Experiment 3, by number of prior repetitions
	Figure 1.13	Results of Experiment 3, by angle and direction of line
	Figure 1.14	Individual data on clockwise operants from Experiment 3
	Figure 1.15	Individual data on counter-clockwise operants from Experiment 3
Paper II	Table 2.1	Repetition ratios and absolute levels required in Experiment 1
	Table 2.2	Operant classes used in Experiment 1
	Table 2.3	Complete individual choice data for Group A, Experiment 1
	Table 2.4	Complete individual choice data for Group B, Experiment 1
	Table 2.5	Symbols that replaced letters in Experiment 2 operant classes
	Table 2.6	Complete individual choice data for Experiment 2
	Table 2.7	Combined individual choice data from Experiments 1 and 2
	Figure 2.1	Picture of the keyboard used in Experiment 1
	Figure 2.2	Choice results during testing for Experiment 1
	Figure 2.3	Choice results for individual operant class pairs in Experiment 1
	Figure 2.4	Degree of bias for less-practiced operant classes in Experiment 1
	Figure 2.5	Close-up of the keyboard used in Experiment 2
	Figure 2.6	Screengrab of instructions displayed in Experiment 2
Figure 2.7	Choice results during testing for Experiment 2	
Figure 2.8	Choice results for individual operant class pairs in Experiment 2	

Paper III

- Figure 2.9 Degree of bias for less-practiced operant classes in Experiment 2
- Figure 2.10 Kinesthetic bias, related to position on keyboard
- Figure 2.11 Kinesthetic bias, related to effort required to emit operant
- Figure 2.12 Kinesthetic bias, related to angle of initial hand movement
- Figure 3.1 Picture of the keyboard used in both experiments
- Figure 3.2 Unique Value for each participant in each session of Experiment 1
- Figure 3.3 Proportion of operants in Experiment 1 longer than the minimum
- Figure 3.4 Percentage of operants in Experiment 1 that were invalid
- Figure 3.5 Proportion of keystrokes in Exp 1 from middle/center of keyboard
- Figure 3.6 Proportion of second keystrokes in Exp 1 adjacent to the first
- Figure 3.7 Unique Value for each participant in each session of Experiment 2
- Figure 3.8 Proportion of operants in Experiment 2 longer than the minimum
- Figure 3.9 Percentage of operants in Experiment 2 that were invalid
- Figure 3.10 Proportion of keystrokes in Exp 2 from middle/center of keyboard
- Figure 3.11 Proportion of second keystrokes in Exp 2 adjacent to the first

Introduction: Characteristics of the Operant Class as a Whole

The field of behavior analysis was launched by Skinner's conception of the operant as a class of behaviors that can be defined by their common effects, rather than their topography. In the earliest experimental analysis of behavior, the operant used was often deliberately kept as simple as possible (a single key peck or lever press), as the topic of study was not the operant behavior itself but its rate or pattern of occurrence under different experimental contingencies (Skinner, 1938). But behavioral researchers soon expanded their focus to investigate characteristics of the operant class itself, and have continued to do so for the last 80 years, resulting in many fundamental discoveries about the nature of behavior. This dissertation represents an attempt to continue deepening our understanding of operant behavior by looking at two related aspects of the operant class: topographic bias, or the consistent emission of one behavior over another based not on the experimental variables but on the organism's biology and prior history; and noncriterial operant variability, or the many ways in which individual responses can differ topographically while still remaining members of the same operant class, subject to the same consequences.

There are two basic terminological distinctions that must be clarified in order to properly position the experimental work presented here. First is the distinction between a) measuring the individual occurrences or instances of operant behavior, and b) measuring characteristics of the larger operant class that require the comparison of aspects of multiple operant occurrences. The second distinction is the one drawn between a) "criterial" and b) "noncriterial" aspects of operant behavior. In both cases my research is concerned primarily with the second category, namely the noncriterial dimensions of the larger operant class.

The first distinction, that between the individual response and the overall operant class, was initially made by Skinner: "An operant is a class, of which a response is an instance or member. The usage is seldom respected" (Skinner, 1969, p.131). Indeed, the term "operant" is often used to refer to both the class as a whole and to the individual instances of behavior that occur, leading to confusion. Part of this confusion is inherent, however. It derives from the nature of the operant as a somewhat flexible or "fuzzy" category, which never exhibits perfect alignment between the class of responses defined by their consequences and the collection of responses actually generated when applying those consequences. We can clarify this confusion by thinking of the operant under study in any given experiment as an overarching class or

category of behaviors in which the individual responses recorded reflect the overlapping effects of multiple contingencies. Different contingencies can operate on different levels of the operant class, sometimes combining their effects and sometimes cancelling them out (Catania, 2013, p.163). One of these levels is that of the noncriterial behaviors associated with the operant class, which leads us to the second distinction referenced above.

Of the behaviors constituting any given experimental operant, some of them can be defined as criterial, i.e., they represent the specific aspects of behavior that determine whether or not the operant is reinforced or punished under the contingencies in effect (Mechner, 1994a, p.8). All of the remaining dimensions of the behaviors involved, including alternative behaviors if emitted, can be classified as noncriterial (Herrnstein, 1966). For example, in a “simple” lever-pressing operant, the criterial dimension of the behavior is the number of degrees the lever must be depressed below its starting position in order for a response to be registered. The noncriterial dimensions of that behavior are the force and speed with which it is pressed, whether the animal chooses to depress the lever with the left or right paw (or even with its nose or body), etc. The other noncriterial behaviors also associated with that operant class might include approach behaviors, both to the lever and to the food chamber, consummatory behaviors if reinforcement is delivered, and many more.

It is the experimenter who determines which behaviors and which dimensions of behavior would be considered criterial and noncriterial, of course, and the division of the behaviors emitted in any given experimental setting into criterial and noncriterial categories therefore reflects his or her ideas about which behaviors should be measured, and about the effects of the experimental contingencies on that behavior. But we must always remember that these kinds of distinctions are imposed by the researcher, rather than derived from the behavior itself. Operant behavior is defined by its consequences, and there are multiple topographic variants of any operant behavior under study that are capable of achieving the same consequences.

Variability

Since variability is a fundamental dimension of all behavior, no two responses are ever exactly alike, whether they are emitted by the same organism or different ones. Again, I will define my terms before proceeding. There are two different types of variance which are found at least to some degree in all experimental data: variation among the participants, which we can call

“inter-participant variation,” and instance-to-instance variation within the behavior of each participant. Conscientious researchers attempt to minimize the former through methodological control of experimental variables, while the latter has become a significant and broad area of study in behavior analysis (see Lee, Sturmey & Fields [2007] or Neuringer [2002] for a review of the literature), and is what we refer to here by the term “variability” (although it also has multiple forms).

Referring to the first type of experimental variance, one of the ways in which Skinner initially explained inter-participant variation among his subjects was with the concept of “drive,” which can exist at different levels in each participant due to previous history (Skinner, 1938, pp. 362-368). Behavior analysts now would more commonly refer to this type of variation as the result of establishing or motivating operations (Michael, 1993). An establishing operation is anything that changes the subsequent reinforcing or punishing status of a stimulus (Catania, 2013, p.440). When different experimental participants have undergone differing establishing operations prior to the beginning of the experiment (such as one rat being food-deprived and another not), the parameters of the relationship between their behavior and the experimental contingencies in effect will be different for each of those participants (Sidman, 1960, p.49). The pre-experimental history of each organism is therefore one place to start analyzing such variation (Sidman, 1960, p.153). The results of the studies presented here all demonstrate inter-participant variation of this type to some degree.

Regarding the second type of variance, which we call variability (variation *within* an individual organism’s behavior), we can say that there are many types of this phenomenon. One is random variability, generated by constantly changing variables in both the environment and the organism’s physical state. By definition, this type of variability cannot be predicted, analyzed or controlled experimentally. Another is variability that has been explicitly shaped through reinforcement, often termed “variability as an operant” (Page & Neuringer, 1985). Yet a third type is due to the organism’s biologically programmed and/or learned variability; this type is evolutionarily adaptive, as it enables adaptation to changes in environmental contingencies. When it is consistent and orderly in its effects, variability can be studied as a phenomenon in its own right, a common and fundamental dimension of the operant class in use. It is this last type of variability that will be measured and analyzed in the experiments discussed here, in addition to the multiple types of topographic bias that were also recorded.

Topographic bias

The novel research methodology used in these experiments resulted in my observation of a separate but related phenomenon that exists alongside variability, one that has received only scant prior research attention, namely bias based on the organism's biology and learning history, or on physical aspects of the experimental operanda. I call this "topographic bias." Topographic bias is another important characteristic of the operant class as a whole, similar to variability, and it too falls into distinct categories, such as anatomical, sensory, verbal, history-based, mechanical, etc. The term "topographic bias" is the result of the expansion of my understanding of the phenomenon that occurred during the course of my research and analysis, and replaces the term "operant bias" that was used in the three papers included in this dissertation; the data in those papers showed me that bias is present in all operant topographies and can have multiple origins. "Operant bias" was originally chosen in order to specify that I referred not just to stimulus bias, but bias measured for or against the entire range of operant behaviors, including the full motor program required for the organism to perform the response as well as discriminate all the associated stimuli. But the term "operant bias" carries no information about the nature or source of such bias, so the more descriptive term topographic bias will be used from now on.

In my research, topographic bias is defined as a consistent preference for one operant class—or *certain noncriterial behaviors associated with an operant class*—over alternative operant classes or noncriterial behaviors that are also available to the organism, which cannot be attributed to the experimental contingencies (at least those known to the experimenter). Many of the biases demonstrated by experimental participants are idiosyncratic, and thus contribute to the first of the two types of experimental variance, variation among subjects. However, it is also possible for *all* participants in a given experiment to demonstrate the same type of biased responding, in which case the bias can be thought of as either systematic and endemic to the population of organisms under study or, alternatively, as a feature of the apparatus or other experimental conditions. It is these last types of bias that were measured in all seven experiments presented in this dissertation.

We can identify at least two different types of topographic bias that can occur in any experimental situation: kinesthetic bias, relating to the motions required to complete the operant, and perceptual bias, relating to the organism's perception of the associated stimuli. Often their effects are combined. Both of these two types of bias can have their roots in either the biology of

the organism, or its pre-experimental learning history (or again, a combination of the two).

Learned bias is particularly prevalent in the case of human participants. In humans this type of bias often (though not always) has a verbal component, and includes the human phenomenon referred to as rule-governed behavior, which will be discussed later in this dissertation.

Topographic bias as a topic of study has been addressed in the experimental psychology literature extremely infrequently. There are a few tangentially-related studies worth briefly mentioning here, however. Hansen, Jensen, Pedersen, Munksgaard, Ladewig, and Matthews (2002) looked at the effects of operant response type on the demand function for food reinforcers in female mink. The animals were willing to work harder pressing a lever than they were pulling a chain, in order to obtain the same reinforcer (and even though the force required to move the lever and chain had been carefully equalized). This suggests a systematic kinesthetic bias of mink for one operant behavior over another. Kono (2017) shaped pigeons' pecking behavior toward a large circular target under a fixed-interval schedule of reinforcement; as the fixed-interval requirement increased, the pigeons' pecking began to show systematic position biases within the target area, moving from the center toward the edges. Interestingly, their behavior also became more variable as well (perhaps due to the effects of ratio strain on the behavior). Variability and bias are often observed together in studies in which the experimental methodology allows them both to be measured, but although they can be closely associated they are separate phenomena. Topographical bias is not merely a type of variability; indeed, one can say that topographic bias is capable of decreasing variability if participants in a given experiment are systematically biased toward certain response variants over others.

In an experiment from the perceptual psychology literature, when asked to estimate the slope of a hill seen in a photograph, human participants showed differing biases depending on whether the answer was given either orally or through a motor response (Shaffer, McManama, Swank, Williams & Durgin, 2014). Verbal responses tend to overestimate the angle of slope, while responses given by means of a "palm board" on which participants rest their hand before rotating up from horizontal to match the observed slope tend to be systematically underestimated. In another perceptual psychology experiment, participants were asked to press one of two buttons to indicate the presence of a grid pattern on a background of visual noise on either the right or left side of the screen (Abrahamyan, Silva, Dakin, Carandini, & Gardner, 2016); they showed a consistent bias toward switching from one side to the other after a non-

reinforced trial. Both of these studies demonstrate how experimental contingencies and procedures can result in alternative topographies among which biases will manifest, an effect seen in my own research as well.

There has been quite a bit of experimental work on the phenomenon of “response bias” in humans—but that phrase refers in the literature to the tendency of human participants to give inaccurate answers on questionnaires or surveys, or in interviews. Since this type of research does not involve clearly defined operant behavior, that literature is not relevant here (but see Furnham [1986] for a review, if interested).

There has been some experimental work on stimulus bias—the pre-existing preference for one stimulus over another, which can form one component of topographic bias. Just to give two examples, in Pisacreta (1982), pigeons were reinforced for pecking in such a way that two to five response keys were all illuminated with the same color; the birds showed definite color preferences even though any color would produce reinforcement. Hamid (1973) showed that human subjects showed initial preference for shapes which carried high “association value” (i.e. when asked in other, previous studies to free-associate to geometric shapes, subjects were able to come up with more associations for those in this category). In these studies, the initial biases could be overcome through conditioning and reinforcement, but that is not always the case, as will be seen when looking at the biases recorded in Papers I and II of this thesis.

From an experimental methodology standpoint, there are certain standard ways of reducing bias which are commonly used in behavior analysis. In matching to sample procedures, side or position bias is normally eliminated at the beginning of acquisition by means of a correction procedure: the experimenter merely repeats the correct stimulus in the less preferred position as many times as necessary until the animal begins to choose the corresponding response reliably (Catania, 2013, pp. 157-158). In this way the researcher gradually shapes some of the initially noncriterial aspects of the operant response (the stimuli themselves) into criterial ones, while others (the spatial position of the stimuli) remain noncriterial. The biases thus are not entirely eradicated, but remain, associated with the noncriterial elements of the operant class, where they can often still be measured (as was done in the experiments in this dissertation).

In stimulus equivalence procedures with humans, stimuli are often chosen from sets of symbols with which the participants are unlikely to have prior experience, and are assumed to be free from pre-existing bias (Green & Saunders, 1998, p.237). For applied behavior analysis

studies which use stimuli to which the participants have already been exposed, the most common methodology involves conducting pre-tests to evaluate the participants' pre-existing stimulus bias, if any, before applying the experimental contingencies (Green & Saunders, 1998, p.237). And applied behavior analysis practitioners frequently conduct preference assessments with their clients in order to determine which stimuli, reinforcers and/or tasks are initially preferred above others, so that they can collect accurate data and plan appropriate interventions. (Please see Cannella, O'Reilly and Lancioni (2005) for a review of the literature on preference assessment with individuals with developmental disabilities; Wine, Reis and Hantula (2014) provide an equivalent review on preference assessment procedures in organizational behavior management).

But regardless of which methods are used to minimize the effect of bias on experimental results, it is not eliminated but still exists as one of the multiple factors that affect the larger operant class. The results of the studies included in this dissertation show that quite clearly, largely due to their measurement of the noncriterial aspects of the operant behaviors used. Variability and bias can affect both the criterial and noncriterial aspects of behavior, but their effects are often more visible in the noncriterial dimensions, and thus they are more easily measured when using an operant methodology that records noncriterial data, as the present one does.

Theoretical Framework for the Experimental Work in the Dissertation

The three papers that make up this dissertation were designed to examine both topographic bias and noncriterial variability in depth, measuring multiple types of these two characteristics of the operant class—and their interaction, if any—and analyzing that bias and variability quantitatively. Since the definition and conceptualization of the operant is relevant to this program of research, this section will therefore discuss it in some detail, along with some of the associated research from the literature.

All operants are sequences of behavior, and all operants are choices

First of all, we can say that 1) all operants, no matter how simply defined, are not single behaviors but sequences or clusters of behaviors (Mechner, 1994a; Mechner, Hyten, Field & Madden, 1997), and 2) all operants represent choices made by the organism between engaging in the behavior being measured, and engaging in alternate behaviors (Baum, 2016). Both of these

facts about the nature of the operant class contribute both to variability in the behavior stream, and to bias in the choice of those behaviors. Since alternative behaviors are always available to the organism no matter how barren the experimental environment, we can think of behavior as an allocation of the organism's limited time between competing alternatives, of which the operant behavior being studied by the experimenter may only be a very small subset (Baum, 2016). Variability is therefore created by the organism's choices among multiple behaviors, particularly noncriterial behaviors—choices that occur at every occurrence of the operant.

Even though criterial behaviors are usually the only ones recorded in the traditional experimental setting, it is valuable for the experimenter to consider the existence of the noncriterial ones in the operant sequence: "Failure to recognize that interactions are taking place between recorded and unrecorded behavior is not equivalent to eliminating such interactions from the data" (Sidman, 1960, p.179). Noncriterial behaviors often vary independently of criterial ones, for example, leading to a greater variety in the responses constituting the functional operant class. Skinner drew attention to this fact in the very first chapter of *The Behavior of Organisms*:

The present point is that when the reinforcement depends upon such a property as 'pressing a lever,' *other properties of the behavior may vary widely, although smooth curves are still obtained* (Skinner, 1938, p.38). [The emphasis is his.]

To highlight just one example of the importance of noncriterial behaviors, in situations of delayed reinforcement many experiments have confirmed that a stereotyped sequence of noncriterial behaviors can develop during the delay between emission of the criterial behavior and presentation of the reinforcer (Sidman, 1960, p.380). This is often referred to as mediating behavior, but like the overall division between criterial and noncriterial, that is a distinction made by the experimenter, not the participant, for whom the entire sequence of behaviors is reinforced as a unit.

If all operant behavior represents a choice between alternatives, then the experimental literature on behavioral choice can potentially provide insight into some of the factors affecting operant behavior in general. Since Herrnstein's (1961) formulation of the matching law, most behavioral research on choice has been based on some form of his theory. The matching law states that relative responding on two differently-reinforced alternatives matches the relative reinforcement received (Catania, 2013, p.229). Organisms most often match rather than

optimizing; optimization refers to behavior that produces the highest amount of reinforcement given the contingencies in place, which would thus lead the organism to choose the more-reinforced alternative exclusively (Catania, 2013, pp. 229-230). Davison and McCarthy (1988) provide a comprehensive history and research review of the matching law in their book of the same name.

But matching is rarely perfect. “The matching law converts separate response rates to ratios. Sometimes the ratios must be adjusted because the two responses differ in some way that biases the responding organism to one or the other” (Catania, 2013, pp. 229-230). In 1974 Baum formulated the power function version of the matching law, which is commonly known as the generalized matching law, in order to account for experimental data that deviated from strict matching. The parameter k (sometimes written as c) in the generalized matching law was designed to measure a hypothetical bias for one of the two responses relative to the other, and can thus be thought of as a measure of pre-existing or inherent preference in the organism (Jensen & Neuringer, 2009). As such, a power-based matching law equation does a good job of measuring the level of bias for one of two alternative responses in an experimental choice situation (which represents a relatively small subset of the potential effects of topographic bias on behavior). For example, the bias parameter has been shown to be consistently affected by using more- and less-preferred food reinforcers for the two responses (Sumpter, Foster & Temple, 1995). Often, however, undermatching can be eliminated entirely by taking response duration into account (Mechner, 1994a), in which case the bias represented by c can be thought of as the choice among response durations.

