Choice Under Conditions of Uncertainty: Methods for Analyzing the Structure of Behavior

Greg Jensen¹

¹Columbia University

ABSTRACT

All behavior varies, and behavior under a wide range of circumstances displays high levels of variability. This is especially true under many concurrent schedules of reinforcement, despite those schedules not being designed to elicit unpredictable behavior. A generalized matching analysis was performed to measure the effects of reinforcement on 5-item choice under two conditions: Probabilistic concurrent schedules and threshold-based operant variability schedules. Behavior was analyzed in terms of conditional probabilities, incorporating trial-by-trial response dynamics into the model. Performing this analysis meant overcoming a major difficulty: Obtained reinforcement is not a valid independent predictor of behavior the two are causally interlinked. The method of instrumental variable estimation is utilized to overcome the “endogeneity” of reinforcement, which permits unbiased estimation of the causal influence of reinforcement on responding. The analysis revealed a simple relationship between choices made and the distances traveled to make them. Subjects were more willing to travel through the chamber under the Concurrent schedule than under the Variability schedule. As a result, with respect to predicting a subject’s next response, Concurrent schedules elicited higher levels of behavioral variability than did Variability schedules. However, longer-term behavior under the Variability schedules better resembled a steady-state random process.

Keywords: Generalized matching, concurrent schedules, instrumental variable estimation, operant variability, rats.

INTRODUCTION

In the context of operant conditioning, reinforcement is usually associated with a decrease in behavioral variability. If, under some schedule of reinforcement, a particular behavior yields reinforcers, then that behavior is expected to increase in frequency. In most cases, this decreases the entropy of sequences of behavior, which is another way of saying that the organism’s next choice becomes easier to predict.

Unlike Pavlovian schedules (in which reinforcer delivery is independent of behavior), the defining feature of an operant schedule is the cyclical relationship between reinforcers and responses. When a response yields a reinforcer, that response (and responses that resemble it along some relevant set of dimensions) are expected to arise more frequently in the future, in turn yielding more reinforcement. The end-goal of operant learning models is thus to predict, given some history of reinforcement, what the resulting pattern of behavior should be.

Despite the simplicity of this model, the traditional question of “How do reinforcers control behavior?” is deceptively difficult to pin down. For example, although generalized matching models (Baum, 1974; Davison and McCarthy, 1988) provide an overall description of choice proportions as a function of overall obtained reinforcers, they do not describe trial-by-trial probabilities of a response, which are often instead characterized by a high probability of switching (MacDonall, 2009). In general, a subject’s next choice is often highly dependent on its current context, which in turn depends on its previous choices (Jensen,

2014a). This makes behavior difficult to describe succinctly, especially when subjects have many choices at their disposal.

The uncertainty in behavior can be accounted for in part by considering the state the subject was in when the choice was made. For example, knowing where in the operant chamber the subject was when it made the choice tells you a lot about which option it is likely to make next. This approach lends itself to being described in terms of a Markov process, in which behavior is described not in terms of response probabilities, but instead in terms of conditional probabilities (Machado, 1997). This manuscript will show that conditional probability descriptions can fruitfully be combined with a generalized matching model.

Another deceptively difficult problem facing analysts is the cyclical relationship between reinforcers and behavior. As described below, obtained reinforcers may not ordinarily be used as predictors of behavior in regression models of operant behavior because reinforcement is not strictly an “independent variable” with respect to the behavior it is used to predict (Jensen, 2014b). Fortunately, bypassing this difficulty can be accomplished in a straightforward manner using a technique called “instrumental variable estimation” (IVE), originally developed by econometricians (Davidson and MacKinnon, 1993). As shown below, making use of this technique requires no special software or techniques — if you can run a multiple regression, you can use IVE to obtain unbiased estimates of reinforcement’s effect on behavior.

Variability in Behavior and Models

Whether behavior is described in terms of overall proportions or in terms of conditional probabilities, every model of behavior has variability built into it in some fashion. At a minimum, any model that falls short of perfect prediction must have residual error. This differentiates learning models from those in classical economics, which generally focus on normative behavior. A behavior analyst recognizes that, very often, subjects will make non-optimal choices, and the aim is to understand *why* choices are made, rather than trying to calculate which choices *should have* been made. Furthermore, there are many dimensions along which behavior might vary that are not considered by the schedule of reinforcement.

In general, those dimensions of behavior that are not refined by operant conditioning display considerable variability even after a behavior is well-learned (Stokes and Balsam, 1991). For example, although subjects adjust their proportions of response to the various operanda made available under concurrent schedules of reinforcement, the response-by-response pattern of their choices nevertheless display high levels of entropy (Boren et al., 1978; Odum et al., 2006; Racey et al., 2011; Kono, 2017). The variability of behavior under concurrent schedules is interesting precisely because it is unnecessary: In many studies, subjects could have received similar rates of reinforcement with much more stereotyped strategies (e.g. involving higher rates of repetition). Variation also appears when an animal has difficulty acquiring a given schedule of reinforcement (Machado, 1992). In the absence of feedback specifically constraining behavior, it appears as though variation should be expected by default.