Normally experimenters conducting experiments on matching attempt to minimize the bias in their results as much as possible, thus there has been very little direct investigation of bias as a phenomenon in the matching literature, and none at all looking at bias in noncriterial aspects of the operant. A very few studies, however, have attempted to study matching using deliberately different types of operant behaviors, thus making the bias caused by inequivalent operanda their direct topic of interest, rather than merely a parameter in the equation used to analyze the data. Using hens as their subjects, Sumpter, Foster and Temple (1995) compared an operant consisting of key pecks to one that required the bird to push a “door” (two rods between which the birds could put their heads and push with their shoulders). The sequence of key pecks was trained first, over many sessions, then door-pushing was trained, again over several sessions,

before both operants were presented together and a changeover delay was instituted. All conditions were run until a stability criterion was met. The differing responses acted as a constant source of bias; in other words, preference for the key peck over the door push stayed relatively the same despite different reinforcement rates, and could be predicted based on the results from previous conditions (Sumpter, Foster & Temple, 1995). This study shows that topographic bias regarding the physical components of the response required can affect participants' choice data, a result also seen in the experiments that make up this dissertation.

Jensen and Neuringer (2009) reworked the generalized matching law to allow the calculation of bias for or against more than two alternative operant responses. They allowed rats to respond freely on five different possible operanda (three keys on one wall of the chamber, and two levers on the opposite, with the food tray between them), each with a different probability of reinforcement, with no changeover delay. Overall, the levers were consistently preferred to the keys; in addition there was a strong bias toward the left lever (perhaps due to the location of the water bottle) and a weak one against the center key. There were considerable individual differences among subjects, however (Jensen & Neuringer, 2009). Again, in this study we see biases that cannot be eradicated by the programmed experimental contingencies, similar to those reported in the experiments included here.

Finally, in a very interesting experiment that applied the matching law to non-laboratory behavior, Critchfield, Meeks, and Stilling (2014) compared the choice of either running or passing plays in football games to the "reinforcers" (yards gained) obtained therefrom, and found that the results could be accounted for by the matching law, with systematic biases for either passing plays or running ones depending on how many yards were needed at the time, and whether the team was losing or winning. This study is an interesting example of translational research, or the type of research which applies discoveries made in the laboratory to the real world. Translational research facilitates innovation in a science, and collaboration between basic and applied researchers is important for the future of behavior analysis as a discipline (Mace & Critchfield, 2010). The research presented in this dissertation, although it is not translational research per se, makes use of a type of experimental methodology that is designed as an analogue of real-world contingencies present in the realm of human performance and skill.

This type of real-world behavior is often seen in the study of motor learning, another area of research that both depends on and helps provide evidence for the theory of the operant as a

sequence of choices. We can think of a motor program as equivalent to Skinner's definition of the operant class based on consequences, complete with topographical variation in the individual occurrences (Mechner, 1994b, pp. 3-4). For example, writing one's name with a pen on a piece of paper (horizontally) and writing one's name with chalk on a blackboard (vertically) involve different movements of many different muscle groups, but produce the same effect: the name, written in what is recognizably the same handwriting. Skill learning, shaping, modeling, etc., thus have to do with manipulation and shaping of the operant class itself, including non-criterial characteristics thereof. There is a large literature showing that motor programs are encoded in the central nervous system, and each instance of a given motor program is unique, due to differences in parameters such as force, speed or rhythm (Mechner, 1994b, pp. 4-10). Motor programs thus represent sequences of behavior, just as operants do. "Skilled performance always involves an organized sequence of activities" (Fitts & Posner, 1967, p.1). And at each step of the sequence the organism makes a choice regarding the next one. These choices are made with the help of stimuli generated by the environment and by previous responses, normally referred to as feedback (Fitts & Posner, 1967, p.2).

All motor programs, *if* they are executed slowly enough, are affected during their emission by corrective feedback (Mechner, 1994b, pp. 5-6). Kinesthetic cues, produced by the body movements required for each step in a sequence of motions, are an especially important form of feedback in the performance of motor programs: this type of operant behavior, in which each motion provides the cue for the next, is often referred to by behavioral psychologists as a chain (Mechner, 1994b, p.20). Often what the experimenter thinks of as different operant units can become chained together in the organism's behavior, so that invalid responses as well as valid ones—i.e. unnecessary and necessary links in the chain—are being reinforced equally (Sidman, 1960, p.360).

With enough repetition, the behaviors in a chain can become linked into a single response unit through the process of automatization. Automatization is the process of learning a motor routine by which the behaviors in the sequence fuse together, becoming faster and more accurate/efficient, with shorter pauses between them (Mechner, 1994b, p.9). When that happens, kinesthetic cues, as well as other sources of feedback, become less and less necessary; this process is also known as chunking (Mechner, 1994b, pp. 20-21). Behavioral chunks are defined

as sequences of events that are too fast or fluent for each link to provide the stimulus for the next one, causing them to function as integrated units (Catania, 2013, p.129).

The concept of the automatization of a repetitive task goes back as far as William James. In 1899 Bryan and Harter, studying telegraph operators of differing degrees of skill, pinpointed the development of automaticity as key to the successful acquisition of a complex skill. (For a review of experimental work on automatization see Logan [1985].) Again, Skinner was aware of this concept as early as 1938, although he did not use that name for it: “As a rather general statement it may be said that when a reinforcement depends upon the completion of a number of similar acts, the whole group tends to acquire the status of a single response...” (Skinner, 1938, p.300).

Automatization normally results in increased stereotypy, greater speed and increased resistance to disruption, in addition to other changes in the performance (Mechner, 1994b, p.21). Its effects are commonly referenced in the motor and skill learning literature:

During the final phase of skill learning, component processes become increasingly autonomous, less directly subject to cognitive control, and less subject to interference from other ongoing activities or environmental distractions (Fitts & Posner, 1967, p.14).

Automaticity is not without a cost to the organism, however, as an overreliance on automatized motor routines can lead to more than one type of serious performance error (Toner, Montero & Moran, 2015). Automatized motor programs cannot benefit from correction from kinesthetic feedback for each sub-response in the sequence, as the speed with which they are performed causes the feedback from past responses to overlap with the performance of present ones (Fitts & Posner, 1967, p.116). And automatized motor routines lose their behavioral flexibility and become far more resistant to shaping. The countervailing benefits of automatization, on the other hand, allow an organism to perform a complex sequence of behaviors as a single operant, thereby freeing other body structures for other tasks.

All operants occur in biological and historical context

The second theory of the operant class which ties together the experimental data presented in this dissertation involves the fact that no behavior exists in isolation; it is only possible to study any operant behavior within the larger context of the life of the organism

producing that behavior. “Stimuli always occur in some kind of context both spatial and temporal. Responses to a new stimuli [sic] are always superimposed upon ongoing behavior...” (Fitts & Posner, 1967, p.83). There are three different types of context that are particularly important in the experimental analysis of complex human behavior: 1) the genetic endowment and physical capabilities of the organism, which place natural limits on behavior; 2) the individual’s behavioral history during their lifespan prior to the study, which determines the effects of current contingencies, and 3) the verbal history and conditioning that human beings accrue, which affect how they perform under experimental conditions. These three categories of pre-experimental contingencies will be discussed here in some detail in order to provide the appropriate background for the articles in this dissertation.

The organism’s genetic blueprint—the way different species are wired, so to speak—is what helps to determine which movements they find easy or hard and which stimuli they find pleasant or aversive. Marian and Keller Breland, in their famous article on the “misbehavior of organisms” (1961), pointed out that there are significant phylogenetic constraints on an animal’s behavior. For a given organism, some operants are more susceptible to shaping than others. And these constraints vary systematically from species to species: “...a species is characterized in part by the positive and negative reinforcers to which it is sensitive and the kinds of topography which are within reach...” (Skinner, 1969, p.201).

Humans, in particular, have multiple genetic limits on their physical capacities which affect their performance in many different ways; these have been studied extensively in the domains of perceptual psychology and motor learning (see Pew [1974] for one review). Some of these constraints have to do with the nature of human sensory capabilities: in order for a stimulus to play any role in operant behavior, it must be capable of being both detected *and* recognized by the experimental participant (two different things), *within the context* of the experimental environment (Fitts & Posner, 1967, p.43). In addition, the mere fact that a stimulus is detected and recognized does not determine the response it will evoke; humans have certain built-in perceptual processes, such as pattern recognition and spatial processing, that control their ability to select certain aspects of the environment over others (Fitts & Posner, 1967, p. 50). Motor programs also have constraints imposed by physics and anatomy; there are physical limitations on the accuracy and speed with which a motion can be performed, and on the precision of timing

which can be achieved (among others), and there is always a minimum delay between the occurrence of a stimulus and the initiation of a response (Fitts & Posner, 1967, pp. 72-73).

One particularly important type of genetic constraint on human operant behavior involves the existence of coordinative structures, or linkages between certain muscle groups. If we had to control each of the many hundreds of muscles in our body separately in order to perform a complex movement, we would never be able to do anything at all, which is why the human body (as in all animals) depends on these coordinated linkages between muscles and the neural structures that control them (Mechner, 1994b, pp. 31-32). These lead to natural compatibilities between certain groups of movements which are not present between others, and can operate independently of environmental contingencies (Catania, 2013, p.268). “Motor routines are generally executed by linked muscle groups that work cooperatively. The linkages are both mechanical (at the joints and tendons) and neural” (Mechner, 1994b, p.32). Neural linkages, in addition to those deriving from the organism’s genetic inheritance, can also be modified: through the process of motor learning, humans create new neural linkages and thus new and modified coordinative structures (Mechner, 1994b, pp. 32-34). In other words, organisms are predisposed to perform certain types of complex movement patterns far more smoothly and with less effort than others.

For just one example of the many highly complex coordinative structures that exist in the neuromuscular system, we can look at Loeb, Giszter, Saltiel, Bizzi, and Mussa-Ivaldi (2000). The authors measured the stability of all possible combinations involving the excitation of the 16 different leg muscles of the common bullfrog. Only a small number of these combinations (23 out of 65,536, to be exact) were able to stabilize the leg in position, thus showing that only certain groupings of muscle activations are capable of coordinated movement. Those combinations of muscle activations will thus occur more easily and naturally for the organism, and operant responses that require them will be learned more easily and naturally as well.

A particularly important subclass of coordinative linkages are those that exist between eye movements and other movements (Mechner, 1994b, p.35). These linkages are commonly grouped together under the generic term of “eye-hand coordination,” and represent the combined effects of both perceptual and motor response constraints on human behavior. (Of course, visual perception itself, as a highly complex skill, already depends on multiple coordinative structures maintained through neural linkages [Mechner, 1994b, pp. 35-36].) The importance of eye-hand

coordination shows us that genetic constraints on sensory processing and motor skills, in addition to affecting behavior singly, are also interactive: certain types of stimuli pair more naturally with certain types of responses, and vice versa, a phenomenon termed stimulus-response compatibility. This has implications for experiments that measure human performance, as stimulus-response compatibility has a significant effect both on response time and on the number and type of errors that occur (Fitts & Posner, 1967, pp. 22-24).

Since not all differences in individual performance are due to genetics, the next important type of context that must be taken into account when studying any operant is the organism's lifelong learning history. We can say that *all* experimental analysis of behavior is the study of history effects on current behavior (Tatham & Wanchisen, 1998). The imposition of an independent variable in a behavioral experiment can be viewed as the installation of a history whose effect is then studied. But in a larger sense, all organisms also enter the laboratory (or the applied behavior analysis setting) having already experienced an extensive history of learning; the "blank slate" does not exist, even in a newborn animal. The learning history of an organism can be thought of as a special type of establishing operation, sometimes termed a CEO (conditioned establishing operation), rather than an unconditioned one such as hunger or thirst (Michael, 1993).

History variables are often not investigated directly in behavior analysis studies because they are inaccessible for contemporary identification; like other aspects of the experiment which are beyond the researcher's control, such as bias and inter-participant variation (of which the participants' history is one source), they are often minimized through methodological means. Several researchers, however, such as Wanchisen (1990), have instead proposed that we should study history variables directly by including a limited subset of them as part of the experiment in order to measure their effects before the participant is exposed to the particular contingency under study, a research methodology used extensively in our laboratory (Jones & Mechner, 2007; Mechner & Jones, 2011, 2015). In this way we can compare the effects of a controlled experimental condition on a subsequent condition, manipulating the variables as needed (Tatham & Wanchisen, 1998).

The analysis of history effects on behavior can be valuable in many ways: it can provide insight into systematic differences between humans and non-humans, it has important implications for translational research and thus clinical work, and it may eventually lead to

needed changes in the theory and methodology of experimental behavior analysis (Wanchisen, 1990). Furthermore, the study of what might be called “artificial” or “experimentally installed” history effects—those created deliberately in the laboratory—can help us understand, and thus perhaps control for, the unknown effects of the participant’s learning history *before* they begin participation in the study. The experiments included in this dissertation were designed with this in mind; the specific variables used to shape the participants’ learning history in the laboratory will be discussed in the General Methodology section below.

The study of history variables in mainstream psychology dates back all the way to Ebbinghaus (1885/1913). More recently, a small amount of behavioral work has been done on the effects of laboratory-controlled history variables on operant behavior. For one example, Hirai, Okouchi, Matsumoto, and Lattal (2011), in a series of studies with human participants, found that the stimulus previously associated with a fixed ratio schedule produced higher rates of responding in a subsequent experimental component than the stimulus previously associated with a differential reinforcement of low rate (DRL) schedule, even though current contingencies for the two were the same, and the participants had not been exposed to either stimuli for some time (the experiment was four weeks long). These differing rates of responding were not stable, but gradually converged with continued exposure to the current contingencies (Hirai et al., 2011).

History effects, as observed in the laboratory, are often transient in this way, and they are often strongest when behavior itself is in transition, for example during acquisition, extinction, or when response rates have not yet reached stability after some change in the contingencies (Tatham & Wanchisen, 1998). However, we cannot assume that current contingencies will ever completely override history effects (Tatham & Wanchisen, 1998). Indeed, Hirai et al. (2011) found that the effects of the differing experimental histories of the two stimuli reemerged later in the experiment, even after having undergone an extinction process. This, of course, is an example of resurgence, the phenomenon whereby a previous behavior reemerges after a period of non-occurrence (Lattal & St. Peter Pipkin, 2009; Mechner & Jones, 2011).

Resurgence, and the effects of history variables generally, are common in motor learning. The effect of prior learning upon complex motor skills is remarkably persistent (Fitts & Posner, 1967, p.20). Even after automatization of a new motor response has taken place, any disruption or stress placed on the performance can cause resurgence of previous ones, and any variants thereof. Even invalid motor routines that have been inadvertently practiced frequently reoccur

later, especially under stressful conditions (Mechner, 1994b, pp. 39-40). In addition to resurgence of specific responses that have been previously performed, some experimental work done with punishment and avoidance responses in both monkeys and pigeons seems to show that history effects can also generalize to novel response topographies, thus making them even more likely to influence current behavior (Tatham & Wanchisen, 1998).

Prior to beginning the program of experiments that form the basis of this dissertation, Dr. Mechner and I conducted a comprehensive and multi-experiment investigation of the effects of two different basic learning history variables on human skill performance. In particular, we were interested in (A) the effect of differing numbers of repetitions of distinct operant classes on human participants' subsequent performance of those operant classes (Jones & Mechner, 2007; Mechner & Jones, 2015); and (B) the effect of sequence or order of learning on later performance of different operant classes (Mechner & Jones, 2011). In the experiments reported in these publications we also measured the resurgence of the noncriterial dimensions of the operant sequence. The generic operant class used in all of these experiments was a sequence of keystrokes typed on the computer keyboard; some of the keystrokes were criterial (chosen by the experimenters) and thus defined a specific operant class, and some were noncriterial (allowed to vary) (Jones & Mechner, 2007; Mechner & Jones 2011, 2015). This type of operant is known as a revealed operant, as it permits tracking of both the criterial and certain noncriterial attributes of each individual instance of the operant (Mechner, 1994a). The revealed operant consists of a sequence of "sub-operants" whose beginning and end are marked by unique behavioral events, in order to prevent the beginnings and ends of successive operants from fusing into a single behavioral unit (Mechner, 1994a; Mechner et al., 1997). It will be discussed in greater detail later, as it is also a central element of the general methodology used for the experiments presented in all three articles included here.

How often an organism has previously performed a given behavior (or encountered a given stimulus) can be thought of as a fundamental learning history variable, and has been studied in many different forms, both behavioral and otherwise, for well over a hundred years, dating back to Ebbinghaus's (1885/1913) self-experiments on memory. In our experiments looking at this topic, we found that operant classes which had been previously repeated more often were then more likely to be chosen over those repeated less often, under forced-choice conditions (Jones & Mechner, 2007; Mechner & Jones, 2015); this is an example of bias induced

by learning history occurring *within* the experiment, rather than prior to it. In addition, we also found distinct threshold effects within this general trend; there seems to be a minimum absolute number of prior repetitions that will cause the more-practiced operant class to be consistently preferred later in the experiment, under different experimental contingencies. Finally, we found that resurgence in the noncriterial keystrokes of the operants occurred under “stressful” conditions (i.e. when the experimental contingencies changed), with consistently higher levels of noncriterial resurgence in invalid operants than in valid ones, as well as higher levels of novel noncriterial patterns; these results indicate that many types of performance errors in skilled behavior may be instances of resurgence (Mechner & Jones, 2015). Again, some of these effects are due to history-based biases, in this case the history “installed” in the laboratory as part of the experiment.

History variables are of particular importance in the experimental analysis of human behavior, as humans (even young children), have accumulated far more extensive learning histories before entering the researcher’s laboratory than most animal subjects. Which leads us to the third type of context that is crucial to the effective analysis of human operant behavior: the relative complexity of laboratory research using human participants, who share a history of verbal and cultural conditioning relating even to the behavior of participating in an experiment.

Conducting research in the experimental analysis of behavior using human participants requires modifications to traditional experimental designs used with animals. Humans, unlike animals, have language and engage in verbal behavior almost constantly. By its very nature, their verbal behavior, both overt and covert, creates an additional contingency that is in effect at the same time as the programmed experimental ones. This property of human operant behavior is often addressed by behavior analysts by saying that humans, unlike animals, can engage in rule-governed behavior, and are therefore less sensitive to current contingencies. Rule-governed behavior, sometimes called verbally governed behavior, is defined as behavior mainly under the control of verbal antecedents, which creates different properties than those found in contingency-shaped behavior (Catania, 2013, p.327). Skinner defined rule-governed behavior as being controlled by “prior contingency-specifying stimuli” (Skinner, 1969, p.147). Those stimuli, once established, function as discriminative stimuli for future behaviors (Skinner, 1969, pp. 121-122). Rule-governed behavior in humans can function as a special type of history variable, and a potent source of learned bias.

Humans often generate rules for themselves, particularly in novel situations for which they do not already have behavioral guidelines. Self-instruction can be a form of mediating behavior, especially during acquisition of a response (Cerutti, 1989). If used in this way, self-instruction causes human experimental participants to act initially under the control of their own self-generated rules rather than the contingencies that are in place (Palmer, 2012)—in other words, it creates a source of bias. If such biased responding is maintained consistently, it may prevent the participants from *ever* encountering the actual experimental contingencies, thus making it very difficult to adapt to them (Catania, 2013, p.330).

One specific kind of rule governed behavior is particularly important in experiments with human participants: instructional control. Instructional control occurs when behavior is governed not just by the participants' general verbal history or their self-generated rules, but also by the instructions given to the participant by the experimenter. Such instructional control also involves the participant's history of social conditioning with respect to obeying similar instructions in similar situations. "Following instructions" is a complex higher-order operant class. We can think of instructions as establishing operations that then affect the acquisition of a response, or the salience of a reinforcer (Catania, 2013, p.328).

Instructional control, like most rule-governed behavior, is an example of dual or multiple control, in which behavior is simultaneously under the control of multiple contingencies. "One contingency produces the initial form of responding through instructional control. The second contingency produces collateral consequences only after compliance is generated" (Cerutti, 1989). It is important to recognize the existence of the multiple contingencies, as both contribute to the final form of the behavior (Cerutti, 1989). The study of bias can help us in this task, since if systematic bias is present it necessarily means that at least one of the multiple contingencies in effect has a different function from the programmed experimental contingencies.

For example, Galizio (1979) exposed humans to a history of multiple sessions in which the instructions given (to perform one of a set of avoidance responses in order to prevent loss of earnings for a certain amount of time) were initially accurate. When the contingency changed, instruction-following persisted *if* it did not cause the participants to come into contact with the new contingency (i.e., did not cost them anything), but was quickly extinguished when it resulted in immediate monetary loss, and did not reoccur when the loss-inducing contingency was no longer in effect (Galizio, 1979). It is possible, however, for operant behaviors generated through

instructions to become generalized response classes, which can then transfer to novel stimuli, thus extending instructional control; instructions may even establish previously neutral stimuli as conditioned reinforcers which will then work to strengthen behavior in completely different experimental conditions (Cerutti, 1989). This type of generalized instructional control can then serve as a subsequent source of biased responding.

Interestingly, instructional control of operant behavior varies depending on whether the instructions cover the criterial or noncriterial aspects of the behavior; it is often the noncriterial form, or topography, of instructed responding that is most sensitive to the molecular contingencies (Cerutti, 1989). For a demonstration of the opposite, however, Svartdal (1989, 1991) asked his participants to press a key a specific number of times. Unbeknownst to them, the reinforcers given were based on the speed of their pressing, not the number of presses in each sequence. Although they were unaware of this contingency and continued to follow the instructions given, speed of pressing (the criterial aspect of the behavior) was successfully shaped by reinforcement, and that aspect of the operant came under the control of the programmed contingency. This literature demonstrates the fact that noncriterial (in these cases, formal) properties of behavior can be converted into criterial ones—and vice versa—through differing applications of the experimental contingencies in effect, of which the instructions given are always one. Of course, formal (noncriterial) properties of behavior are not necessarily under the control of either instructions *or* current contingencies (Palmer, 2012).

When studying complex human behavior, contingencies of instructional control are simply unavoidable: even if the operant behavior under study is shaped, there must be at least some verbal interaction between experimenter and participant. But rather than being an impediment, this fact can be thought of as a strength, since it means that human laboratory research more closely resembles most human learning outside of the laboratory, in which the process of shaping is normally cut short by obtaining some form of information about the target behavior—instructions, imitation, etc. (Palmer, 2012). However, to accurately analyze laboratory data with human participants, the existence of instructional control must be recognized and, more importantly, *used* in planning the experimental design. Adult humans in modern culture, unless they have led an extremely sheltered existence, have learned how to behave in an experimental setting, thus creating a systematic bias which affects their response to the programmed contingencies. When a particular type of learning history is ubiquitous within a

culture, it is often referred to as a population stereotype, and researchers then attempt to measure and control for its biased effects on the baseline behavior of their participants (Fitts & Posner, 1967, p.21). The three articles included in this dissertation represent a comprehensive attempt to do just that.

General Methodology of the Studies Conducted for the Dissertation

Purpose of the dissertation

The initial motivation for the program of research that comprises this dissertation was my desire to systematically test the effects of certain independent variables on a wide range of different complex operant classes under experimental conditions of choice. During all of my experiments, as many aspects of each operant class as possible were measured in order to determine the potential effects of both topographic bias and noncriterial variability on their use. In other words, I wished to look “below the surface” of the traditional methods of recording and examining operant occurrences, collecting data on some of the historical and biological contingencies which are always operating on human laboratory behavior at the same time as the programmed experimental contingencies. In conducting such a systematic program of research into aspects of operant behavior that are not often subject to detailed study, I hoped specifically to fill the gap that exists in the experimental operant behavior literature with respect to bias, and to link my research on that topic with the existing bodies of research on behavioral variability, topography, motor behavior, and anatomy. This type of research requires an experimental design that incorporates several specific and often unusual methodological elements, which will be presented in this section.

The revealed operant

First of all, in order to study multiple operant classes in greater detail than traditional recording methods permit, it is necessary to use an operant consisting of behavioral sequences that are programmed and tracked by the experimenter: this is the “revealed operant” (abbreviated as rO) mentioned earlier (Mechner, 1994a, Mechner et al., 1997). As we have discussed, all operants, no matter how simply their criterial elements may have been specified, actually consist of a sequence of behaviors; the revealed operant is “revealed” because it allows some of those additional behaviors to be measured, thus opening up the interior structure of the operant

occurrence to analysis and potentially revealing both biases and variability within the participants' responding that may not be visible if the operant is defined as a single unit. The behaviors making up the revealed operant sequence can be referred to as sub-operants, and each rO begins with an initiating sub-operant and ends with a terminating sub-operant, both of which are designed to be unique actions within the flow of the behavior (Mechner, 1994a). In this way the beginning and end of each revealed operant are behaviorally marked as anchor points that are distinctive for the organism, thus preventing multiple rOs from becoming automatized and fused into one long sequence of behavior, which is always a risk in research with operants that consist of sequences of actions.

The use of different rO sequences, all performed on the same equipment, to study choice among multiple operant classes replaces the need for multiple operanda, for which the amount of effort and motor coordination required is difficult to equalize (keys, buttons, touchscreens, levers, etc.). The rO also allows the researcher to easily program both criterial and noncriterial behavioral options into the operant sequence, which, as has been discussed, is an excellent way to reveal characteristics of the operant class such as bias and variability. All of the operants used in the seven experiments included in the three articles of this thesis were revealed operants, which made both criterial and noncriterial dimensions to the behavior recordable and measurable with the experimental software.

Use of the rO also generates far more data for analysis than a simple all-or-nothing operant response, enriching the researcher's understanding of the phenomena under study. Each sub-operant in the sequence can be analyzed separately, and the relationships among sub-operants can also be quantified and measured in many different ways (such as timing, spatial biases, etc.). The results of many different types of analyses performed on the data from the experiments making up this thesis are presented in the three articles included here, revealing systematic patterns in the behavior of the participants that operate on multiple levels, from simplest (single keystrokes) to most complex (trends in variability visible across several different experimental sessions). It is this wealth of data that enabled the discovery of the topographic biases that are the primary subject of this dissertation, alongside concurrent variability effects.

The study of learning history variables in the operant laboratory

In the studies included here the decision was made to measure certain effects of learning history on behavior by generating and tracking the appropriate acquisition history within the experiment. These types of programmed history variables are one way of studying learning history directly, as discussed earlier, and can provide data which can be extrapolated to history effects in general, as well as making bias effects more apparent. There are two common ways to measure history effects within an experiment: by assigning different values of the independent variable to different groups of participants (a between-groups design, in which statistical significance testing is commonly used to determine whether there is an effect), or by assigning different values of the independent variable to each individual in the presence of different stimuli (a within-individuals, or single-subject design). The latter is preferred, as it allows one to assess differences in responding due to history variables even when there are individual differences between participants (Hirai et al., 2011). Almost all of the studies included here went one step further than that, programming multiple different values of the independent variable (in all cases, number of repetitions) for different operant classes, all of which were performed by each individual participant. These are examples of what is known as a parametric experiment, in which the independent variable is manipulated over some range of values.

Using a wide range of the independent variable with different operant classes allows us to find commonalities across subjects even if those subjects show significant individual differences (Perone, 2018). It also makes bias, if present, immediately noticeable, and allows consistent measurement and analysis of variation both among participants and within each individual's own data. In the case of parametric experiments, the term "single-subject" is really a misnomer; the entire effect is not explored within one subject's data. Rather, each subject in the group shows his or her own version of the effect, which, as mentioned, may reflect considerable individual variation. This variation does not negate the overall trends in the data, however: "If it can be shown that a given factor produces the same kind of lawfulness in the individuals of the population, then the finding has great generality, in spite of the fact that quantitative differences are still observed" (Sidman, 1960, p.51).

In this type of experiment, while individual data is always presented, data is often also averaged across participants in order to make the general trends (i.e., the shape of the function) more apparent. This is *not* the same type of statistical significance testing commonly done in

non-behavioral psychology experiments with large *n* values. The idea of determining whether a given scientific finding is valid or not through significance testing dates from Sir Ronald Fisher's work on statistics in the 1920s and 1930s, although descriptive and inferential statistics did not become the predominant research model in psychology until the 1950s, long after the development of behavioral research techniques (Blampied, 2013). One of the problems with inferential statistics is that it assumes that differences among participants are due to random variation, so-called "background noise." However, as we have seen, much of the variation among participants' data is not random but contains valuable information, particularly regarding individual biases. And looking at all the participants' data, rather than just at averages, can reveal relationships and correlations not visible otherwise. There is a well-known quote from the APA's report on statistical methods: "Before you compute any statistics, look at your data" (Wilkinson & APA Task Force, 1999). Even for effects that turn out to be statistically significant, data should be presented in such a way that the shape of the function is visible.

In addition to allowing each participant to experience a sufficiently wide range of the independent variable within the experiment, another feature necessary for the study of learning history variables in the laboratory is the use of operant behaviors with which the participants do not already have lengthy learning histories. For the seven experiments reported here, every effort was made to design rOs that would require relatively novel behaviors—at least as much as possible given that the participants used standard computer equipment as operanda. Due to practical necessity, most behavioral experiments involving human participants use common operant behaviors that can be recorded and tracked by computers—pressing keys on the keyboard, touching a touchscreen, moving the mouse, speaking into the microphone. Such technology is accessible, affordable and allows for efficient data collection via computer software. All potential human participants in such experiments already possess a vast learning history involving this type of operant behavior, merely by virtue of living in our modern, highly computer-focused society, predisposing them to biased responding.

In the case of the experiments reported here, two different types of manipulations were applied to disrupt the participants' normal behavior involving the operanda in use, changing the form of the response in order to make it less likely for them to draw on their vast learning history with related behaviors. The three studies reported in the first paper used a drawing behavior as their operant, but it was very different from the common behavior of drawing with which almost

all humans have a learning history dating from early childhood. Participants were required to draw with a stylus on a computer graphics tablet, a technology that was relatively new at the time these experiments were performed and thus not one with which our participants were familiar. The stimuli associated with the experiments appeared on the computer screen in the vertical plane before the participants, while the physical actions required for the performance of the operant were performed on the horizontal plane, with the graphics tablet placed flat on the desk in front of them. Thus the perceptual cues normally associated with drawing while looking at the paper were not present. In one of the three studies, participants were also required to draw with their nondominant hand, to further disrupt their history of practiced drawing behavior. In the studies making up papers two and three, which involved typing letters on the computer keyboard, a custom-made particleboard mask was superimposed over the keys, leaving only a small subset of letter keys uncovered and available for use (see Figure 1, Paper II for a photograph of this mask). These letters were chosen to fall between the normal position of the right and left hands when typing, thus preventing participants from using the highly practiced and automatized touch-typing motor routines which almost all adults in Western societies now possess.

Multi-phase experimental design

In order to program parametric learning history studies involving complex, novel behaviors such as these, it is necessary to have long, multi-session experiments in order to generate the quantity of data required to assess multiple operant classes over a wide range of values. All of the experiments reported in the three articles included in this dissertation consisted of multiple sessions (the exact number depended on the number of operants that the participants were asked to emit), and all had the same basic structure, which might be thought of as “acquisition followed by test.” In other words, after the operant classes in use had been emitted enough times by the participants for the performance to achieve some level of stability, a major change in the experimental contingencies was introduced in order to observe the effects of the installed learning histories on the behavior. These new contingencies were designed to place the established performance under increased “stress.” The term stress is commonly used in the motor learning literature to denote the number and level of demands made by both the task and the current environment, whether high or low (Fitts & Posner, 1967, p.33). Biases that are

suppressed by the experimental contingencies during conditions of behavioral stability can emerge when those contingencies are suddenly changed.

For all of the studies except one (presented in the third paper) positive feedback for correct operant responses was provided continuously during the acquisition phase of the experiment in the form of a large (3 by 3 inch) bright green square that flashed briefly on the computer screen after the completion of each valid revealed operant sequence. This corrective feedback helped shape the formation and automatization of the motor program required to execute the operant (discussed earlier in the section on motor learning). The so-called “test” phase of each experiment (i.e., the contingency change) then involved (again, for all experiments except one) imposing a group of related contingencies designed to be “stressful” for the participants. They included termination of the established feedback stimulus and its replacement with highly-salient reinforcement (money), punishment of invalid operants through loss of money coupled with aversive stimuli (unpleasant noises), and the imposition of time limits that required participants to work more or less continuously without pausing to rest. These contingencies mimic those that control human behavior during a real-world performance situation, in which a motor program learned through practice with feedback (such as an athletic skill, a speech, or a musical performance) must then be emitted under stressful conditions. They represent an attempt to create a laboratory analogue of common human behavior outside the laboratory and are described in far more detail in the Method section of the three papers.

The one experiment that did not use this design (included in Paper III) was programmed instead using intermittent reinforcement during the acquisition phase, followed, after the performance had stabilized over the course of several experimental sessions, by un signaled extinction as the contingency change designed to disrupt the established behavior stream. This more traditional example of behavioral laboratory design was chosen for one of the two studies in the third paper as a deliberate contrast to the other model described above, in order to pinpoint which of the variability and bias effects observed in those studies were due to the specific experimental contingencies in effect, rather than the higher-level molar contingencies controlling the participants’ general behavior under those experimental conditions.

With human participants, conducting long, multi-session experiments entails many kinds of special problems that must be solved. One of these problems is the experimenter’s lack of control over the participants’ environment during the time between sessions. In these studies

every effort was made, therefore, to maximize the participants' engagement, and to minimize the disruption caused by sessions taking place on different days, with time elapsed between them. First of all, participants were required to select a specific time slot; all of their sessions thus took place at the exact same time of day on consecutive days. In an attempt to equalize their pre-session establishing operations (and the potential bias generated by them) participants were also required to sign an agreement to keep the amount of sleep they got the night before a session, whether or not they had eaten and how much caffeine they had consumed as consistent as possible from day to day. They were asked to record these variables every day on a worksheet filled out when signing in at the beginning of each session.

Obviously merely asking participants to keep their own establishing operations constant prior to each session does not guarantee that they did so, but filling out the worksheet does require them to think about it every day, making it more likely that they will at least attempt to honor their commitment. Of course, even if each participant completely standardized their own establishing operations prior to every session, this variable would still vary from participant to participant due to their differing daily habits. Care must be taken to quantify and analyze the variance in the data as well as to identify and present the central tendencies; the two studies included in Paper III of this dissertation were designed to do exactly that.

The other problematic aspect of conducting behavioral research with humans in the laboratory is the presence of overarching contingencies of instructional control, as discussed previously in this dissertation. Unfortunately instructions (and thus instructional control) are unavoidable for an operant as complex as those used in these experiments; revealed operants, consisting as they do of multiple different behaviors, all of which must be completed in order to produce the programmed consequences, simply cannot be shaped without at least some direct instructions given to the participants. In the studies reported here, every effort was made to keep these instructions to the minimum required in order for the participants to acquire the response, in order to avoid biasing their behavior any more than necessary. They were delivered by means of a script so as to be the same for all participants, and never included any direct instruction on the specific dependent variables under study, only instruction on the general specifications of the response. Of course, even if the participants' behavior had come to be entirely under the control of the experimental contingencies rather than experimenter instructions or self-generated rules (which can never be known), the overarching group of molar contingencies controlling the

behavior of “participating in a psychology experiment” would still be in effect, as is unavoidable in all research with human participants. In the case of the current thesis, however, which concerns participant bias and variability under such conditions, it can be argued that that set of molar contingencies is crucial to the behavior under study.

Ethical considerations

The control exerted over participants by the established social contingencies that affect participation in an experiment can be considerable (the famous Milgram experiments [1963] represent only one example), and as such is rightfully the subject of serious ethical concerns regarding the participation of human volunteers in scientific research. Ever since the Helsinki Declaration (World Medical Association, 2013) was first adopted in 1964, it has been considered the gold standard for ethical treatment of human research participants, and the seven experiments that make up this dissertation were designed to comply with its guidelines. In particular, the Helsinki Declaration asserts respect for the individual as the primary principle of all research involving humans, meaning that the researcher must always protect the rights of volunteers to make informed decisions regarding their participation in the study, both before beginning, and at every point throughout. In the studies reported here, potential participants were provided with all information necessary to give fully informed consent before starting the experiment, and had the ability (of which they were informed both verbally and in writing) to withdraw from the study at any time. Furthermore, the operant behavior under study did not involve deception of any kind and did not include any sensitive or potentially upsetting content; it was also not onerous or difficult, and carried no possibility whatsoever of harm to the participants.

Participants were recruited by means of flyers posted on local university campuses, providing a number to call if interested in participating in scientific research. During the initial telephone call a preliminary screening was carried out and prospective participants were fully informed about the nature of the tasks required during the study, the number of sessions, and the compensation provided for participation. They were thus given an accurate idea of what it would be like to participate in the experiment before committing themselves, and had the opportunity to make a truly informed decision on whether or not to sign up for the study. Those who did sign up had no obligation to attend the first session; those who did arrive for the first session were given a thorough briefing on the study and asked to sign a comprehensive statement of informed

consent before they were required to do anything else. If they decided at any time during the first session that they would prefer not to participate in the research they were immediately released from all obligation to the study. If they decided at any point during the multi-day experiments that they did not wish to continue participation in the study, no action was required on their part; they merely declined to return to the laboratory for their next scheduled session. Withdrawing from these experiments was thus made as easy as possible.