Proposals for why variability persists have been presented at several different levels of analysis. On the one hand, maintaining a certain level of behavioral variability often confers a strategic advantage. For example, if an organism forages in an unpredictable way, it is more difficult for predators to anticipate where to lie in wait to ambush them. Multiple competitors can all be driven to variable behavior in order to avoid being taken advantage of by one another, as demonstrated by game theory’s Nash equilibria (Glimcher, 2005). Outside of competitive scenarios, variability is beneficial if it facilitates the discovery of new opportunities or new behaviors that yield better outcomes than those in an organism’s repertoire (Balsam et al., 1998; Stokes et al., 2008). This has, in some cases, been characterized as “information foraging,” because behaviors that appear materially counterproductive can simultaneously be highly informative (Manohar and Husain, 2013; Glaze et al., 2018).

Another perspective, however, is that much of behavior’s apparent variability is induced by variation

in the environment. Not only are most environments dynamic, but organisms themselves greatly increase that dynamism by moving through the environment, leading to a constantly changing stream of proximal stimulation to the senses as their vantage point changes. Braitenberg (1986) demonstrated that this agent-environment interaction can lead to surprisingly complex and unpredictable behaviors even in extremely simple agents navigating environments with only a handful of cues. This possibility is consistent with the theoretical skepticism that has traditionally been a part of behavior analysis: it can be misleading to describe behavior in high level terms (e.g. as “strategic”) when the possibility exists that the underlying mechanism is relatively simple and inflexible.

To better understand the character of behavioral variability, several lines of research have sought to induce variability directly. Under the framework of “operant variability” (Neuringer, 2002), the measurable entropy of a response sequence can be thought of as a dimension of behavior, and reinforced accordingly. Several varieties of reinforcement schedule have been designed with this aim in mind, and have been shown to reliably increase the entropy of response sequences relative to control conditions. The two most popular methods in recent years have been lag procedures, which reinforce behavior that differ from the last n observed behaviors (e.g. Ward et al., 2008), and threshold procedures, which reinforce behaviors if their recent observed frequency is less than some target threshold (e.g. Neuringer et al., 2001).

However, disagreement remains over the theoretical implications of these results. Machado and Tonneau (2012) propose that “operant variability schedules” like those cited above merely provide differential reinforcement to those behaviors that have not appeared recently. Under this “balance hypothesis” (Barba, 2014), the apparent variability of behavior would not be the result of an intrinsic random process, but would instead be a product of changing reinforcement provided over time. The behavior would thus be “steered” by the schedules of reinforcement, and because the schedules themselves award reinforcers in an unpredictable fashion, the resulting behavior displays high entropy.

At the heart of this debate is the question, “What controls behavioral variability?” That is, by what mechanism, and in relation to what feedback from the environment, does behavioral variability emerge? Furthermore, does the incidental variability observed under concurrent schedules differ in kind from the variability that arises from attempts to reinforce it directly? These are difficult questions to answer, in part because most traditional analyses of behavior have focused on the means of their models, without considering the residual variability to be an object of statistical interest as anything other than an error term.

Direct comparisons are also complicated by a lack of standardized experimental procedures. Each study differs from every other study in a wide range of particulars, from the implementation of the reinforcement schedule to the design of the operant chamber. While many studies include a control condition of some kind, those control conditions usually place no requirement on what behavior the animal performs, so long as it continues to respond. This may not be strictly comparable to a scenario in which reinforcers are informative cues.

Models of Behavior and Their Variability

In this study, rats experienced both a concurrent schedule of reinforcement and a threshold-based variability-reinforcing schedule, within-subject and counterbalanced. Both schedules were experienced in the same operant chambers, and the respective schedules were designed to elicit similar distributions of responses across the operanda. This provides an opportunity to analyze behavior under both schedules using a common approach.

If concurrent schedules engendered a distinct flavor of variable behavior from schedules that seek to reinforce variability directly, then it suggests that the variability routinely observed across all behavior need not arise from some common mechanism. Instead, it may simply be an error term that manifests

differently from one task to the next. This would raise theoretical concerns about whether one could make a systematic study of behavioral variability as a cohesive phenomenon. If, however, subjects undertake both tasks similarly, responding similarly to feedback as it becomes available, then the case for a common mechanism of behavioral variability can still be made. Furthermore, such a comparison could provide clues as to the character of that mechanism.

METHODS

Subjects

Twenty 35-week-old male Long-Evans rats had been previously trained to press the two levers and push the three keys present in an operant chamber to obtain food pellets. Beyond that training, the rats were experimentally naïve. All procedures were reviewed and approved by Reed College IACUC.