Since the authors were not affiliated with a university but rather an independent research foundation at the time the research took place, course credit was not provided to compensate the volunteer participants for their time. Instead they were paid a modest fee for traveling to the experimental laboratory to participate, in addition to receiving the amount of the monetary reinforcers earned during the study. On the whole this method of compensation can be thought of as a less problematic one; allowing university students to participate in experiments for course credit poses both an ethical dilemma, as they may be consenting under coercion if they need the credit for a good grade (Brody, Cluck & Aragon, 1997), and a scientific one, as the data acquired under these conditions may not be representative (Bowen & Kensinger, 2017). However, many scientists and students of scientific ethics have also raised issues with the practice of paying participants, in particular the possibility that the amount paid may constitute “undue influence” (Latterman & Merz, 2001; Williams & Walter, 2015). If one is to pay one’s experimental participants, it is necessary to be very careful to do so in such a way that it does not affect their ability to freely and fully consent. The present experiments were designed to accomplish this goal as much as possible, in an imperfect world in which money can be a strong reinforcer.

Participants in all the studies included in this dissertation were paid a flat “participation fee” per each session attended, regardless of whether or not they completed the session. These participation fees were set at a moderate level of compensation for their time (roughly between \$5 and \$10 per hour depending on each participant’s speed at emitting the operants required): enough to be reasonable, but not so high that it would tempt someone to participate who was not otherwise willing to do so. (For comparison, the minimum legal wage in New York City at the time these experiments were conducted ranged from \$5 to \$7 per hour.)

In addition to the modest per-session fees earned by all the participants in these experiments (which did not depend on their performance but only on their presence), money was also used as the reinforcer throughout all experiments. This adds another layer of ethical

considerations, of course, which will be addressed here in detail. In addition to respecting the participant's right to informed consent, ethical research must also be effective research, with a reasonable likelihood of providing results that benefit the population studied (World Medical Association, 2013). Money can be a very strong conditioned reinforcer for adult human beings, and thus exerts far more reliable control over their behavior than something like points earned, or praise from the experimenter. It is also divisible in a way that course credit or praise are not, thus allowing the amount of any given reinforcer to be varied in fine increments. For these reasons, the use of money as a reinforcer adds immeasurably to the efficacy and thus the scientific value of the experiment.

The individual monetary reinforcers earned in these experiments were small, ranging from 8 to 80 cents per valid operant, but due to the length and complexity of the experiments the total earned in this manner by any given participant could be considerable, ranging as high as \$343. (However, this particular amount was for participation in a study that consisted of 14 sessions, each approximately an hour in length, taking place at the same time each day on 14 consecutive days including weekends, so it represented a considerable investment of time and effort on the part of those participants.) The amounts were set at this level for purely practical reasons, based on what had been learned in previous experiments regarding the amount of money required in order to recruit enough participants. There was certainly a risk, of course, that total payments at this high level might represent a kind of coercion for participants to continue their participation in the study; however in this case the judgement was made that the benefit obtained, in the sense of greater representativeness of the behavior measured using a more salient reinforcer, outweighed any such risk. Furthermore, the fact that the monetary reinforcers to be earned in future sessions were not guaranteed to the participants made it far less likely that they would continue participating in a task they found onerous.

Finally, there is definitive evidence that the payment of participants in all of these studies did *not* constitute undue influence, and that they truly felt able to withdraw their informed consent at any time: the fact that so many of them did exactly that. Participants dropped out of all of the studies reported here in droves, at many different points during the duration of the experiments. Specifically, for the three experiments in Paper I a total of 40 participants signed up but only 27 completed the studies, while the other 13 withdrew. 7 of the 26 original participants in the two experiments in Paper II withdrew during the course of the study, and 11 of

the 28 participants initially recruited for the two experiments reported in Paper III did not complete the experiments. The participants in all seven experiments included in this dissertation clearly were quite capable of making their own informed decisions regarding participation in the research despite the amount of money potentially available to them.

Summaries of Papers Included in the Dissertation

Paper I: Systematic operant bias observed in human participants during research on choice. DOI: <https://dx.doi.org/10.1080/15021149.2013.11434462>

Paper I is the result of a series of three pilot studies, originally designed to develop and test a new type of revealed operant for use in the ongoing program of research on learning history variables we conducted (mentioned earlier in this dissertation). In this sense the three experiments included in Paper I did not achieve their initial purpose. But they provided crucial data on the existence of systematic (i.e. not idiosyncratic but shared by all the participants) topographic bias affecting this type of operant behavior, data that can give us insight into the nature of such bias as a general phenomenon in human experimental participants.

The particular revealed operant chosen for these three pilot studies was a line or shape drawn on the computer graphics tablet using a stylus. The three experiments all followed the structure discussed above in the General Methodology section, each consisting of one or more sessions of acquisition (with continuous feedback in the form of the green square), followed by a final “test” session in which the interrelated group of contingencies designed to stress the performance was in effect.

In Study #1 the operant behavior consisted of drawing a line from one of three targets arranged vertically on the left side of the screen to one of three similar targets on the right side. There were thus nine possible operant classes; during the single acquisition session the 10 participants were required to emit three of those nine 100 times each, three others 200 times each, and the remaining three 400 times each. During the second session, the test session, they were allowed to choose which operant class to perform from among a rotating subset of three of the nine, while earning monetary reinforcers for each valid operant and being punished with the loss of money (coupled with a loud noise that served as an aversive stimulus) for each mistake, as well as for pausing too long at any time during the session. The nominal independent variable in this experiment was thus the number of times each operant class had been emitted during

acquisition, while the dependent variable was the number of times each was chosen by the participant during the test. The results showed that this experimental contingency was *not* controlling the participants' behavior, and that they were instead choosing which operant class to emit based on systematic biases relating apparently to both the perception of the experimental stimuli involved and the motor coordinations required to perform the behavior.

Study #2 was then designed specifically to test a modified type of operant that would hopefully not be affected by these biases. This operant consisted of drawing a triangle starting in a target in the middle of the screen, through two of the nine targets arranged in a circle around it, back to the starting point. There were thus also nine different operant classes, as in Study #1, and during the four sessions in the acquisition phase the 13 participants were required to emit three of them 100 times each, three others 300 times each and the remaining three 900 times each. During the test session, as in Study #1, they were again allowed to choose which operant class to perform from among a rotating subset of three of the nine, under the same set of "stressful" performance contingencies. Despite the methodological changes made and the increase in the ratio of repetitions forming the three levels of the nominal independent variable, again participants' choices under testing did not reflect the effects of prior number of repetitions but instead showed the presence of systematic biases, in this case again relating to the kinesthetic components of the operant behavior.

Study #3 was a replication of Study #2 in which 7 participants were required to perform the same operant with their non-dominant hand. This was a final effort to see if the biases already measured could thus be eradicated, but it too was unsuccessful, producing results almost identical to Study #2 (with a few key differences). In all three studies, the bias observed was quantifiable, and analysis showed that all of the participants' choice behavior under test conditions showed the same effects, albeit to differing degrees reflecting significant variation among them.

Paper II: Kinesthetic operant bias in keyboard-based choice behavior

DOI: <https://dx.doi.org/10.1080/15021149.2015.1093796>

During the same period when the three pilot studies whose results are reported in Paper I were being conducted, work was also continuing on our laboratory's long-term program of research designed to pinpoint the effects of learning history variables (such as number of prior repetitions, used in the pilot studies) on choice (Jones & Mechner, 2007; Mechner & Jones,

2011, 2015). As mentioned, all of those studies used an operant that consisted of typing a non-word sequence of letters on a modified computer keyboard. Although systematic bias was not measured in any of those experiments (which were designed in such a way as to minimize or avoid its effects), given the strong topographic biases measured in the three drawing experiments in Paper I the decision was made to look more closely at operants involving typing behavior, conducting an exhaustive test of possible preference for certain groups of letter keys over others in order to determine if there were biases at work here too. The two experiments included in Paper II were designed to do just that.

In Experiment #1 the operant behavior involved typing a sequence of eight specific keystrokes on the modified keyboard, with all letter keys other than the twelve letters T, Y, U, I, G, H, J, K, V, B, N and M covered by a custom-made particleboard mask (again, please see Figure 1, Paper II for a photograph). The first keystroke of each operant class was always the space bar and the last one was the enter key, thus providing a behaviorally distinct starting and ending sub-operant for each sequence. The six letters pressed in between the space bar and enter key were different for each of the operant classes used in this study. As the purpose of the study was to test as many non-word sequences of letters as possible, a total of 54 unique six-letter operant classes were learned by the 13 participants in Experiment #1, over the course of 14 sessions taking place on 14 consecutive days.

Due to the large number of operant classes used and the subsequent length of the experiment, in this study the test phase, in which participants chose which operant classes to emit, was not left until the last session. Instead, the operant classes were learned two at a time, with the test phase for each pair immediately following the acquisition phase. In this way each session of the experiment consisted of multiple acquisition and test phases, each following the general format discussed previously in the General Methodology section. Participants were always required to emit one of each pair of operant classes more than the other during acquisition, so that each test phase always represented a choice between a high-repetition and a low-repetition operant class. Although on average the results showed a very slight effect of prior repetition on choice (operants that were emitted more frequently were generally more likely to be chosen under test conditions), there was a great deal of inter-participant variation in this study, far more than in previous studies in our laboratory using a keyboard operant or in the three studies using a graphic operant that are included in Paper I. In addition, for twelve of the 27

operant class pairs (a significant minority of those tested during the study), participants were more likely to choose the low-repetition one of the pair than the high-repetition one, a result directly contradictory to that predicted by the nominal independent variable and thus evidence of systematic bias toward typing those particular groupings of letter keys over others.

Experiment #2 was designed to see if the bias toward those particular sequences was due to participants' verbal history affecting their preference for certain letters of the alphabet over others. Consequently it was a replication of Experiment #1, but with all the letters on the keys covered by stickers with simple symbols printed on them, to remove the effects of verbal cues. However, the results of the 6 participants in Experiment #2 were very similar to #1, showing not only that there was systematic topographic bias affecting typing behavior in our participants but that it was largely kinesthetic in nature, due to the motions of the hands. An in-depth analysis of the less-practiced letter groups for which bias was measured during choice showed that the bias was correlated with ergonomic aspects of keyboard behavior such as the spatial position of individual keys and the distance between the keystrokes in the sequence. Furthermore, the nature of these biases was similar, although not identical, to the systematic biases involving drawing behavior that had been measured in the three experiments included in Paper I, thus providing further insight into the fundamental topographic biases that exist in adult humans with regard to small, precise movements of the hands.

Paper III: Noncriterial behavioral variability and related operant bias in humans

DOI: <https://dx.doi.org/10.1080/15021149.2020.1745526>

The experiments reported in Papers I and II all measured systematic topographic bias present in adult human participants with respect to such common behaviors as typing and drawing. In addition, however, the data from all of those experiments also showed significant variation, both inter-participant variation and variability within each participant's data. Such high levels of variance are not surprising when one considers that the behavior in these studies was not well controlled by the experimental contingencies but instead reflects the effects of contingencies occurring previously in the participants' histories, unknown to the experimenters and thus not subject to experimental manipulation. Variability and bias, as discussed earlier in this dissertation, are associated phenomena, and often occur together in this way. The two experiments that make up Paper III were designed to allow the participants' behavior to vary

freely and thus generate a large amount of noncriterial behavioral variability to be analyzed alongside any topographical bias that might occur under the same experimental conditions.

Both studies in Paper III utilized a revealed operant consisting of 14 or more keystrokes typed on the same modified computer keyboard that was used for the experiments in Paper II, with the same 12 letter keys available for use. As with the operant sequence in those experiments, this one always began with the space bar and ended with the enter key, which served as the starting and ending sub-operants of the sequence, respectively. In order to perform a valid operant, participants were required to emit *at least* 12 letter keystrokes in between the space bar and enter key, but unlike the revealed operant sequences learned in Paper II, the specific keys were not mandated, nor was there a maximum length. Participants could thus type any combination of 12 or more letters from the letter keys available throughout the entire experiment, with certain restrictions.

The acquisition phase of the first experiment included in this paper consisted of nine sessions of intermittent monetary reinforcement on a VR4 schedule, with the reinforcer amount varying pseudo-randomly between 11, 19 or 26 cents. During the tenth session (the condition change designed to stress the established performance), unsignaled extinction was in place. Eight participants took part in 10 sessions taking place on 10 consecutive days. In the second experiment the same experimental design used in all the other studies making up this dissertation was used: the first nine sessions involved continuous feedback using the green square stimulus, while during the tenth or “test” session the “stressful” group of related contingencies was implemented. Nine participants completed the 10 sessions of this experiment.

The results of both experiments showed high levels of operant-to-operant variability and the existence of systematic topographic biases regarding choice of letter keys on the keyboard, although the size of the effects varied considerably among the different participants, and there were significant differences in results caused by the two different experimental designs. Most importantly, the data showed correlations between measurements of variability and one type of bias, which covaried for some participants, thus providing quantifiable evidence that the two phenomena can be linked. The nature of the revealed operant allowed for very wide-ranging data analyses, looking at both variability and bias in multiple different ways.

General Discussion

The results in summary

All three papers in this dissertation report experiments that permitted the investigation of bias affecting the participant's choice of operants and of noncritical elements within the operant class, and all found quantifiable evidence of multiple such biases. The experiments suggest the existence of this previously unconceptualized and rarely mentioned phenomenon that is always present to some extent in noncritical aspects of the operant (even when not visible in the critical ones), and exists alongside traditional behavioral variability, about which my research also provided significant data. To the extent that bias has a relationship to variability and interacts with it, it seems that it would tend to decrease it, as it restricts the range of topographies when it occurs. The major findings from all seven studies will now be summarized here.

First of all, data from all three articles supports the existence of a basic human topographic bias toward objects in the center of the visual field and/or the middle of the operanda in use, at least while engaged in tasks that require fine motor control. In the first study in Paper I, using the drawing operant, participants preferred to draw from the center of the three targets visible on screen. In both experiments in Paper II, participants were biased toward operants containing more keystrokes from the middle row of the keyboard (as stated earlier, in all experiments using the keyboard operant, as opposed to those using the drawing operant, no stimuli were provided on screen while participants were performing each operant, meaning that they were free to look down at their own hands on the keyboard while they typed if they wished). And in Paper III, the data on bias support an overall preference for keystrokes from both the center horizontal row of the keyboard, and the center vertical section of the letter keys available, although these preferences are weak, and not seen in all participants. Interestingly, although the additional data analysis was not included in Paper II due to issues of length, operants for which bias was observed in Experiment #2 of Paper II also contained more keystrokes from the center vertical section of the keyboard than did their high-repetition alternatives, thus aligning this data more consistently with that from Paper III.

What is the possible source of this persistent bias in adult human participants? Considering the diversity of the different operant behaviors involved, it appears to have both a perceptual and a kinesthetic component. In the first experiment in Paper I, participants were drawing on a tablet placed flat in front of them while looking at the associated stimuli on a

screen in the vertical plane in front of their eyes. The middle starting point available for each line was thus in the middle of their visual field, but required a slightly farther reach with the hand in order to draw a line beginning at that point than a line drawn from the bottom (closest) starting point. Therefore, since lines drawn from the middle target carried a slightly higher cost in terms of effort and efficiency of movement than those drawn from the bottom one, it seems reasonable to assume that the bias toward the middle target is the result of a perceptual bias in humans that leads them to preferentially notice and choose objects in the center of their visual field over those on the edges. This bias is well-known in the perceptual psychology literature, in which it is commonly referred to as “center bias” (Tatler, 2007; Tseng, Carmi, Cameron, Munoz & Itti, 2009).

In both Papers II and III, participants were free to look at the keyboard while emitting the operants in use (although they were not required to), meaning that center bias may have been affecting their choice of operants here as well. However, there is also evidence from ergonomic studies of typists that typing is easier and faster on the middle of the three rows on a standard typewriter keyboard than on either of the other two rows (Dvorak, Merrick, Dealey, & Ford, 1936; Gilbreth & Gilbreth, 1920). This indicates that there may be a kinesthetic component to this type of bias as well as a perceptual one. Topographic biases, of course, can stem from many different possible sources as we have discussed, especially the types of biases that depend in complex ways on the both the biological and learned contingencies that formed them.

All seven experiments included in the three papers also found evidence of systematic bias amongst our participants regarding the initial motions of their hands in the small space on the horizontal plane immediately in front of them. In all three studies in Paper I, participants were far *less* likely to choose operants that required initial drawing motions from right to left, or from farther away from themselves toward their own torsos. These biases did not disappear when participants in Experiment #3 were required to draw with their non-dominant hand, but were still observed. The results in Paper II also showed that participants showed bias against motions that required an initial movement (between typing the first key in the sequence and typing the second one) either from right to left, or toward the body.

In Paper III, the data on bias showed that almost all participants were more likely to choose both the first and last letter keys in each sequence from among those located at the edges of the space available on the keyboard. Since they were also slightly more likely to choose more

of their *total* keystrokes from the center of the keyboard, the middle keystrokes of each operant were therefore weighted toward the center keys, indicating that there was a general bias toward performing the action with an overall motion across the space available. Although the data was omitted from Paper III due to length, in those experiments there is also a very slight tendency (not significant, and not present in all participants) to choose the first keystroke of each operant from the bottom row and/or left vertical section of the keyboard, and the last keystroke from the top row and/or right vertical section, indicating that at least some participants were performing the operant sequences in Paper III while moving their hands in a left-to-right, near-to-far motion similar to that used for biased operants in Paper II.

This type of bias would seem to be largely kinesthetic in nature, and can possibly be attributed to the contingencies in place during each participant's lifelong learned history of writing and drawing. English speakers (which all of our participants were) write from left to right, of course, so by the time they reach adulthood they are well-practiced in moving their dominant hands in that direction while performing complex motions (Morikawa & McBeath, 1992). In addition, many children spend time throughout their childhoods drawing or painting with materials that will smear if they drag their hands across the line they have just produced, a contingency which reinforces motion of the hand away from one's own body in the horizontal plane and punishes the opposite motion. The contingencies in place when humans are learning to type favor motion from a lower row of keys toward a higher one, at least before the typist has learned to touch-type, while he or she is still looking down at the keyboard, as the hands on the keyboard hide the letter keys on the rows below the one on which one is typing. The typist is thus more likely to be accurate when moving quickly from a lower to a higher key; movement toward oneself while typing involves "aiming blind," so to speak, at a key one cannot see, and carries a higher risk of mistakes.

Finally, in five of the seven studies included in this dissertation, evidence was found for a kinesthetic bias in favor of efficiency of hand motion, i.e. toward choosing operants that involved slightly smaller movements than the alternative choices available (in Experiments 2 and 3 of Paper I, the amount of motion involved was precisely equalized for all operants learned, removing this potential source of bias for those studies). In Study #1, reported in Paper I, in addition to preferring to draw lines either directly from left to right or angled away from themselves, participants also showed a preference for drawing shorter lines (these two aspects of

the operant overlapped, making it hard to determine the source of this bias, one reason Studies #2 and 3 were designed to remove this confound). The results presented in Paper II show that participants were far more likely to choose less practiced operants containing more keystrokes immediately adjacent to each other, over more practiced operants containing more keystrokes farther apart. And in Paper III, half the participants in Experiment #1 and all the participants in Experiment #2 overwhelmingly chose second letter keys adjacent to the first letter they had just typed, almost all the time.

This may seem an obvious preference, but these types of motor programs are executed with so little physical force and involve such relatively small movements that common sense might lead one to think of them as practically interchangeable in terms of response effort. To press one key rather than another, when the keys are located only fractions of an inch apart, would seem, at first glance, to involve equivalent effort. I did not find that to be the case in my experiments. Either the effort required to emit the different operants in these experiments, although it seemed comparable, was not sufficiently equalized, or the other kinesthetic and perceptual variables, discussed earlier, were exerting control over the choice of movements.

The results of all seven experiments discussed here also show relatively high levels of variation, both inter-participant variation and variability, or variation within each participant's results. Although variance among experimental participants can obscure the trends visible in their data and thus make it harder to reach conclusions regarding the results, in the case of experiments that measure bias such variation is practically unavoidable, due to the weak control exerted over behavior by the experimental contingencies, which is what allows pre-existing bias to override them and affect the behavior strongly enough that the effect can be observed in the results. Topographic bias itself, which is present to a greater or lesser degree in different participants based on their unique learning histories, can also be a source of some of the inter-participant variation measured in our experiments, as well as affecting the variability within each participant's own behavior. In all five studies making up Papers I and II, even though the participants showed the same general effects, their behavior varied widely, both from participant to participant and within each participant's behavior.

Since variability is often more visible, and can be analyzed in more detail, when the noncriterial behaviors that are part of the operant sequence are examined, the results of the two experiments in Paper III suggest some general findings regarding noncriterial variability,

particularly noncriterial variability that occurs together with measurable bias. In those two studies, the level of variability itself varied, with a wide range observed among the participants, perhaps reflecting the differing degrees to which the specific topographic biases measured were exerting control over their behavior. In particular, participants in Experiment 1 (with the VR4 followed by extinction design) who were more variable in their behavior were also much more likely to choose the letter keys making up their noncriterial sequences from the middle of the keyboard than those who were less variable. This was true even when looking at only the first and last keystroke of each revealed operant, for which there was a general bias *against* keys from the middle of the space available. This pattern of responding suggests that high-variability participants were using the center of the keyboard as their “home base,” so to speak, while low-variability ones were more likely to move their hands repeatedly back and forth across the space available (as discussed earlier). But overall, in both experiments in Paper III operant to operant variability was relatively high throughout despite not having any effect on reinforcement. Even those participants whose data shows the lowest overall level of variability still never became completely stereotyped after 10 sessions, as one might assume would be the case with such a repetitive task, although their variability levels were higher during the first one or two sessions, while still acquiring the task, and decreased thereafter.

In addition, Paper III provided significant evidence that this type of behavioral variability is linked with errors. For all participants in both studies, invalid operant sequences in every session were found to be much more variable than valid ones. In Experiment #1, participants who were more variable (the same ones who showed evidence of a stronger bias toward the keys at the center of the keyboard) also made more errors overall than those who were less variable. Since the response required was repeated many hundreds of times and was thus subject to automatization, emitting a unique set of sub-operants in any given sequence (even if this occurred frequently) can be considered a minor disruption of the behavior’s fluency, leading to a very slightly greater probability of making a mistake. In all the experiments we conducted as part of our long-term program of research on learning history variables (discussed earlier), noncriterial variability for keystrokes chosen within invalid operants was always higher than variability measured within valid operants (Jones & Mechner, 2007; Mechner & Jones, 2011, 2015). As we have discussed, automatized or chunked sequences of motor behaviors are prone to disruption if one or more of the sub-operants change from their fluent and practiced form. The

in-depth study of non-criterial variability therefore provides us with insight into the nature of errors made while performing a repetitive operant task.

Finally, the comparison of two different experimental designs in Paper III provided important findings regarding the effect of the specific experimental contingencies on how both variability and bias are expressed in the results. Participants in Experiment 2, with an experimental design similar to that used in Papers I and II, were overall more variable than those in Experiment 1, under conditions of VR reinforcement followed by extinction. However, the imposition of the group of related “stressor” contingencies in the final (test) session of Experiment 2 caused a sharp change in participants’ behavior across the board, affecting both variability levels and measures of bias, which was not the case for extinction in Experiment 1. Participants in Experiment 2 were also more variable in the spatial biases demonstrated in their choice of letter keys, but were almost uniform in conserving effort by choosing a second keystroke immediately adjacent to the first, far more so than Experiment 1 participants.

These differences indicate that traditional behavioral laboratory contingencies, such as variable ratio reinforcement and extinction, affect the emission of behavior that shows bias as well as various types of variability in different ways than contingencies designed to mimic real-world skill learning. In particular a VR schedule of monetary reinforcers clearly exerts more control over behavior in some participants than does continuous feedback, such as that provided to a learner practicing a skill. It can suppress variability and allow for the establishment of more stereotyped response sequences. Extinction, however, which often disrupts established behavior, does not have as strong an effect as the “test” session contingencies used in our work. The imposed conditions of “extinction” may often be overridden, and their effect thereby weakened, by self-administered verbally formulated “rules,” as discussed earlier.

Theoretical findings of the present work

I believe that the findings of these seven studies fill a gap in the operant behavior literature by providing quantifiable data on the nature of multiple pre-existing topographic biases of various origins in human participants, some of which can be correlated with noncriterial variability, and the way in which those biases can affect their behavior in studies of operant learning and choice. In addition, the results from the performance learning analogue “practice followed by test” model of experimental design, especially when compared directly with those

from a more traditional behavioral design in Paper III, help us see more clearly how both bias and variability can potentially manifest in human behavior in the real world.

The complex nature of the specific operants used in my experiments, however, means that the results may not generalize to other operant responses. This is potentially true of all research on bias or variability, however: many different factors can influence response generality, and when studying these types of complex phenomena, the specific topography of the response can matter a great deal. For example, Cruvinel and de Azebedo Pires Sèrio (2008) found that variability trained on one response (lever pressing) did not generalize to another (nose poking), even though it involved the same criterial dimension of the operant, namely duration. In any experiments on human performance the results depend to some extent on how the researcher has chosen to define the behavior; no measurements of any phenomenon can ever be completely generalizable (Fitts & Posner, 1967, p.84).

In addition, due to the large amount of variation among participants found throughout all my experiments, the results are often not definitive; we cannot necessarily see the entire effect in each participant, as in the ideal of behavioral single-subject research design. Again, however, significant variation among participants is a common occurrence in human behavioral research, especially when the behavior is measured not under “steady state” conditions but, as it is here, under acquisition or when there is a significant change in the experimental contingencies, putting the performance under stress of some kind. In dealing with the variance found in my experimental results, I have followed Sidman’s (1960) advice to assume that it is orderly in some way, rather than random, and to look for functional relationships between the varying forms of my data as expressed by the different participants, who have all undergone different establishing operations prior to their participation.

All of the experiments performed as part of this dissertation provide evidence of one kind or another to support the theory that in experimental operant behavior research, the operants are under the control of multiple contingencies in addition to those programmed by the researcher. Many of these reside in the organism’s past history, and thus necessarily A) vary from organism to organism and B) are not known to the experimenter. By demonstrating that much of what controls the behavior of human experimental participants is unknown, my research forces us to the conclusion that there is simply no such thing as a “neutral” or behaviorally “equivalent” operant (or set of operants, in the case of research on choice). *All* operant responses are affected

to some degree by biases formed due to contact with contingencies prior to the experiment, and these are often just as powerful or more so than those imposed by the experimenter.

One implication of the findings presented in the three articles included here is that each occurrence of an operant is also a unique and complex motor program, learned and executed as such by the organism involved. The study of motor learning and performance can give us valuable insight into the physiological constraints on human behavior and thus the kinesthetic and perceptual biases that influence it. Furthermore, ergonomic studies of repetitive motion, in particular those required to use technological apparatus such as touchscreens, tablets, etc., can potentially provide data on the kinesthetic contingencies that have shaped such behavior before participants ever enter the behavioral laboratory.

The next notable implication of these findings is that such topographic biases, those based on perceptual and kinesthetic contingencies shaped by both genetics and past history, can potentially be far stronger than the experimentally-imposed contingencies. Not only were they strong enough to override the contingencies put in place in all experiments in Papers I and II, but in Paper II the kinesthetic biases regarding motion of the hands on the computer keyboard were as strong or stronger than any verbal biases relating to the actual letters on the keys. After analyzing the results from Experiment 1 in Paper II, we originally assumed that the bias observed could be eradicated by removing the verbal cues associated with each operant (i.e. the letters on the letter keys), but that was not the case, and more bias was observed in Experiment 2 than in Experiment 1. This has additional implications that are discussed below.

The studies making up Paper III also have implications for the broader theory of automatization, discussed previously. Increasing stereotypy of the response is often thought of as one of the criteria for assessing whether automatization has taken place, but in fact the literature is divided on this issue (Jonides, Naveh-Benjamin, & Palmer, 1985). Furthermore, there is evidence that stereotypy can develop without automaticity, in a response sequence with programmed delays (Schwartz & Reilly, 1983). In all participants in both experiments reported in Paper III, the performance of the keystroke operant sequence certainly became faster, more fluent and less prone to errors over the first few sessions of the experiment, as it presumably became automatized, but for a majority of participants it did *not* become more stereotyped, and even those whose variability levels dropped to the lowest levels still emitted dozens of completely unique operants per session, without any experimental contingency in place requiring

them to do so. It appears then, based on the results from these studies, that stereotypy of a sequence of behaviors is not a necessary condition in order for some degree of automaticity to be considered as having taken place (although it frequently does occur in conjunction with that phenomenon). Variable behavioral sequences can become automatized alongside stereotyped ones, at least when they have some stereotyped elements in common (in the case of our experiments, the space bar and enter key, which were the same in every operant).

This form of automaticity without complete stereotypy may be thought of as a type of strategy during motor learning, which maximizes later reinforcement during performance. (Please note that the term “strategy,” as used both here and in Paper III, does not imply conscious awareness or decision-making on the part of the behaving organism, whether human or animal. It merely defines a group of higher-level, interrelated behaviors that have a combined effect on the complex ways in which reinforcement can impact learning.)

Variability during practice can increase the adaptive flexibility of the performance. The inevitable unpredictable changes in the performer’s body and environment create varying circumstances to which the performance must adapt. (Mechner, 1994b, p.53)

In other words, practicing a motor program in such a way as to de-emphasize the kinesthetic cues prompting each link in the chain (i.e. in a variable manner) helps avoid later errors under conditions of performance-induced stress, in which the performance may be disrupted; this type of practice gives the performer a sense of confidence and control (Mechner, 1994b, p.54). As a general rule it is important to remember that human responses, especially those made in experimental conditions, are always potentially affected by groups of historical contingencies, some verbal in nature, that might collectively be described as strategy.

Finally, the results of Paper III, particularly the correlations between measurements of variability and measurements of certain of the spatial biases regarding the motion involved in choosing keys to press, have implications for theories regarding the sources of behavioral variability. We know from the literature on variability that humans are unable to emit truly random behavior, or perceive it accurately (Griffiths & Tenenbaum, 2003; Lopes & Oden, 1987). The high levels of variability observed in the majority of participants in the experiments in Paper III have an underlying order. However, participants in these experiments do *not* show evidence of higher-order stereotypy, i.e. the emission of stereotyped sequences of responses, which is also

often theorized as an explanation of behavioral variability. Since levels of variability, at least in one of the two studies, are correlated with degree of center bias, we can theorize that at least part of the variability observed in studies of that phenomenon reflects the overlapping influence and control of differing historical contingencies.

Implications for current and future research

The existence of measurable topographic bias for such common small hand movements as drawing lines and pressing keys, especially when it is also shown to influence experimental variability, indicates that pre-existing bias, derived from both genetic/physical constraints on behavior and our shared history with respect to the use of computer equipment, can potentially be a factor in any experiment involving human participants. Bias effects always hover in the background, whether acknowledged or not. The findings of the present experiments may therefore be relevant to the design of many other kinds of behavioral experiments.

For example, it is possible that the spatial biases found in the current research can apply to experiments that require their participants to touch icons on a touchscreen computer, which constitute a very large segment of current behavioral research. There is at least some recent evidence that they do. Hansen and Arntzen (2017) found, using eye-tracking technology, that in matching-to-sample tasks, when the sample stimulus was at the center of a screen displayed vertically in front of the participants, the comparison stimuli at the top and to the upper right of the sample were more likely to be visually scanned first than those at any other position on the screen, while participants were least likely to look first at those stimuli at the very bottom of the screen. In general, participants had a tendency to move their gaze in a rotating, circular path from the starting point at the center of the screen.

Another possibility is that the spatial biases found with respect to typing letter keys on the computer will also apply to pressing buttons in general, even those not situated on a standard keyboard. Again, there is at least one study, although not recent, which suggests that they might. Salzinger, Feldman and Portnoy (1964) reinforced human participants for either saying one of five different nonsense syllables out loud or pressing one of five buttons, arranged in a horizontal row, on which those same syllables were printed. Their intention was to compare verbal operants to nonverbal ones and see if the means of responding made any difference (which it

did—response rates were higher for nonverbal operants than verbal ones), but they also discovered significant position preferences for some keys over others.

The Salzinger, Feldman and Portnoy article is part of a long tradition of experimental work in psychology, behavioral or otherwise, using “nonsense syllables” or other non-word groupings of letters, as stimuli and/or responses. Much experimental effort has been expended over the last ninety years to attempt to find nonsense syllables (sometimes called CVC trigrams, indicating that they are made of a consonant, vowel, and another consonant) that are neutral and equivalent for use in studies of memory or choice, often by rating them on their “familiarity,” “association value,” or “meaningfulness” (see Jenkins [1985] for one review of the literature). The use of nonsense syllables in research always carries the potential for learned bias effects, as all letter sequences entail associations that create preferences. In related work, Fields and Arntzen (2018) have shown that the association values of pictorial stimuli affect equivalence learning. But all of this research considers only the potential biases toward the stimuli themselves, *not* the specific motor responses required of the participant when selecting those stimuli. Given my findings that kinesthetic bias involving the motions used in pressing keys is as strong as or stronger than any possible bias related to the letters on those keys, it seems that any measurement of bias or association value of stimuli must also include potential bias related to the physical form of the response, as Salzinger et al. did in 1964.

In general, the results of my experiments, in particular the comparison between the results of the two different experimental designs used in Paper III, also have implications for the potential effects of experimental design when studying complex human behavior, especially behavior that involves a skilled motor learning component such as that used in those studies. A type of design which has become standard in the behavioral laboratory, such as variable ratio reinforcement followed by extinction, used in Experiment 1 of Paper III, can have specific effects on the degree to which pre-existing biases affect the data, or on the levels of variation among participants. As was discussed earlier, we know that conditions where reinforcement is uncertain or intermittent can potentially lead to more variability in human participants (Lee, Sturme & Fields, 2007). But there are many different laboratory methods for creating conditions of reinforcement uncertainty, and the potential generalizability of the results should be taken into account when choosing the appropriate experimental design for the variable we wish to investigate.

When it comes to the operant response itself, the theoretical implications of the present work suggests that there is value in a more detailed analysis of the specific motions involved, how they may be subject to topographic bias when initially encountered, and how they may change as the response is acquired, whether due to automatization or higher-order strategies to minimize effort and maximize reinforcement. There is also value in considering both criterial and noncriterial aspects of the response, as they often co-vary. For example, speed and accuracy are related attributes, and when performing a motor program the organism is always capable of trading one for the other (Fitts & Posner, 1967, p.109). If accuracy is a criterial component of a given operant behavior and speed is not, biases affecting the latter, such as the participants' history with the operanda involved, may well affect the former, and vice versa. These types of biological constraints on human behavior affect the choice of operant behavior for laboratory experiments.

Concluding remarks

The concept of the operant is a flexible and incredibly useful tool for defining and studying the vast panorama of human behavior, whether that behavior takes place in the laboratory under the control of programmed experimental contingencies or outside in the wider world. As a researcher I have attempted to use this tool to look more closely at laboratory behavior that is *not* under the control of experimental contingencies. The methodologies used in my experiments allowed me to shine a spotlight on the noncriterial components of operant classes involving common human behaviors such as drawing and pressing keys. This gave me the ability to observe behavioral details that normally go unobserved, and led to the discovery and measurement of multiple types of topographic bias, some of which can interact with variability.

Based on the results of the studies included in this dissertation, we can think of all human behavior as under multiple control at all times. We are always simultaneously under the influence (whether strong or weak) of many previous contingencies with which we have a history, whether genetic or learned, and which can interact to produce compound effects. The effects of these contingencies can be measured, as I have attempted to do in this dissertation, by looking at the internal structure of operant occurrences, including quantifiable measurements of variability within the data as well as stability.

Skinner taught us to study behavior by establishing a set of environmental contingencies, observing the resulting behavior and then looking for orderly changes. For 80 years behavioral researchers have done exactly that, discovering incredible things along the way. Now we are in an era where experimental technology is evolving at dizzying speeds, and as scientists we should be committed to harnessing those improvements to investigate the nature of the operant in ever increasing detail, which will no doubt reveal many more new insights.

REFERENCES

- Abrahamyan, A., Silva, L. L., Dakin, S. C., Carandini, M., & Gardner, J. L. (2016). Adaptable history biases in human perceptual decisions. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, *113*(25), E3548-E3557. doi:10.1073/pnas.1518786113
- 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. doi:10.1901/jeab.1974.22-231
- Baum, W. M. (2016). Driven by consequences: The multiscale molar view of choice. *Managerial & Decision Economics*, *37*(4-5), 239-248. doi:10.1002/mde.2713
- Blampied, N.M. (2013). Single-case research and the scientist-practitioner ideal in applied psychology. In G. Madden (Ed.). *Handbook of Behavior Analysis Vol 1*. (pp. 177 – 197). Washington, DC: American Psychological Association.
- Bowen, H. J., & Kensinger, E. A. (2017). Cash or Credit? Compensation in Psychology Studies: Motivation Matters. *Collabra: Psychology*, *3*(1), 12. doi:10.1525/collabra.77
- Breland, K., & Breland, M. (1961). The misbehavior of organisms. *American Psychologist*, *16*(11), 681-684. doi:10.1037/h0040090
- Brody, J. L., Cluck, J. P., & Aragon, A. S. (1997). Participants' understanding of the process of psychological research: Informed consent. *Ethics & Behavior*, *7*(4), 285-298. doi:10.1207/s15327019eb0704_1
- Bryan William, L., & Harter, N. (1899). Studies on the telegraphic language: The acquisition of a hierarchy of habits. *Psychological Review*, *6*(4), 345-375. doi:10.1037/h0073117
- Cannella, H. I., O'Reilly, M. F., & Lancioni, G. E. (2005). Choice and preference assessment research with people with severe to profound developmental disabilities: A review of the literature. *Research in Developmental Disabilities*, *26*(1), 1-15. doi:10.1016/j.ridd.2004.01.006
- Catania, A. C. (2013). *Learning (5th Edition)*. Cornwall on Hudson, NY: Sloan Publishing.