Apparatus

Each of 10 identical chambers measured 28 cm wide, 27 cm deep, and 29 cm tall (Figure 1). The chambers contained five operanda: three inset keys (3 cm diameter, 9 cm above the floor) on the left wall and two levers (5 cm wide, 1.5 cm deep, 1 cm thick, 6 cm above the floor) on the right wall. These are identified as the Left Key (henceforth LK), Center Key (CK), Right Key (RK), Left Lever (LL), and Right Lever (RL). Each key had the option to be illuminated from behind, and each lever had a light mounted on the wall directly above it. These lights will henceforth be referred to as “operandum lights.” Food pellets (45 mg) were delivered to a tray (3 cm diameter, 2.5 cm above the floor) on the right wall, located between the levers. Water was continuously available from a bottle 18 cm above the chamber floor and 3 cm forward from the back wall. The front door, back wall, and ceiling were composed of clear Plexiglas, while the left and right walls housing the operanda were made of aluminum. The floor consisted of evenly-spaced steel rods. All mechanical components of the chambers were manufactured by Gerbrands. Chambers were housed within sound-dampening enclosures, and a one-way mirror allowed the chamber to be viewed through the Plexiglas door. An Apple eMac computer running Macintosh OS 9 controlled each chamber, running programs written in TrueBasic.

Procedure

Each of the 20 rats chose among five continuously available operanda under two schedules: A “Concurrent” schedule and a “Variability” schedule. These schedules were presented across 12 separate phases, each lasting 10 sessions. The schedule parameters used in a given phases is hereafter referred to as a “configuration.” The experiment employed eight configurations: Four in the Variability condition (denoted as V1 through V4) and four in the Concurrent condition (denoted as C1 through C4). These configurations were presented in a counterbalanced fashion, with 10 subjects initially experiencing Variability and then later experiencing Concurrent, while the other 10 subjects experienced the reverse. Table 1 specifies the order in which both groups experienced these configurations across the twelve phases. To control for order of experience, half of the subjects experienced the four Variability configurations twice, and half experienced the four Concurrent configurations twice.

Table 2 specifies the schedule parameters in each configuration. The four Variability configurations and the four Concurrent configurations were designed to closely resemble one another in terms of proportions of response. For example, preliminary pilot research led us to expect that V1 and C1 would engendered similar distributions of responses, which in turn facilitated a direct comparison of the schedules.

Unless otherwise noted, the Variability and Concurrent schedules were as similar as possible. Each response was followed by a 0.5 s interresponse time (IRT), during which a brief 2.1 kHz tone was presented,

Phase	Subjects 1-10	Subjects 11-20
1	V1	C1
2	V2	C2
3	V3	C3
4	C1	V1
5	C2	V2
6	C3	V3
7	V1	C1
8	V2	C2
9	V3	C3
10	V4	C4
11	C4	V4
12	V4	C4

Table 1. Schedule configurations for each phase of data collection.

	Probability of Food Setup (Concurrent Schedule)				
Configuration	LK	CK	RK	LL	RL
C1	.022	.082	.018	.173	.034
C2	.034	.022	.082	.018	.173
C3	.173	.018	.034	.082	.022
C4	.06568	.06568	.06568	.06568	.06568
	Target Response Distribution (Variability Schedule)				
Configuration	LK	CK	RK	LL	RL
V1	.105	.236	.088	.424	.147
V2	.147	.105	.236	.088	.424
V3	.424	.088	.147	.236	.105
V4	.2	.2	.2	.2	.2

Table 2. Task parameters for each schedule configuration

Figure 1. Layout of the operant chamber, with the five operanda (Left, Center, & Right Keys, plus Left & Right Levers), and other important components of the apparatus.

the operandum lights were darkened, and additional responses during the IRT were not effective other than to reset the 0.5 s interval. Food delivery was preceded by the response-generated tone, followed by an ascending sequence of tones (2.6 kHz, 2.85 kHz, 3.1 kHz) lasting 1 s in total and concluding with the delivery of a 45 mg food pellet. Other than during IRT and food delivery, responses were effective throughout the session and were not demarcated by trial cues. That is, with the exception of the 0.5-s IRT, both Concurrent and Variability were free-operant schedules. Sessions terminated after 90 min.

The Concurrent Schedule

The concurrent schedule of reinforcement was probabilistic and was evaluated every time a response was made (Lau and Glimcher, 2005; Jensen and Neuringer, 2009). Each response (on any operandum) caused a different random number to be generated for each of the five operanda. If the number generated for a particular operandum was equal to or less than its corresponding pre-specified “setup probability,” then a food pellet set up and awaited collection, to be delivered on the next response to that alternative. Each operandum set up its own pellet, such that a pellet could set up on one, several, or no operanda following any given response. Once a pellet set up, it remained available until a response to the associated operandum resulted in pellet delivery. Only one pellet could be set up on any one operandum, such that the pellets could not ‘stack.’ The setup probabilities for each phase are listed in Table 2.

It is important to distinguish this concurrent schedule from variable ratio (VR) schedules in which there

is a fixed (and non-cumulative) probability of reinforcement when a response is made to an operandum (e.g. Baum and Aparicio, 1999). Instead, as in variable interval (VI) schedules, pellets could set up on an operandum even in the absence of a response to that particular operandum. Jensen and Neuringer (2008) distinguish between these two types of schedule in terms of their probability of ‘reinforcer hold;’ the present experiment has a hold probability of 1.0, as is the case in most VI schedules, whereas traditional VR schedules typically have a hold probability of 0.0.

Under these concurrent schedules, food delivery did not depend upon response sequences. All that was necessary was that food had set up and that a response was made to the associated operandum. Changeover delays (Silberberg and Fantino, 1970) were not employed. Unless noted, the schedule was identical to that used by Jensen and Neuringer (2009).

The Variability Schedule

The Variability schedule was intended to reinforce stochastic-like behavior that closely approximated a target distribution. This was accomplished using thresholds for response sequences: If the conditional probability of a given sequence occurred less often than its assigned threshold, it was reinforced. Subjects maximized their obtained rate of reinforcement if their behavior took the form of a stationary Markov process with conditional probabilities identical to those defined by the schedule.

Different phases of the experiment used different target distributions, each corresponding to a distinct stationary Markov process based on the three most recent responses. The target probabilities for individual responses are listed in Table 2. Since most variability schedules in the literature (see Neuringer, 2002; Neuringer and Jensen, 2013, for review) reinforce behaviors that approach equiprobability, the V4 configuration was analogous to a traditional variability schedule.

The schedule examined a moving window of the last three responses, which (given five alternatives) could be any of the 125 possible response sequences. For example, given a response order LK, LK, LL, RK, RK, the first triad would be LK-LK-LL, followed by LK-LL-RK, followed by LL-RK-RK, and so on.

A memory array tracked the relative frequency of these 125 response triads and delivered a reinforcer whenever the most recent response triad was emitted less frequently than its corresponding threshold. Each triad’s threshold was calculated separately, consisting of a conditional probability, plus a small “pad” equal to .0032. The conditional probability was calculated from each configuration’s target distribution. For example, if the target probability for LK was .147, the threshold for the triad LK-LK-LK would be $(.147 \times .147 \times .147) + .0032 = .0064$. Thus, if the triad LK-LK-LK occurred with a probability of less than .0064 in the memory array, it would be reinforced. The pad made the schedule more lenient, to ensure that food was received with similar frequency under both the Variability and Concurrent schedules.

In order for the schedule to be responsive to changing patterns of behavior, it was important that the schedule base reinforcement primarily on recent responses. To achieve this effect, the memory array that tracked past response triads used an “amnesia” coefficient to give recent responses greater weight than less recent ones (Denny and Neuringer, 1998). The emission of a triad increased the value associated with that triad in the memory array by one; then, an algorithm moved the distribution in memory multiplicatively toward the current target distribution. The result of this algorithm was that responses occurring farther back in time contributed less to the memory array, making the schedule differentially sensitive to recent responses. In particular, this prevented behavior early in the phase (when subjects poorly approximated the target distribution) from reducing reinforcement later in the phase.

In overview, a response was followed by food only if it terminated a triad of responses whose relative frequency (with amnesia applied) was less than a specified threshold. This Variability schedule thereby reinforced patterns of response that were stochastically emitted according to a distribution that was similar to the schedule’s prescribed target distribution.

ANALYSIS

One of the most fundamental objectives of behavior analysis is to “describe what happened.” To this end, there is a long tradition of reporting the results of experiments in terms of descriptive statistics. Skinner’s emphasis on cumulative records is perhaps the purest manifestation of this instinct: Just the facts, in as granular a form as can be put on paper.

A more difficult aim, however, is to “explain what happened.” That is, beyond merely describing behavior, an explanation seeks to illuminate the underlying processes and mechanisms that gave rise to that behavior. Ideally, the end goal of this approach is a model of behavior that permits the analyst to make predictions about future behavior, including behavior in novel contexts. This provides the grist for the scientific mill, as model predictions (especially surprising ones) can be compared to real behavior in future experiments.

The exercise of model-building is far more treacherous than mere description, however. On the one hand, specifying a model requires that the analyst make many simplifying assumptions about the behavior and its causes. If these assumptions are made carelessly, essential characteristics of behavior can get “simplified away” and remain invisible to the analyst. On the other hand, a model with sufficient detail to make specific predictions may require the analyst to employ a broader range of mathematical tools than was covered by many psychologists’ formal training. The historical limits of the psychologist’s toolkit force many analysts to use tools that impose too many assumptions and conceal critical aspects of behavior.

In order to build a model that both describes observed behavior and makes specific predictions about future behavior, three of statistical tools will be presented below. Although conceptually distinct from one another, each will help solve a specific problem presented by these data.