- Cerutti, D. T. (1989). Discrimination theory of rule-governed behavior. *Journal of the Experimental Analysis of Behavior*, 51(2), 259-276. doi:10.1901/jeab.1989.51-259
- Critchfield, T. S., Meeks, E., & Stilling, S. T. (2014). Explanatory flexibility of the matching law: Situational bias interactions in football play selection. *The Psychological Record*, 64(3), 371-380. doi:10.1007/s40732-014-0064-5
- Cruvinel, A. C., & de Azebedo Pires S erio, T. M. (2008). Variabilidade comportamental: A produ o de variabilidade da dura o da resposta [Behavioral variability: The production of variability of response duration]. *Acta Comportamentalia*, 16(1), 5-23.
- Davison, M. C., & McCarthy, D. (1988). *The matching law: A research review*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Dvorak, A., Merrick, N. L., Dealey, W. L., & Ford, G. C. (1936). *Typewriting behavior: Psychology applied to teaching and learning typewriting*. New York: American Book Company.
- Ebbinghaus, H. (1885/1913). *Memory: A contribution to experimental psychology*. New York: Teachers College, Columbia University.
- Fields, L., & Arntzen, E. (2018). Meaningful stimuli and the enhancement of equivalence class formation. *Perspectives on Behavior Science*, 41(1), 69-93. doi:10.1007/s40614-017-0134-5
- Fitts, P. M., & Posner, M. I. (1967). *Human Performance: Basic Concepts in Psychology Series*. Belmont, CA: Wadsworth Publishing Co.
- Furnham, A. (1986). Response bias, social desirability and dissimulation. *Personality and Individual Differences*, 7(3), 386-400. doi:10.1016/0191-8869(86)90014-0
- Galizio, M. (1979). Contingency-shaped and rule-governed behavior: Instructional control of human loss avoidance. *Journal of the Experimental Analysis of Behavior*, 31(1), 53-70. doi:10.1901/jeab.1979.31-53

- Gilbreth, F. B., & Gilbreth, L. M. (1920). *Motion Study for the Handicapped*. London: Routledge.
- Green, G., & Saunders, R. R. (1998). Stimulus equivalence. In: Lattal K. A., Perone M. (eds) *Handbook of Research Methods in Human Operant Behavior*. Boston, MA: Springer.
- Griffiths, T. L., & Tenenbaum, J. B. (2003). Probability, algorithmic complexity, and subjective randomness. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 25, No. 25).
- Hamid, P. N. (1973). Exposure frequency and stimulus preference. *British Journal of Psychology*, 64(4), 569-577. doi:10.1111/j.2044-8295.1973.tb01383.x
- Hansen, S. W., Jensen, M. B., Pedersen, L. J., Munksgaard, L., Ladewig, J., & Matthews, L. (2002). The type of operant response affects the slope of the demand curve for food in mink. *Applied Animal Behaviour Science*, 76(4), 327-338. doi:10.1016/S0168-1591(02)00008-4
- Hansen, S., & Arntzen, E. (2017). *Scan-pattern in sequentially arranged matching-to-sample tasks*. Paper presented at the Association for Behavior Analysis International Ninth International Conference, Paris, France.
- Herrnstein, R.J. (1961). Relative and absolute strength of responses as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behaviour*, 4, 267–72. doi: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. New York: Appleton-Century-Crofts.
- Hirai, M., Okouchi, H., Matsumoto, A., & Lattal, K. A. (2011). Some Determinants of Remote Behavioral History Effects in Humans. *Journal of the Experimental Analysis of Behavior*, 96(3), 387–415. doi:10.1901/jeab.2011.96-387
- Jenkins, J. J. (1985). Nonsense syllables: Comprehending the “almost incomprehensible variation.” *Journal of Experimental Psychology: Learning, Memory and Cognition*, 11, 455–460. doi: 10.1037/0278-7393.11.3.455

- Jensen, G., & Neuringer, A. (2009). Barycentric Extension of Generalized Matching. *Journal of the Experimental Analysis of Behavior*, 92(2), 139–159. doi:10.1901/jeab.2009.92-139
- Jones, L. D., & Mechner, F. (2007). The effect of number of prior operant repetitions on choice behavior. New York: The Mechner Foundation. Retrieved from <http://mechnerfoundation.org/category/downloads/basic-research/>
- Jonides, J., Naveh-Benjamin, M. & Palmer, J. (1985). Assessing automaticity. *Acta Psychologica*, 60, 157-171. doi:10.1016/0001-6918(85)90053-8.
- Kono, M. (2017). The effects of fixed-interval schedules on variability of pigeons' pecking location. *Journal of the Experimental Analysis of Behavior*, 108(2), 290-304. doi:10.1002/jeab.276
- Lattal, K. A., & St. Peter Pipkin, C. (2009). Resurgence of previously reinforced responding: Research and application. *The Behavior Analyst Today*, 10(2), 254-266. doi:10.1037/h0100669
- Latterman, J., & Merz, J. F. (2001). How much are subjects paid to participate in research? *American Journal of Bioethics*, 1(2), 45-46. doi:10.1162/152651601300169040
- Lee, R., Sturmey, P., & Fields, L. (2007). Schedule-induced and operant mechanisms that influence response variability: A review and implications for future investigations. *The Psychological Record*, 57, 429-455. doi:10.1007/bf03395586
- Loeb, E. P., Giszter, S. F., Saltiel, P., Bizzi, E., & Mussa-Ivaldi, F. A. (2000). Output units of motor behavior: An experimental and modeling study. *Journal of Cognitive Neuroscience*, 12(1), 78-97. doi:10.1162/08989290051137611
- Logan, G. D. (1985). Skill and automaticity: Relations, implications, and future directions. *Canadian Journal of Psychology/Revue canadienne de psychologie*, 39(2), 367-386. doi:10.1037/h0080066
- 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. doi:10.1037//0278-7393.13.3.392

- Mace, F. C., & Critchfield, T. S. (2010). Translational research in behavior analysis: Historical traditions and imperative for the future. *Journal of the Experimental Analysis of Behavior*, *93*(3), 293–312. doi:10.1901/jeab.2010.93-293
- Mechner, F. (1994a). 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*. Cambridge, MA: Cambridge Center for Behavioral Studies.
- Mechner, F. (1994b). *Learning and practicing skilled performance*. New York, NY: Behavioral Science Applications, Inc. Retrieved from <http://mechnerfoundation.org/category/downloads/educational-innovation/>
- 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*, 45-68. doi:10.1007/bf03395212
- Mechner, F., & Jones, L. D. (2011). Effects of sequential aspects of learning history. *Mexican Journal of Behavior Analysis*, *37*, 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*, 63–83. doi:10.5514/rmac.v41.i2.63729
- Michael, J. (1993). Establishing operations. *The Behavior Analyst*, *16*(2), 191-206. doi:10.1007/BF03392623
- Milgram, S. (1963). Behavioral Study of obedience. *The Journal of Abnormal and Social Psychology*, *67*(4), 371-378. doi:10.1037/h0040525
- Morikawa, K., & McBeath, M. K. (1992). Lateral motion bias associated with reading direction. *Vision Research*, *32*(6), 1137-1141. doi:10.1016/0042-6989(92)90014-A
- Neuringer, A. (2002). Operant variability: Evidence, functions and theory. *Psychonomic Bulletin & Review*, *9*(4), 672-705. doi:10.3758/bf03196324

- Page, S., & Neuringer, A. (1985). Variability is an operant. *Journal of Experimental Psychology: Animal Behavior Processes*, *11*(3), 429–452. doi:10.1037/0097-7403.11.3.429
- Palmer, D. C. (2012). The role of atomic repertoires in complex behavior. *The Behavior Analyst*, *35*(1), 59-73. doi:10.1007/BF03392266
- Perone, M. (2018). How I learned to stop worrying and love replication failures. *Perspectives on Behavior Science*. doi:10.1007/s40614-018-0153-x
- Pew, R. W. (1974). Human perceptual-motor performance. In B. H. Kantowitz (Ed.), *Human information processing*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Pisacreta, R. (1982). Preferences among stimulus matches in the pigeon. *Journal of The Experimental Analysis of Behavior*, *38*(2), 191-199. doi:10.1901/jeab.1982.38-191
- Salzinger, K., Feldman, R. S., & Portnoy, S. (1964). The effects of reinforcement on verbal and nonverbal response. *Journal of General Psychology*, *70*(2), 225-234. doi:10.1080/00221309.1964.9920592
- Schwartz, B., & Reilly, M. (1983). Response stereotypy without automaticity in pigeons. *Learning and Motivation*, *14*(3), 253-270. doi:10.1016/0023-9690(83)90017-6
- Shaffer, D. M., McManama, E., Swank, C., Williams, M., & Durgin, F. H. (2014). Anchoring in action: Manual estimates of slant are powerfully biased toward initial hand orientation and are correlated with verbal report. *Journal of Experimental Psychology: Human Perception and Performance*, *40*(3), 1203-1212. doi:10.1037/a0036217
- Sidman, M. (1960). *Tactics of Scientific Research*. Boston, MA: Authors Cooperative, Inc.
- Skinner, B. F. (1938). *The Behavior of Organisms: An Experimental Analysis*. Oxford, England: Appleton-Century.
- Skinner, B. F. (1969). *Contingencies of Reinforcement: A Theoretical Analysis*. New York: Appleton-Century-Crofts.
- Stewart, I., Barnes-Holmes, D., Roche, B., & Smeets, P. M. (2002). Stimulus equivalence and nonarbitrary relations. *The Psychological Record*, *52*(1), 77-88. doi:10.1007/bf03395415

- 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. doi:10.1901/jeab.1995.63-151
- Svartdal, F. (1989). Shaping of rule-governed behaviour. *Scandinavian Journal of Psychology*, 30(4), 304-314. doi:10.1111/j.1467-9450.1989.tb01093.x
- Svartdal, F. (1991). Operant modulation of low-level attributes of rule-governed behavior by nonverbal contingencies. *Learning and Motivation*, 22(4), 406-420. doi:10.1016/0023-9690(91)90004-R
- Tatham, T. A., & Wanchisen, B. A. (1998). Behavioral history: A definition and some common findings from two areas of research. *The Behavior Analyst*, 21(2), 241–251. doi:10.1007/bf03391966
- Tatler, B. W. (2007). The central fixation bias in scene viewing: Selecting an optimal viewing position independently of motor biases and image feature distributions. *Journal of Vision* 7(14):4, 1–17. doi:10.1167/7.14.4
- Toner, J., Montero, B. G., & Moran, A. (2015). The perils of automaticity. *Review of General Psychology*, 19(4), 431-442. doi:10.1037/gpr0000054
- 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):4, 1-16. doi:10.1167/9.7.4
- Wanchisen, B. A. (1990). Forgetting the lessons of history. *The Behavior Analyst*, 13(1), 31–37. doi:10.1007/bf03392515
- Wilkinson, L., & Task Force on Statistical Inference, American Psychological Association, Science Directorate. (1999). Statistical methods in psychology journals: Guidelines and explanations. *American Psychologist*, 54(8), 594-604. doi:10.1037/0003-066X.54.8.594
- William, L. B., & Harter, N. (1899). Studies on the telegraphic language: The acquisition of a hierarchy of habits. *Psychological Review*, 6(4), 345-375. doi:10.1037/h0073117

- Williams, E. P, & Walter, J. K. (2015). When does the amount we pay research participants become “undue influence”? *American Medical Association Journal of Ethics*, *17*(12), 1116-1121. doi:10.1001/journalofethics.2015.17.12.ecas2-1512
- Wine, B., Reis, M., & Hantula, D. A. (2014). An evaluation of stimulus preference assessment methodology in organizational behavior management. *Journal of Organizational Behavior Management*, *34*(1), 7-15. doi:10.1080/01608061.2013.873379
- World Medical Association (2013). Declaration of Helsinki: Ethical principles for medical research involving human subjects. *Journal of the American Medical Association*, *310*(20), 2191–2194. doi:10.1001/jama.2013.281053

Paper I

Systematic Operant Bias Observed in Human Participants During Research on Choice. Jones, L. D. & Mechner, F. (2013) *European Journal of Behavior Analysis* 14, 295-311.

DOI: <https://dx.doi.org/10.1080/15021149.2013.11434462>

[Article not attached due to copyright]

Paper II

Kinesthetic Operant Bias in Keyboard-Based Choice Behavior. Jones, L. D. & Mechner, F. (2015) *European Journal of Behavior Analysis* 16, 202-228. DOI: <https://dx.doi.org/10.1080/15021149.2015.1093796>

[Article not attached due to copyright]

Paper III

Noncriterial Behavioral Variability and Related Operant Bias in Humans.

Jones, L. D. & Mechner, F. (2020) European Journal of Behaviour Analysis.

DOI: <https://dx.doi.org/10.1080/15021149.2020.1745526>

This is an original manuscript / preprint of an article published by Taylor & Francis in European Journal of Behavior Analysis on 07/04/2020, available online: <https://www.tandfonline.com/doi/full/10.1080/15021149.2020.1745526>.

Title: Noncriterial behavioral variability and related operant bias in humans

First Author: Laurilyn D. Jones

Affiliation: Oslo Metropolitan University and the Mechner Foundation

Second Author: Francis Mechner

Affiliation: The Mechner Foundation

Suggested Running Head: Noncriterial variability and operant bias

Corresponding Author: Laurilyn D. Jones

The Mechner Foundation

200 Central Park South, 18E

New York, NY 10019

(917) 292-4871

ldj@mechnerfoundation.org

Noncriterial behavioral variability and related operant bias in humans

Laurilyn D. Jones

Francis Mechner

All operant behaviors have multiple dimensions in addition to those designated by the experimenter as criterial for reinforcement, and behavioral variation occurs in all of those dimensions. In addition, all dimensions of an operant can reflect possible bias due to the pre-experimental histories of the participants. In two experiments designed to measure both noncriterial variability and operant bias, human participants performed an operant consisting of typing 14 or more keystrokes on the computer keyboard. The first and last keystrokes were mandated, while the middle 12 (or more) were allowed to vary. The first experiment involved nine sessions of monetary reinforcement on a variable ratio schedule followed by one session of extinction, while the second required nine sessions of continuous feedback (without monetary reinforcement) followed by a final “test” session with multiple contingencies designed to disrupt the participants’ behavior. There were significant differences in variability among the individual participants, as well as systematic effects of the different experimental designs. Despite not being required for reinforcement, operant-to-operant variability was high overall. The “test” session of Experiment 2 resulted in a much larger increase in variability than did extinction, in Experiment 1. Looking at operant bias, there was an overall preference for the letter keys in the center of the keyboard. However, the participants also showed a strong bias *against* center keys when choosing either the first or last letter in each operant sequence. There were also correlations between measures of variability and bias, showing them to be related.

Keywords: Behavioral variability, noncriterial variability, stereotypy, operant bias, bias, human participants

The operant can be thought of as the basic unit of the experimental analysis of behavior, but each operant is also more than a single instance of behavior. Every

individual occurrence of an operant 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 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 dimensions of the operant (such as the duration of the lever press) are free to vary, and necessarily do. No two operant occurrences can ever be exactly alike.

All operants, therefore, no matter how simple, can be thought of as sequences of behaviors (Mechner, 1994; Mechner, Hyten, Field & Madden, 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, or as examples of response induction or generalization, 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 affect the reinforcers received and to be affected by them in turn (Neuringer, 2002). For example, in the case of “superstitious” behavior, noncriterial aspects of operant behavior (despite their not being required for reinforcement) can be shaped by reinforcement, then affect the form of the organism’s response and thus come to control access to reinforcement. An experimental methodology that permits one to measure at least some of the noncriterial dimensions of each operant response thus allows an experimenter to analyze behavior in far greater depth than that permitted by methodologies that only record the emission of each operant as a single data point. This

is particularly helpful when studying complex topics such as behavioral variability vs. stereotypy, and the existence of participant bias for or against a given operant behavior (referred to as operant bias, to distinguish the phenomenon from the more commonly-referenced stimulus bias). All organisms have biases, of course, in the sense that they prefer certain behaviors over others. Operant bias, as used here, refers to any consistent systematic preference for one *equivalent and equally available* behavior over another, which cannot be explain by the current contingencies. It can have many different causes (such as anatomical variables, learning history, topographical constraints on behavior, to name only a few), which may or may not—usually not—be known to the experimenter.

The revealed operant is such an in-depth methodology; it allows tracking of both the criterial and certain noncriterial attributes of each individual occurrence of an operant class (Mechner, 1994; Mechner et al., 1997). The revealed operant unit consists of a sequence of recorded actions (sub-operants) whose beginning and end are marked by behaviorally distinct events. Some of these sub-operants are mandated by the definition of the operant (criterial), and others are not (noncriterial). The experiments presented in this paper used a revealed operant consisting of a sequence of keystrokes on a modified computer keyboard in order to measure both variability and operant 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, Foster & McEwan, 2009; Margulies,

1961; Morgan & Lee, 1996; Neuringer, Kornell & Olufs, 2001, Notterman, 1959). It is also possible to generate higher levels of noncriterial variability by changing the response requirement (Mechner et al, 1997; Tatham, Wanchisen & Hineline, 1993). In general, intermittent reinforcement often results in greater noncriterial variability than continuous reinforcement (Antonitis, 1951; Boren, Moerschbaecher, & Whyte, 1978; Eckerman & Lanson, 1969), however the literature also contains examples of contrary results (Gates & Fixen, 1968; Herrnstein, 1961; see Lee, Sturmey & Fields, 2007, for a review).

In contrast, there has been hardly any comprehensive experimental work done on operant bias, particularly with human participants. Researchers generally attempt to implement experimental methods to reduce potential bias for or against the responses they use in their studies, just as they generally attempt to reduce uncontrolled effects on their independent variables that can generate variability (Sidman, 1960). But bias is not often studied directly. It is, however, measured in the experimental analysis of matching, as a variable in the equation indicating how closely the relative rate of different responses match their relative reinforcement rates (Baum, 1974).

Researchers looking at variability occasionally also examine noncriterial operant bias superficially, 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 operant 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 operant bias can have interactive effects on behavior. For these reasons, the experiments presented in the present paper use the revealed operant methodology to generate the relevant data.

In both experiments, human participants learned to execute a lengthy 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 was exactly the same, 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, operant 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 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 experienced in previous multi-day studies.

Participants received a flat fee of \$10 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 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 prevent dropouts.

All participants were informed when they signed up that they were free to withdraw from the experiment at any time, keeping the already-earned \$10 per session participation fees, but would receive the final payment only if they completed all ten sessions. Of the 12 participants initially recruited, 4 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. Unfortunately, due to the nature of laboratory experiments with human participants, these variables simply cannot be controlled in any more systematic way, but requiring documentation of their behavior between sessions, at the very least, increases participants' awareness of such behavior and the need to keep it constant. All

participants received a written explanation of all requirements of the study (as well as their right to withdraw at any time) at the beginning of the first session and signed a document giving their informed consent to participate in the study. 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]

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 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. A feedback stimulus was programmed to help participants keep track of the operants they emitted: the entire screen turned blue when the space bar was pressed to begin each operant and returned to black with the press of the enter key, remaining black until the participant began to perform the next operant. 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 middle part of each operant were not mandated and were allowed to vary freely except for one restriction: a modified within-operant-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-operant variability. 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. Thus in both experiments, a certain minimum level of variability within each operant sequence was a criterial dimension of the operant, but the level of variability among complete operants was completely noncriterial and did not affect reinforcement in any way.