Conditional Matching

Generalized matching is a regression approach to describing proportions of response made to a series of operanda in terms of a log-log regression model. One of the assumptions built into this approach is that every response made to an operandum should be treated as equivalent. This assumption is often mistaken, however, particularly when many operanda are concurrently available. Conditional matching, presented below, takes into account not only which operandum was selected, but also what path a subject had to travel to reach that operandum. This turns out to have a profound effect on behavior that is normally invisible under the generalized matching approach.

In its classical form, generalized matching is presented as follows (Baum, 1974):

$$\frac{B_i}{B_j} = \frac{k_i}{k_j} \left(\frac{R_i}{R_j} \right)^s \quad (1)$$

Here, B_i refers to the proportion of choices made to operandum i , while R_i refers to the proportion of reinforcement earned from that operandum. Hereafter, B_i and R_i will only denote proportions, such that $\sum_i B_i = \sum_i R_i = 1$. The parameter k_i is a unitless bias term, indicating the preference for operandum i relative to the other operanda. Because k_i is unitless and relative, we can define $\prod_i k_i = 1$. Finally, s is the “sensitivity to reinforcement” (where a zero indicates no influence of reinforcement whatsoever). Further generalizing to any number of alternatives can be accomplished using “compositional matching” (Jensen, 2014a,b):

$$\frac{B_i}{GM(\mathbf{B})} = k_i \left(\frac{R_i}{GM(\mathbf{R})} \right)^s \quad (2)$$

Here, $GM()$ refers to the geometric mean, and boldface symbols ($\mathbf{B}$, $\mathbf{R}$, etc.) each refer to a vector of all possible operanda. For example, $\mathbf{B}$ is a vector containing the proportion of responses made to every alternative, and $GM(\mathbf{B})$ is the geometric mean of those values. Note that because $\prod_i k_i = 1$, it necessarily follows that $GM(k) = 1$, so the geometric mean of the biases need not be included explicitly in Equation 2.

With five operanda, the model has six parameters (one bias term k_i for each operandum, and a sensitivity s). For model comparison purposes, however, it has only five free parameters because the relative character of k_i removes one degree of freedom from the model. In principle, the parameters of Equation 2 can then be estimated using log-log-transformed linear regression, here written two different ways:

$$\begin{aligned} \log \left(\frac{B_i}{GM(\mathbf{B})} \right) &= \log(k_i) + s \cdot \log \left(\frac{R_i}{GM(\mathbf{R})} \right) + \epsilon_B \\ LR(B_i) &= LR(k_i) + s \cdot LR(R_i) + \epsilon_B \end{aligned} \quad (3)$$

Hereafter, $LR(X)$ will denote $\log \frac{X}{GM(\mathbf{X})}$, such that the shorthand $LR(B_i)$ should be read as “the log-ratio of response proportion B_i , divided by the geometric mean of all response alternatives in the choice vector $\mathbf{B}$.” Throughout this manuscript and the electronic supplement, it is presumed that $\log_e(x)$ (i.e. the natural log) is being used for all transformations.

In the event that one would want to reverse this transformation, the log-ratios can be converted back to proportions using the “softmax” rule (Luce, 1959), again assuming that the natural log has been used for all log-ratio calculations:

$$B_i = \frac{\exp(LR(B_i))}{\sum_x \exp(LR(B_x))} \quad (4)$$

One of the critical simplifications that most matching analyses (compositional or otherwise) have historically made is to presume that every response to a particular operandum is effectively equivalent. If a subject presses the left lever after having traveled across the chamber, that is considered equivalent to pressing the left lever immediately after having already pressed it. Put another way, matching analyses have generally treated the “response” as the moment the operandum’s circuit closes, without considering the sequence of events leading to that choice. Such analyses are routinely described as “molar” (Baum, 2013). There is no reason to think, however, that subjects consider such responses equivalent. If selecting an operandum requires a long journey to reach it, then such a choice is plausibly less likely than if the operandum is nearby. The journey is part of the response.

To take these factors into account, one may extend the compositional matching paradigm into the realm of conditional probabilities. Let $B_{i|j}$ denote the proportion of responses made to operandum i when the subject’s prior location in the chamber is j , and let $\mathbf{B}_j$ denote a vector of responses made to all operanda starting from location j . As such, $\sum \mathbf{B}_j = 1$. Table 3 gives an example of how a table of response counts for each conditional scenario is converted to conditional proportions $B_{i|j}$, and then to log-ratios $LR(B_{i|j})$.

Note that one of the rows in Table 3 is labeled “Rein,” a shorthand for “reinforcement.” Because subjects must travel to the pellet tray to collect reinforcers, choices made following reinforcement can reasonably be expected to display a different preference profile than those following a non-reinforced response. Consequently, the conditional probabilities used in this analysis are always framed in terms of “previous location” rather than “previous response.” The proportion $B_{LK|LK}$, for example, indicates that the

		Conditional Response Counts						Conditional Response Proportions $B_{i j}$						Conditional Log-Ratios $LR(B_{i j})$				
		LK	CK	RK	LL	RL		LK	CK	RK	LL	RL		LK	CK	RK	LL	RL
Previous Location	LK	10	168	24	141	162	LK	.020	.333	.048	.279	.321	LK	-1.83	1.00	-0.95	0.82	0.96
	CK	277	180	245	560	85	CK	.206	.134	.182	.416	.063	CK	0.21	-0.23	0.08	0.91	-0.98
	RK	83	188	14	267	29	RK	.143	.324	.024	.460	.050	RK	0.17	0.99	-1.61	1.34	-0.88
	LL	138	904	376	2002	320	LL	.050	.330	.137	.366	.117	LL	-1.14	0.74	-0.14	0.84	-0.30
	RL	86	325	26	288	39	RL	.113	.425	.034	.377	.051	RL	-0.11	1.22	-1.31	1.10	-0.90
	Rein	44	115	34	1702	33	Rein	.020	.051	.015	.757	.157	Rein	-1.29	-0.33	-1.55	2.37	0.79

Table 3. Demonstration of the conversion between conditional response counts, conditional probabilities, and geomean-centered conditional log-ratios.

previous choice made was LK *and* that the consequence of that choice was no reinforcement (because, if the previous choice had been reinforced, the subject would have traveled to the pellet tray next, making the subsequent choice a $B_{LK|Rein}$ case).

The resulting matching model appears deceptively simple:

$$LR(B_{i|j}) = LR(k_{i|j}) + s \cdot LR(R_{i|j}) + \epsilon_B \quad (5)$$

Although still at its heart a slope-and-intercept model, the present experiment seeks to model six sets of conditional probabilities (one for each operandum, as well as one for choices made after reinforcement), each consisting of five alternatives. Thus, although the model still contains a single slope s , its list of intercept terms $k_{i|j}$ consists of thirty parameter values. Using conditional reinforcers $R_{i|j}$, Equation 5 provides a simple way to predict the estimated conditional log-ratios of responses $LR(\hat{B}_{i|j})$. Equation 4 can, in turn, be used to convert those log-ratios back into estimated proportions $\hat{B}_{i|j}$, which can be used to populate a table of conditional probabilities like that in Table 3.

Instrumental Variable Estimation

The most popular model-building technique in psychology is regression (often, but not always, of the linear variety), and behavior analysis is no different. Indeed, regression models are particularly important to behavior analysis, because they are a powerful way to link the chief subject of the field (namely, behavior) with its primary theoretical cause (namely, reinforcement). Under theories of conditioning, the effect of a reinforcer is to strengthen a particular response class or dimension of behavior. Therefore, in a valid analysis, the parameters relating reinforcement to responding should reflect the degree of influence that reinforcement exerts on responding.

This approach is especially appealing because both reinforcers and responses have meaning across many different schedules of reinforcement. A lever press is topographically the same response under any schedule, just as a food pellet yields a comparable number of calories. In order for behavior analysis to meaningfully generalize beyond the operant chamber, the models of behavior we construct should identify general relationships between behaviors and their consequences. Such models should *not*, however, be applicable only to a narrow subset of contrived experimental preparations.

Despite regression's popularity, it is not widely appreciated among behavior analysts that "obtained reinforcement" is very frequently not a valid predictor for use in regression models. This is because reinforcement often has a reciprocal relationship with responding: distributions of response depend on reinforcement, and distributions of reinforcement also depend on responding. This causal loop is at odds with the fundamental requirement that the predictors in a regression be independent of the outcome (hence their being "IVs"). Consequently, most regression models that predict responses in terms of obtained

reinforcers give rise to biased and misleading parameter estimates. When a predictor is confounded in this way, it is called an “endogenous” predictor. Contrastingly, predictors that are genuinely independent of influence from outcome variables are “exogenous” predictors. Regression models yield unbiased estimates if and only if all predictors in the model are exogenous.

The severity of the bias caused by endogeneity depends on the strength of the outcome’s influence on the endogenous predictor. In concurrent VI-VI schedules, for example, using intervals between reinforcer setups that are much longer than the inter-response intervals results in a distribution of obtained reinforcers that is largely independent of the response distribution. However, in concurrent VR-VR schedules, the distribution of obtained reinforcers is necessarily always strictly dependent on the response distribution, and consequently the distribution of reinforcers cannot be used as an independent variable to predict the response distribution. This is acknowledged directly in the matching literature:

Because there is a fixed molar relation between response rate and reinforcer rate in concurrent-ratio schedules, a sensitivity of 1.0 always occurs trivially in generalized matching analyses on these schedules.