If a participant's performance violated this within-operant restriction, the operant was not terminated, but 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. Therefore, it was possible for a participant who realized their mistake mid-operant to simply type additional letters fulfilling the variability restriction before pressing the enter key. There was no maximum limit on the number of keypresses per operant; as long as at least 12 acceptable keypresses were registered before the enter key was pressed, each operant could be as long as the participant chose.

Experiment 1 was designed to see if this type of operant behavior would reach a steady state after many sessions of intermittent reinforcement, before testing the stability of the performance by putting it under extinction. Unpredictable reinforcement, especially in experimental situations which allow for variability, can potentially generate high levels of such variability; 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 much higher rates of operant emission than an interval schedule. For the test condition, extinction was chosen in order to see whether the well-documented phenomenon of extinction-induced increase in variability would occur with this type of operant.

Baseline Condition (with Procedure)

The first nine sessions of Experiment 1 were programmed to be 440 operants each; since only valid operants (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. Since participants worked at different speeds, the duration of each session also varied. Valid operants were reinforced throughout the first nine sessions on a VR4 schedule, programmed in a pseudo-random pattern which repeated every 80 operants, so as to be unpredictable to the participants while still keeping the number of reinforcers earned relatively consistent throughout the session. The reinforcer used was money, with the actual amount of each reinforcer delivered also programmed to vary among 11, 19 or 26 cents on a pseudo-randomized schedule. Again, this was done to increase the unpredictability of reinforcement. The 110 reinforcers programmed during each of the first nine sessions added up to \$21.34 per session, making the average reinforcer amount 19.4 cents.

Reinforcement was signaled by a 440Hz “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 a consumatory response for an edible reinforcer, was required after each reinforcer presentation, consisting of typing the amount earned on the number keypad, followed by 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 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 operants of the first session only, the computer screen also displayed a large (3” x 3”) green square for .5 seconds after *every* valid operant as feedback, regardless of whether or not the operant

received reinforcement. This was done in an attempt to ensure that participants initially learned how to produce valid operants consistently and independently of reinforcement.

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 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 and so forth. The 12 or more letters can be anything you want, except 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 then 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 operants 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 operants, again with only valid operants being counted toward this total. This session was programmed as unsignaled extinction: 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, thus providing the participants with no way to discriminate the new contingency. No additional 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.”

Results

Noncriterial Variability

An operant consisting of a lengthy sequence of discrete, highly variable behaviors that are precisely tracked, such as the operant used in these experiments, offers an excellent 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 is the proportion of each participant’s operants in each session that are unique; i.e., the fraction of the total remaining if one removes all repetitions of operants previously emitted during the same session. This might be termed the “Unique

Value,” and it serves as a useful analogue for the commonly used U-value measure of variability, as it too is a single value falling somewhere between 0 (perfectly stereotyped) and 1 (perfectly variable). There are certainly limitations to the information provided by this statistic (just as there are for U-value—see Barba, 2012; and Kong, McEwan, Bizo & Foster, 2017), but for the current experiment it provides an excellent indication 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 almost immediately to very 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 throughout the study. The double vertical line between Sessions 9 and 10 indicates the change in experimental contingencies. Note that there is *not* a large increase in variability in Session 10, when the condition change was implemented and extinction was imposed.

[Figure 2]

In spite of the significant differences among participants, some fairly consistent effects of elapsed time, repetition and reinforcement on noncriterial variability levels can be seen. First of all, the differences *within* the two distinct groups of participants tend to decrease from session to session, showing that overall variability levels in the two groups

are stabilizing. In addition, 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 operants emitted, since participants were allowed at all times to type as many letters as they wanted within the body of each operant. Figure 3 shows the proportion of operants 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 operants 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 at the highest possible level using a one-way ANOVA ($F(1, 78) = 21.98, p = .00001$). 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 and then decline, while participants in the high-variability group are more variable in operant length from session to session and cover a wider range of values.

[Figure 3]

Finally, although the accuracy with which this type of operant was performed is not a direct measure of operant variability, it is quite closely related to it, as variability in general often occurs in conjunction with mistakes or disruptions in the behavior stream. In all sessions of Experiment 1, the proportion of invalid operants (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 operants in the session as a whole. The average difference in Unique Value between total and invalid operants 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 operants 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 significantly higher than the others, the difference in accuracy between the two groups of participants is still statistically significant at the highest level ($F(1, 78) = 27.64, p = .000001$).

[Figure 4]

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 operant 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 preferred 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 preference for letter keys from the middle of the three rows (letters g, h, j and k), vs. preference for 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 preference for 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, 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 operant. The gray horizontal 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 biases 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.)

[Figure 5]

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, 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 operant than are low-variability participants. Again, these differences between groups are statistically significant at the highest possible level ($F(1, 78) = 84.82, p = 4.23 \times 10^{-14}$ for the middle row, $F(1, 78) = 66.65, p = 4.53 \times 10^{-12}$ for the center section).

The bottom left panel of Figure 5 shows that the same strong 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 ($F(1, 78) = 40.82, p = 1.11 \times 10^{-8}$). However, the proportion of last keystrokes in each operant 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 preference. 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 just as much as the keypresses themselves. Once the participant has chosen the first letter of each operant, 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 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 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.

[Figure 6]

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. Importantly, the results of this study provide data showing that variability and operant bias were linked, creating a combined effect. In particular, it appears from the bifurcated nature of the results that participants in this study responded to the experimental contingencies in either one of two ways: 1) with very high variability overall and a tendency to choose their 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 from the edges of the space available for typing. These opposite patterns of responding suggest that high-variability participants are choosing their keystrokes more pseudo-randomly, while low-variability participants instead develop a more stereotyped operant sequence that involves moving their hands systematically up and/or down and side to side over the letter keys. These results clearly show an interaction between noncriterial variability and operant bias effects.

Overall, it is interesting that operant-to-operant variability levels were so high despite being completely noncriterial. Participants could have chosen to type the exact same 12-letter sequence for every operant repetition during the entire experiment, but none did so; even the low-variability group still emit anywhere between 1.5% and 12% completely unique operants in each session, a relatively robust level of operant-to-operant variability. In addition, with this type of operant we did not see a large increase in noncriterial variability when extinction was imposed, as one might have expected given the results in the literature on extinction-induced variability.

After analyzing the data from Experiment 1, the decision was made to run another experiment using the same operant 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. The contingencies put in place for Experiment 2 were designed to fulfill one of the fundamental aims of translational research by creating a laboratory analog of real-world behavior.

EXPERIMENT 2

Method

Participants

16 adults were recruited in the same manner as that described for Experiment 1, except that they 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 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 performed 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 extinction, 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 a performer, or a student of any skill that includes a physical component (motor learning). 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.

Baseline Condition (with Procedure)

The first nine sessions of Experiment 2 consisted of 400 operants each; as in Experiment 1, only valid operants were counted toward this total. However, in Experiment 2, *all* valid operants during Sessions 1 through 9 received continuous feedback in the form of the 3” x 3” green square described under Experiment 1 above, which was displayed upon the completion of all valid operants. No monetary reinforcers were given. The green square, in addition to providing feedback to the participants, could also potentially be thought of as a type of weak conditioned reinforcer.

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 that 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 which individual one is the cause of changes in the dependent 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 operants 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 operants, whether valid or not—all operant attempts were counted toward this total.

In addition to positive monetary reinforcement of valid operants, invalid operants 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 operant, or between operants, 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 110Hz 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 as a whole. In order to be considered valid, the 12-or-more letter sequence making up each operant 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 extreme levels of stereotypy. (Similar high levels of stereotypy as a reaction to test conditions had been observed in previous, related studies.) The within-operand 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 to a high of \$331.60 (out of a total possible amount of \$368.00). 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.

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 graph shows no clear trend; there is 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 here, nor are the session-to-session trends. There is, however, a noticeable increase in the proportion of unique operants 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 requirement, however the increase recorded is far beyond what would have been necessary to fulfil that minimal requirement.

[Figure 7]

Figure 8 shows the proportion of operants in each session of Experiment 2 that were longer than the minimum 12 letters required (an alternate measurement of noncriterial variability), for all participants. Again, there is a large amount of variation among the participants. Some participants are more likely overall to emit operants longer than the minimum length; amongst these participants, there is a tendency for the proportion of longer operants to drop after the first session and then gradually increase again during the last few sessions. In general, due to these few participants, the average length of operants emitted in the last three sessions of Experiment 2 is much greater than in the corresponding sessions of Experiment 1.

[Figure 8]

As it was in Experiment 1, in Experiment 2 the proportion of *invalid* operants (those not conforming to the criterial requirements) that were unique was much higher than the proportion of unique operants in the session as a whole. The average difference

in Unique Value between total and invalid operants in Experiment 2 is .50; this effect is consistent, occurring in all participants. Figure 9 shows the percentage of each participant's operants 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.

[Figure 9]

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 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) used for Experiment 1. 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, and the gray horizontal line across each panel shows the indifference point.

[Figure 10]

In general, the bias data for Experiment 2 shows the same overall trends as the variability data: we see more variance in preference 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 most 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 points spread across a much wider range of values than found in the Experiment 1 results, and there are two anomalous participants (S101 and S106) who show a completely different bias pattern for keys from the middle row. These two participants 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 preference for 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, 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.

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

As was done for Experiment 1, we also analyzed the transition between the first and second keystrokes for all operants 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.

[Figure 11]

Discussion

Overall, the experimental contingencies of Experiment 2 generated large amounts of noncriterial variability and significant operant 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 develop only 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 within the behavior of individual participants; clearly continuous feedback coupled with the possibility of future reinforcement exerts a relatively weak control over individual behavior, at least compared to a more traditional experimental design. Or, depending on your point of view regarding the function of the green square, continuous reinforcement using a weak conditioned reinforcer exerts weaker control than intermittent and unpredictable reinforcement using a very strong conditioned reinforcer (money) coupled with an overt consummatory response.

There are, however, some noticeable trends to be found 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 operants longer than the minimum length. Given the experimental data showing that continuous reinforcement results in less behavioral variability than intermittent, one might expect variability under these conditions to be lower; apparently in this type of experiment 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 operant preference. This was not the case in Experiment 1.

Finally, data regarding the transition between 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 simply easier and faster to type strings of letter keys that are closer together—was minimized by all participants under conditions of continuous feedback in a way that was simply not the case for the participants under conditions of intermittent reinforcement, who instead were more likely to choose the letter keys making up their operants based on their absolute rather than relative spatial position on the keyboard.

GENERAL DISCUSSION

Of the experimental findings presented in this paper, some are generally applicable, i.e. true for both studies in spite of their different experimental designs, and thus can potentially tell us something about noncriterial variability and operant bias in general, and the way these two linked phenomena interact in shaping complex human behavior. Other findings from the two experiments demonstrate the different results due to their different experimental designs, allowing us to draw tentative conclusions about the effects of intermittent reinforcement and extinction 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* operants—far higher than one might expect from such a highly-practiced, automatized response. 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 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 more response effort than the production of unique sequences. Even performances with the lowest noncriterial variability levels overall never became completely stereotyped but continued to contain unique operants 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 operant) 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—alternation between two responses—when a Lag 1 criterial variability requirement was implemented. 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 operants. With the revealed operant, 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 (alternating between the same two 12-letter patterns).

It is worth noting that the revealed operant, unlike the sequences of behaviors often used as operants in studies of variability, always has a behaviorally distinct beginning and ending sub-operant (Mechner et al, 1997). This feature can help slow or disrupt the automatization or fusion of long strings of operants 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 a single longer routine, usually faster and more fluent than the sum of its 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 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 (response generalization across a range of similar responses)—and then decreases as the criterial response is automatized under differential reinforcement (Jarmolowicz, Hudnall, Darden, Lemley & Sofis, 2015). In addition, in both experiments, we can see that invalid operants are more likely to be variable operants. In all sessions and for all participants, a much greater proportion of invalid operants were unique than valid ones. Variability is thus correlated with disruptions in an otherwise stable behavior stream.

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 preference (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 infinite number of valid options for each operant they emit. Other variability researchers have also found differences among their human participants (Eckerman & Vreeland, 1973; Peleg, Martin & Holth, 2017). 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, Dozier, Briggs and Juanico (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, choosing which letter keys to press in each operant as part of differing response strategies. In Machado's (1992, 1993) experiments on criterial variability with pigeons, they responded to a frequency-based variability requirement either with a higher-order stereotypical sequence alternating between responses, or seemingly randomly, 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, Lee, Ferreira, da Silva and Caramori (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 previously experienced contingencies or verbal rules (strategies). 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 able to emit truly random behavior; they tend to avoid repetition and symmetry and emit too many alternating responses (Lopes & Oden, 1987). Noncriterial variability in human operant responding, therefore, is *not* random, but controlled by the prevailing 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 third general finding of these two studies, therefore, is that significant operant bias regarding letter choice when typing 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 strong evidence of being kinesthetic in origin, relating 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. Since the operant 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.

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 operants. 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. There is a related perceptual phenomenon referred to as center bias: the general fixation of the gaze at the center of the visual field (Tseng, Carmi, Cameron, Munoz & Itti, 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 in Session 10 of Experiment 2 (the "test session"), being present in 8 of the 9 participants. Jones and Mechner (2015) found analogous results when human participants spent multiple sessions typing several different *critical* 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 as a reinforcer. 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 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. 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 operant 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.

Finally, in half of the Experiment 1 participants and in all the Experiment 2 participants, we found a very strong preference for choosing a second letter key adjacent to the first in the sequence. This type of bias would seem to be a kinesthetic preference based on minimizing the motion required to complete the operant—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 preferred those sequences in which the first two letters were adjacent (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 carried throughout the rest of the experiment, contributing to individual differences. This would explain why only half the participants in Experiment 1 show this effect.

Due to their unique methodology, these experiments also allow us to actually measure and analyze the *interaction* between pre-existing operant bias and noncriterial behavioral variability, two of the major behavioral dimensions that shape the larger operant class. And the evidence for this interaction is the final, and perhaps most interesting, finding of general applicability that we can extract from the data of the two experiments presented in this paper. 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 emitted more variable operant sequences, showing that those whose performances were more variable 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 his or her overall level of variability, perhaps due to the effect of intermittent reinforcement and possible superstitious conditioning, while in Experiment 2 it is not.

In Experiment 2, this type of correlation did not occur, but there was a sharp rise in errors for all participants during Session 10 (the “test session) *not* seen for Session 10 (extinction) of Experiment 1, and a corresponding shift in variability level and/or bias recorded for a majority of the participants. The imposition of deliberate “stressors” such as high-value reinforcers, punishment contingencies, and strict time limits in Session 10 of Experiment 2 thus caused a noticeable shift (whether an increase or decrease) in variability, bias and accuracy for the participants, whereas unsignaled extinction, in Experiment 1, did not. Given the well-documented phenomenon of extinction-induced variability, the effect seen in Session 10 of Experiment 1 is much smaller than one might have expected. However, few of the studies done on extinction-induced noncriterial variability have been done with human participants, and those have had somewhat mixed results (Kinloch, Foster & McEwan, 2009; Maes, 2003; Morgan & Lee, 1996). Kinloch, Foster and McEwan (2009) were able to elicit greater variability under extinction in one noncriterial dimension of their operant (interresponse time) but not in a different one (the size of the rectangle drawn).

In all experimental studies of behavior, particularly 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 their interactive effects. The findings here have important potential implications for the design of experiments involving complex operant behaviors and choice among alternative behaviors, particularly studies of behavioral variability. Unlike the types of operants usually used to study variability, the type of revealed operant preparation used in the present experiments is uniquely suited for the measurement of operant bias as well, and the authors hope to see future studies of *both* the interrelated phenomena of variability and operant bias using this or similar methodologies.

REFERENCES

- Antonitis, J. J. (1951). Response variability in the white rat during conditioning, extinction, and reconditioning. *Journal of Experimental Psychology*, 42(4), 273-281.
- Barba, L. S. (2012). Operant variability: A conceptual analysis. *The Behavior Analyst*, 35, 213-227.
- Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior*, 22, 231-242.
- Boren, J. J., Moerschbaeche, 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.
- 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, 600-621.
- 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, 73-80.
- Eckerman, D. A., & Vreeland, R. (1973). Response variability for humans receiving continuous, intermittent, or no positive experimenter feedback. *Bulletin of the Psychonomic Society*, 2, 297-299.
- Gates, J. J., & Fixen, D. L. (1968). Response variability in retarded children. *Journal of Experimental Child Psychology*, 6, 306-320.

- Gilbreth, F. B., & Gilbreth, L. M. (1920). *Motion Study for the Handicapped*. London: Routledge.
- Glencross, D. J. (1973). Temporal organization in a repetitive speed skill. *Ergonomics*, *16*, 765-776.
- Herrnstein, R. J. (1961). Stereotypy and intermittent reinforcement. *Science*, *133*, 2067-2069.
- Herrnstein, R. J. (1966). Superstition: A corollary of the principles of operant conditioning. In W. K. Honig, (Ed.) *Operant Behavior*, pp. 33-51. New York: 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*, 139-147.
- 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.
- 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.
- Jones, L. D., & Mechner, F. (2013). Systematic operant bias observed in human participants during research on choice. *European Journal of Behavior Analysis*, *14*, 295-311.
- Jones, L. D., & Mechner, F. (2015). Kinesthetic operant bias in keyboard-based choice behavior. *European Journal of Behavior Analysis*, *16*(2), 202-228.

- Kinloch, J. M., Foster, T. M., & McEwan, J. S. A. (2009). Extinction-induced variability in human behavior. *The Psychological Record, 59*, 347-370.
- Kong, X., McEwan, J. S., Bizo, L. A., and 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, V. L. (1996). Superstitious location changes by human beings. *The Psychological Record, 46*, 71-86.
- Lee, R., Sturmey, P., & Fields, L. (2007). Schedule-induced and operant mechanisms that influence response variability: A review and implications for future investigations. *The Psychological Record, 57*, 429-455.
- 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.
- Machado, A. (1992). Behavioral variability and frequency-dependent selection. *Journal of the Experimental Analysis of Behavior, 58*, 241-263.
- Machado, A. (1993). Learning variable and stereotypical sequences of responses: Some data and a new model. *Behavioral Processes, 30*, 103-130.
- Machado, A. (1997). Increasing the variability of response sequences in pigeons by adjusting the frequency of switching between two keys. *Journal of the Experimental Analysis of Behavior, 68*, 1-25.
- Maes, J. H. R. (2003). Response stability and variability induced in humans by different feedback contingencies. *Learning & Behavior, 31*(4), 332-348.