—Davidson and McCarthy (1988), p. 100

While this result is reported as “trivial” (in the sense that it occurs without fail in these schedules), the implication of this result is not, because it is endogeneity that explains why this result occurs. That a regression has the possibility of yielding such a trivial result is an indication that the fundamental assumptions of analysis have been violated. This is especially concerning because other scenarios may exist where the distribution of responses only partially determine the distribution of reinforcers. In such cases, the regression will appear to give a good fit but not necessarily because reinforcers are a strong predictor of responding. The goodness of fit, in such cases, may instead be the result of the bias induced by using endogenous predictors, completely concealing the real relationship between reinforcement and responding. Far from being a curiosity, the uselessness of the matching models in VR-VR experiments is the canary in the coal mine signaling that obtained reinforcement may be a compromised predictor in *any* experimental preparation.

Fortunately, a statistical method called instrumental variable estimation (IVE) can be used to isolate a predictor’s causal influence on an outcome, even when the two variables are causally interlinked (Davidson and MacKinnon, 1993). This is used by introducing an additional predictor, called an “instrument,” which is used to build a model of the compromised variable. Using IVE permits unbiased estimation of obtained reinforcement’s effects on behavior by sidestepping reinforcement’s potential endogeneity. This in turn permits behavior under different schedules (with different degrees of causal connection between reinforcement and responding) to be compared to one another.

Suppose we wish to fit the parameters in Equation 4, but suspect that obtained reinforcer proportions $R_{i|j}$ might be endogenous (thus invalidating the regression parameters). To solve this, a new variable called an “instrument” $Z_{i|j}$ must be introduced. For this method to be effective, the instrument $Z_{i|j}$ must be *strictly exogenous* with respect to both the compromised predictor $R_{i|j}$ and the ultimate outcome of the analysis $B_{i|j}$. In other words, neither $B_{i|j}$ nor $R_{i|j}$ can have any causal influence on $Z_{i|j}$. The instrument is first used, not to predict our ultimate outcomes, but instead to predict the values of our compromised predictor:

$$LR(R_{i|j}) = LR(c_{i|j}) + d \cdot LR(Z_{i|j}) + \epsilon_R \quad (6)$$

This model introduces a series of intercept parameters $c_{i|j}$ (one for each conditional transition) and a slope parameter d . Once values for $c_{i|j}$ and d have been fit, they can be used to obtain an “estimate of

obtained reinforcement” $\hat{R}_{i|j}$. Because this estimate depends strictly on $Z_{i|j}$, and because $Z_{i|j}$ is known to be exogenous, then $\hat{R}_{i|j}$ is a valid predictor for use in a second regression model predicting our ultimate outcome:

$$LR(B_{i|j}) = LR(k_{i|j}) + s \cdot LR(\hat{R}_{i|j}) + \epsilon_B \quad (7)$$

Note how similar Equation 7 is to Equation 5. Despite this similarity, Equation 5 has the potential to be badly compromised by reinforcement’s causal link to responding, whereas Equation 7 is guaranteed to provide an unbiased estimate of all parameters $k_{i|j}$ and s .

To recap: Our objective is to estimate the influence of $R_{i|j}$ on $B_{i|j}$, but a conventional regression won’t do the job because $B_{i|j}$ also influences $R_{i|j}$, rendering it endogenous. To sidestep this difficulty, we need to run two regressions. In the first regression, we use an “instrument” called $Z_{i|j}$ to make a prediction about the values of $R_{i|j}$. Based on the resulting model, we can compute the “instrumental variable estimates” of that predictor, $\hat{R}_{i|j}$. Although both $R_{i|j}$ and $\hat{R}_{i|j}$ appear to estimate the same quantity, the values of $\hat{R}_{i|j}$ are exogenous with respect to our final outcome $B_{i|j}$ because they strictly depend on our instrument $Z_{i|j}$. Because $\hat{R}_{i|j}$ is exogenous, we can use it in a second regression that yields predictions about our outcome of interest $B_{i|j}$. By doing these two regressions in sequence, we can obtain the unbiased estimate that we would have been unable to obtain by fitting $R_{i|j}$ and $B_{i|j}$ directly. For this reason, this particular implementation of IVE is also called “two-stage least squares” regression, since it involves running one regression after another (James and Singh, 1978).

In addition to their being exogenous, the success of IVE depends on choosing one or more instruments that correlate at least moderately well with both the endogenous predictor and with the outcome. The stronger the instrument’s predictive power, the better the estimate $\hat{R}_{i|j}$ will be at standing in as an unbiased proxy for $R_{i|j}$. However, if there is only a weak relationship between the instrument and the other variables in the model, then although Equation 7 will yield an unbiased estimate, it will do so at the cost of a considerable amount of power. In some scenarios, obtaining a good estimate $\hat{R}_{i|j}$ will require using multiple instruments (each with their own slope) in Equation 6. This is fine, so long as Equation 6 is itself a valid regression (including avoiding any further causal entanglements). There is no requirement that Equations 6 and 7 use the same number of predictors.