- Margulies, S. (1961). Response duration in operant level, regular reinforcement, and extinction. *Journal of the Experimental Analysis of Behavior*, 4, 41-48.
- 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*. Cambridge, MA: 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, 45-68.
- Morgan, D. L., & Lee, K. (1996). Extinction-induced response variability in humans. *The Psychological Record*, 46, 145-159.
- Neuringer, A. (2002). Operant variability: Evidence, functions, and theory. *Psychonomic Bulletin & Review*, 9(4), 672-705.
- Neuringer, A., Kornell, N., & Olufs, M. (2001). Stability and variability in extinction. *Journal of Experimental Psychology: Animal Behavior Processes*, 27(1), 79-94.
- 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*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Notterman, J. M. (1959). Force emission during bar pressing. *Journal of Experimental Psychology*, 58, 341-347.

- 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.
- Schwartz, B. (1982). Reinforcement-induced behavioral stereotypy: How not to teach people to discover rules. *Journal of Experimental Psychology: General, 111*(1), 23-59.
- Sidman, M. (1960). *Tactics of scientific research*. Boston, MA: Authors Cooperative.
- Skinner, B. F. (1959). The generic nature of the concepts of stimulus and response, *The Journal of General Psychology, 1935, 12*, 40-65. Reprinted in Skinner, B. F., *Cumulative record* (pp. 347-366). New York: Appleton-Century-Crofts.
- 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*, 349-359.
- 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.



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




Fig. 2 Proportion of the total number of operants emitted by each participant in each session of Experiment 1 that are unique (referred to as the “Unique Value”). (Only the first (or only) occurrence of each unique operant is counted toward this value; if it is repeated within the same session by the same participants, the repetitions are not considered.) 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, i.e. when extinction was imposed.




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

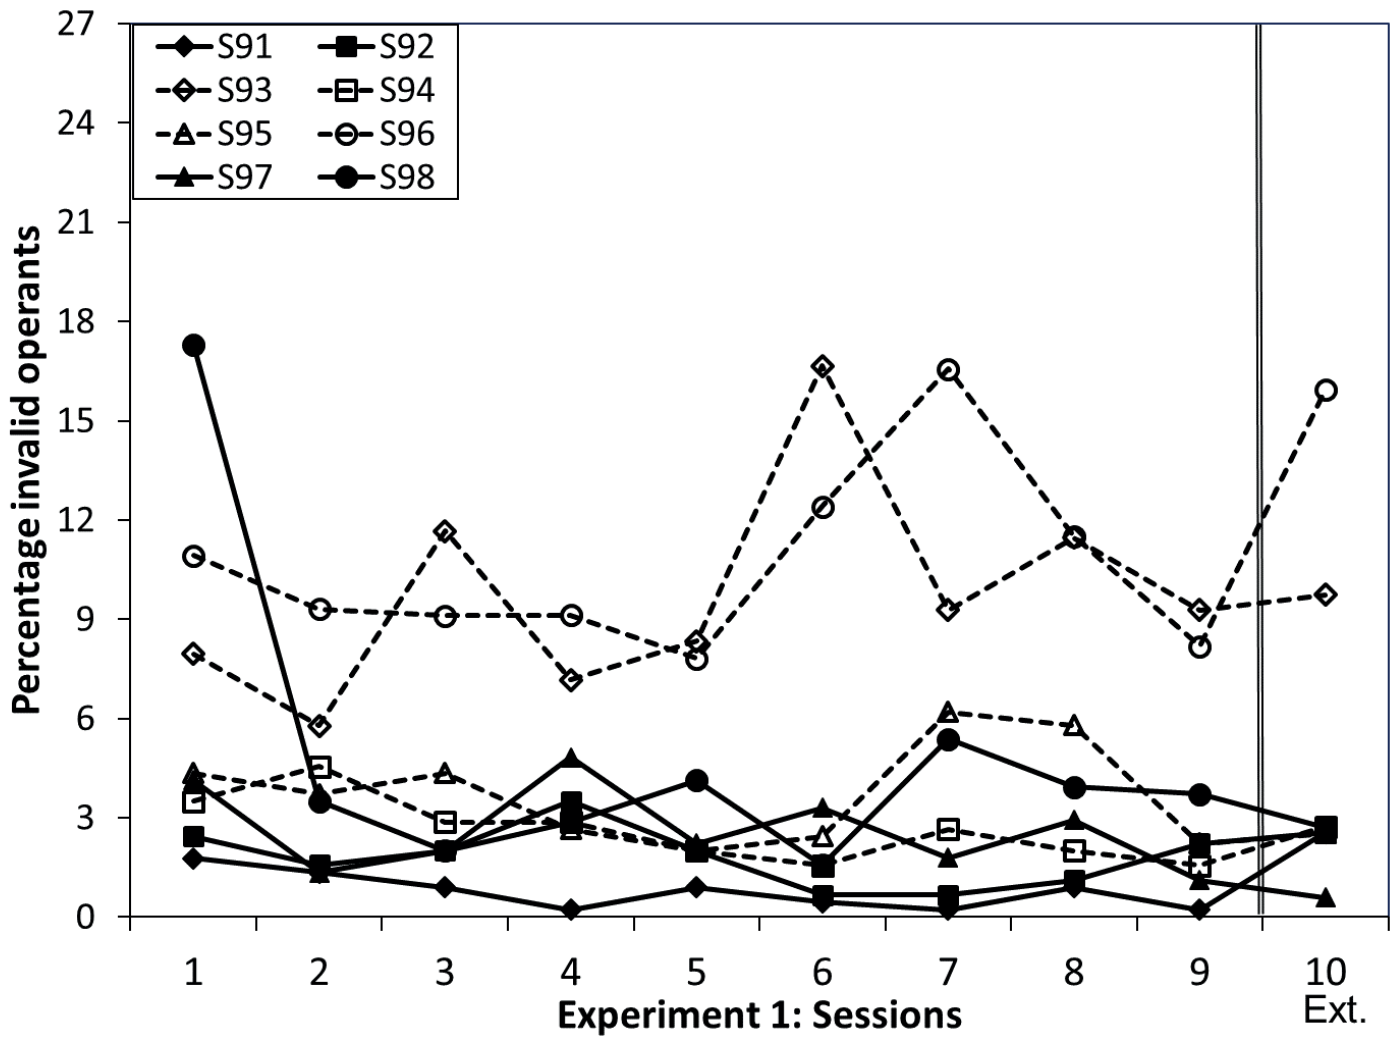


Fig. 4 Percentage of operants emitted by each participant in each session of Experiment 1 that are invalid (not counted toward the variable ratio required for reinforcement).




Fig. 5 Proportion of individual keystrokes (not operants) emitted by each participant in each session of Experiment 1 that are from A. the middle row of three available (the letters ghjk)—the three panels on the left side of the figure, and B. the center section of the keyboard (the letters yuhjbn)—the three panels on the right side of the figure. The top two panels show the proportion of total keystrokes from these spatial areas, the middle two panels show the proportion of the first keystrokes of each operant chosen from these categories, and the bottom two panels show the proportion of last keystrokes of each operant chosen from these categories. The X axis of each panel is the 10 sessions of the Experiment. Each panel also has a gray horizontal line showing where the values would cluster if the participants had shown no bias whatsoever and were indifferent in their selection of letter keys.




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




Fig. 7 Proportion of unique operants emitted by each participant in each session, referred to as the “Unique Value” of Experiment 2.

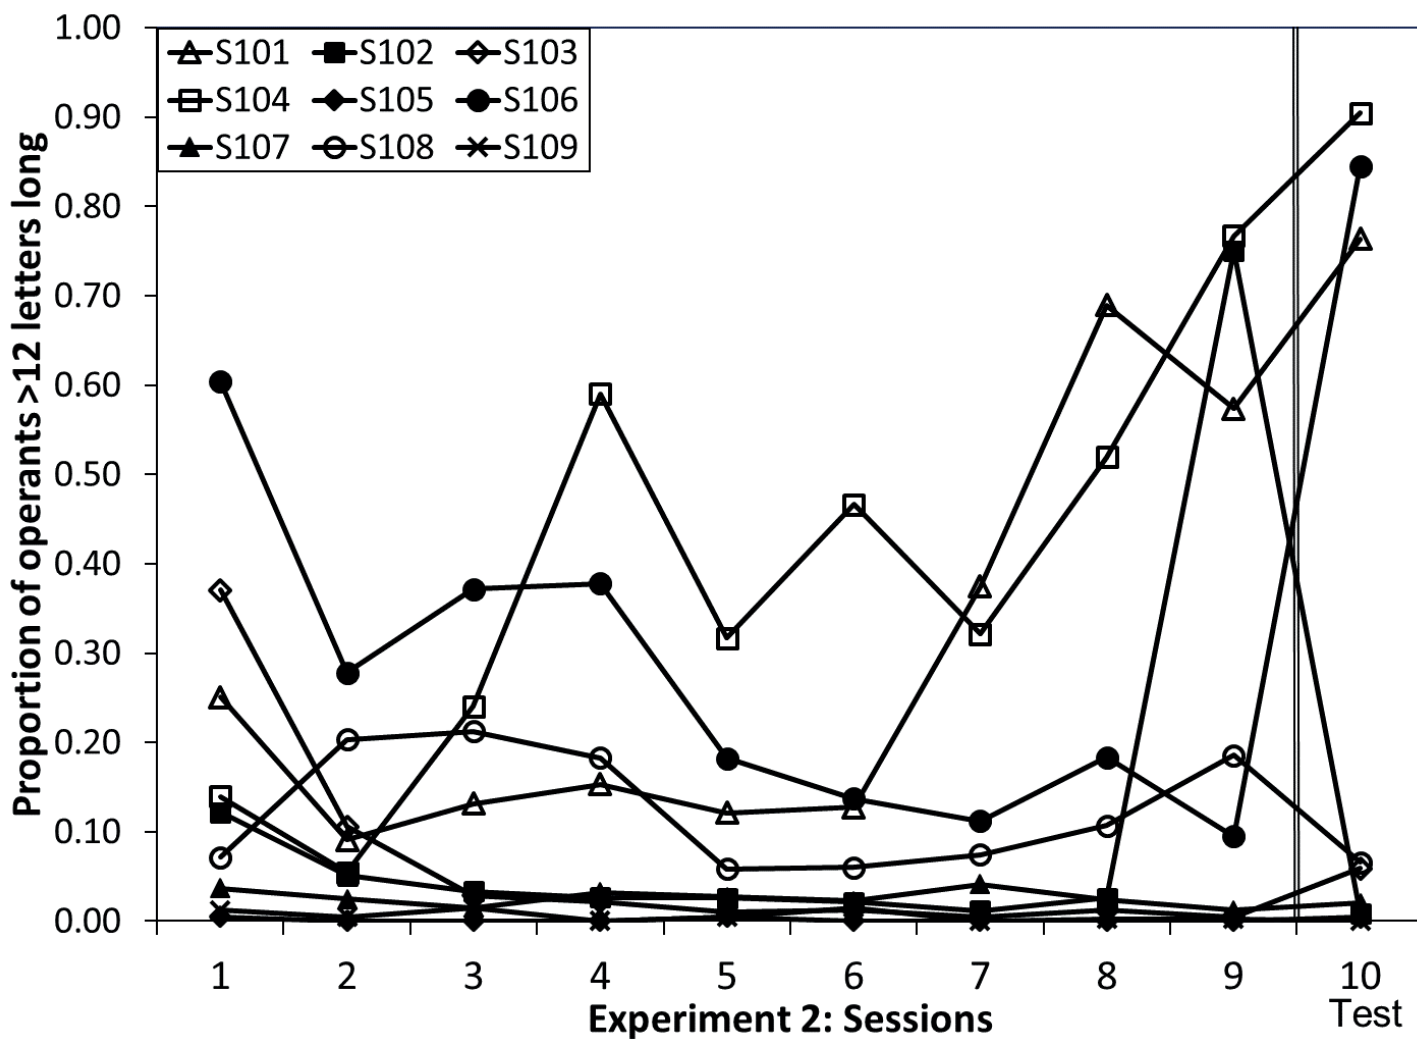


Fig. 8 Proportion of operants more than 12 letters long emitted by each participant in each session of Experiment 2.




Fig. 9 Percentage of invalid operants emitted by each participant in each session of Experiment 2.

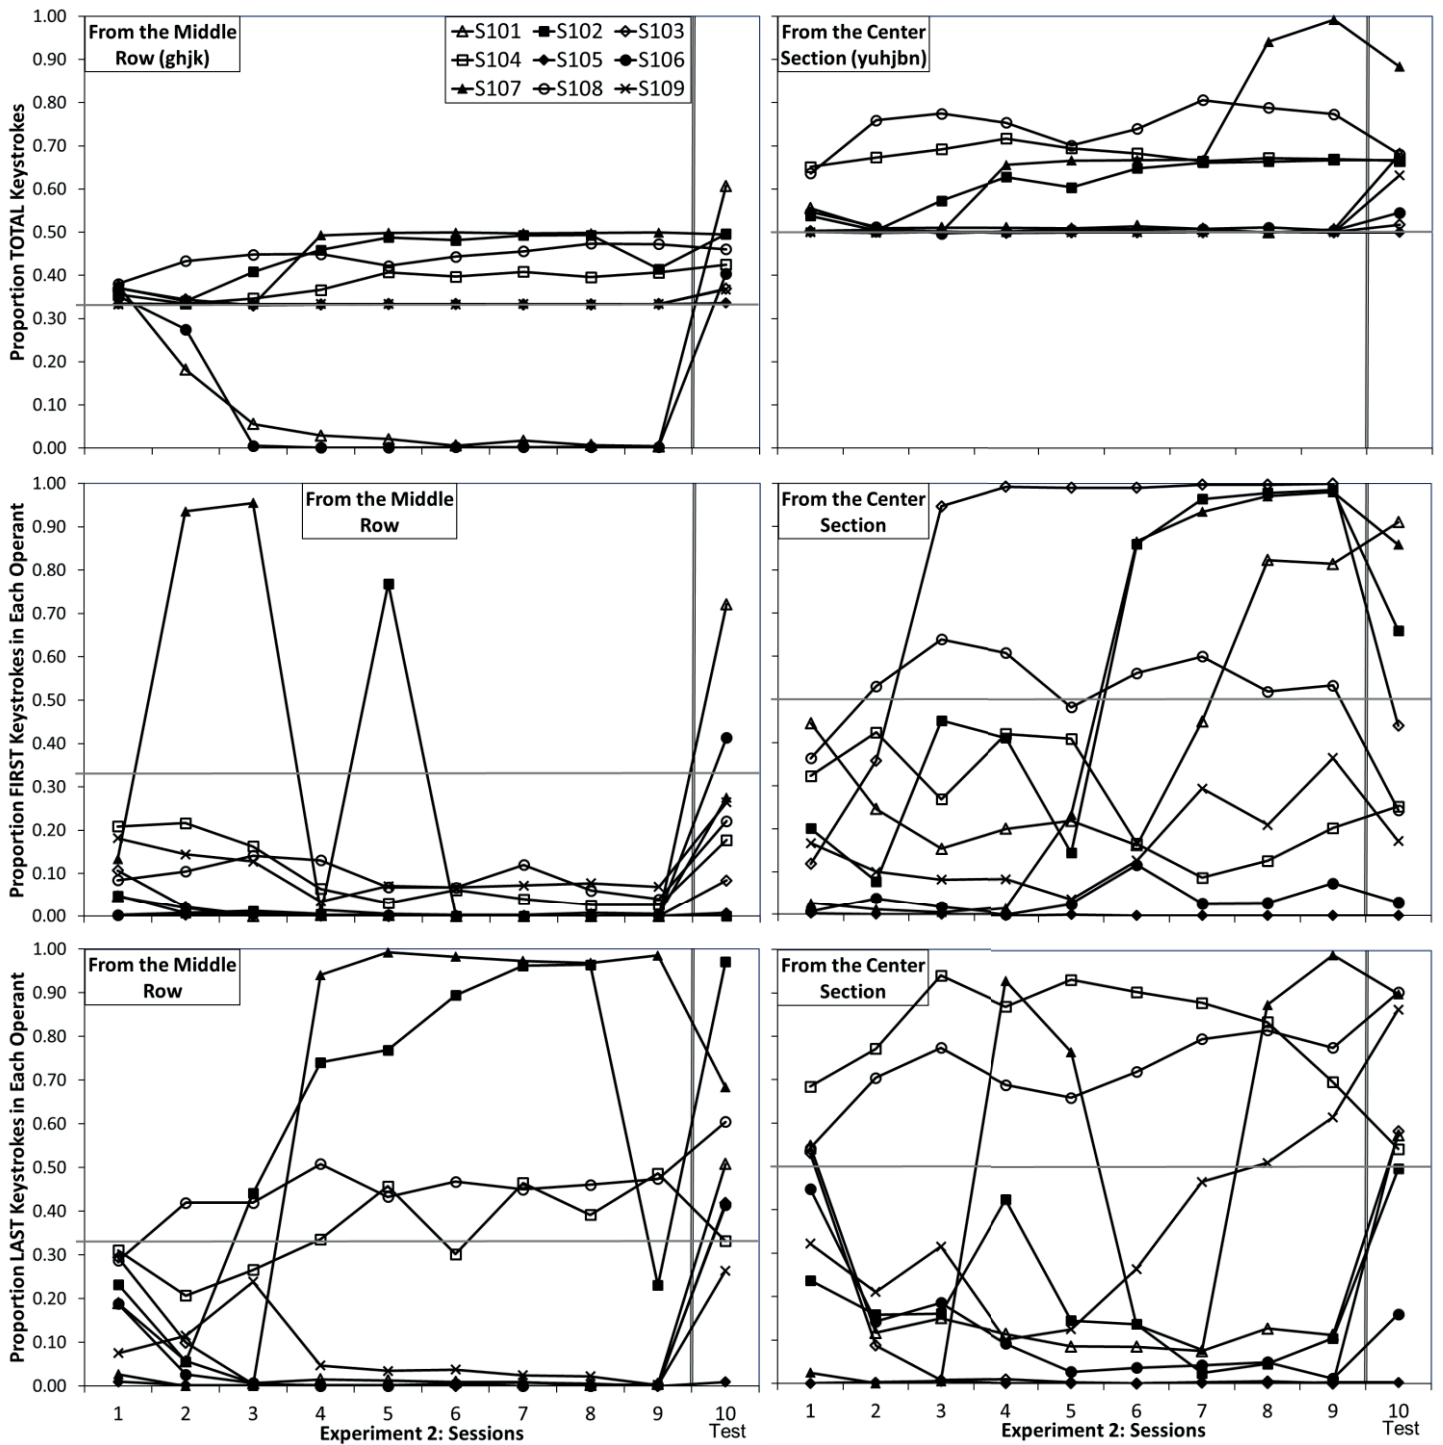


Fig. 10 Proportion of individual keystrokes (not operants) emitted by each participant in each session of Experiment 2 that are from A. the middle row of three, seen on the left side of the figure, and B. the center section of the keyboard, seen on the right side of the figure. The top two panels show the proportion of total keystrokes from these spatial areas, the middle two panels show the proportion of the first keystrokes of each operant chosen from these categories, and the bottom two panels show the proportion of last keystrokes of each operant chosen from these categories.




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