Historically, IVE has mostly been used in fields that work chiefly with observational data, such as econometrics, sociology, and public health. As an example, an analyst might be interested in studying how risk of mental illness contributes to risk of poverty. This cannot be done with a simple regression model, however, because although mental illness increases the risk of financial hardship, financial hardship also increases the risk of mental illness. Therefore, to estimate the (mental illness)→(poverty) effect, one or more instruments must be identified that can create a predictive model of mental illness (genetic factors, demographics, etc.). The analyst first fits the model (instruments)→(mental illness), then uses that model to calculate the values (estimated mental illness), and finally runs a second regression (estimated mental illness)→(poverty). Provided the instruments selected are exogenous with respect to both mental illness and poverty, the estimated effect in the second regression will be unbiased, *despite* mental illness and poverty causally influencing one another. The importance of such methods has long been appreciated in their respective fields. The 2003 Nobel Memorial Prize in Economics was awarded to Robert Engle and Clive Granger largely for their development of methods to combat endogeneity and measure causal influences in such real-world data (Diebold, 2004; Hendry, 2004).

Information Theory & Behavioral Variability

The conditional matching model described above ultimately yields a table of conditional probabilities. Each row consists of probabilities that are conditional on the subject's prior location in the chamber. Given such a table, it would be convenient to be able to immediately proceed to some metric of how variable the resulting behavior is.

A straightforward candidate is information theory's measure of entropy, typically denoted by H (Jensen et al., 2013):

$$H = - \sum_{\text{all } x} p(x) \cdot \log(p(x)) \quad (8)$$

Here, x represents any discrete event that could happen, and $p(x)$ is its probability. In a long string of events, H represents a measure of how difficult it is to predict the next event x . If the next event is perfectly predictable (e.g. because there's only one possible outcome x), then H has a minimal value of zero. H has no upper bound in general, but does have a bound in any specific case that depends on how many different possible outcomes x could happen. If there are n possible outcomes, and all are equally likely, then any single event x is maximally difficult to predict, and the upper bound for H in that scenario is $H = -n \sum \frac{1}{n} \log\left(\frac{1}{n}\right) = -\log\left(\frac{1}{n}\right)$.

A table of conditional probabilities is not sufficient to compute H . The base rates for a subject's starting positions are also needed. For example, if a subject has a probability of 0.2 of choosing LK following a reinforcer delivery (that is, if $p(\text{LK}|\text{Rein}) = 0.2$), then obtaining a probability usable for computing H requires both the conditional probability $p(\text{LK}|\text{Rein})$ and the baseline probability of reinforcement $p(\text{Rein})$. Given a behavior B and a prior location L , the total entropy associated with the conditional probabilities is computed as follows:

$$H = - \sum_{\text{all } L} p(L) \sum_{\text{all } B} p(B|L) \cdot \log(p(B|L)) \quad (9)$$

Note that this is not the same as computing the entropy associated with $p(B \wedge L)$. Equation 9 assumes that the value of L is known, and is being used to predict the subsequent behavior B . This is a different undertaking than having to select a random response from somewhere within the response sequence and guessing both B and L sight unseen.

RESULTS

Figure 2 plots the observed proportions of response and reinforcement under the Concurrent and Variability schedules for each target distribution. At first blush, both the response proportions and the reinforcer proportions appear quite similar, suggesting that the study's experimental objective of making the two conditions as similar as possible was a success. However, a subtle point worth considering is that although the proportions of reinforcement are quite similar to the target distribution in the Concurrent condition, they are quite different from the target distribution in the Variability condition. Instead, it more closely resembles (and is often more extreme than) the observed response proportions. This is an early warning sign that obtained reinforcement might be a highly endogenous predictor, especially in the Variability condition.

Figure 2 also conceals the considerable degree of response structure that is imposed on subjects by the layout of the chamber. To illuminate this, Figure 3 plots the probability of choosing each response

Figure 2. Population means for proportions of responses made to each of the five operanda (black circles) for the eight different schedule configurations experienced by subjects. Also shown are the obtained proportions of reinforcers (gray diamonds) and the relative target distribution for each schedule (white squares).

alternative, conditional on where the subject was immediately beforehand. Since subjects had to travel to the pellet tray to collect reinforcement (and since they did so every time a reinforcer was delivered), all choices following reinforcement were pooled into the “after reinforcement” category. For example, “LK after Reinforcer” (or, formally, $p(LK|Rein)$) refers to all cases in which a subject chose to press the Left Key immediately after having collected a reinforcer, regardless of what the preceding reinforced response was. Contrastingly, “LK after CK” (or, formally, $p(LK|CK)$) refers to all cases in which a subject chose to press the Left Key immediately after having made an *unreinforced* response to the Center Key. Systematically, our analysis suggested that the animal’s position in the chamber at the onset of the next choice was a critical variable in determining that choice.

The most pronounced effect showcased by Figure 3 is that choices made following reinforcement overwhelmingly favored whichever lever was richer, regardless of which schedule was in effect. However, other effects (particularly those that contrast the two schedules) are harder to see in Figure 3 because of the varying target distributions. There is a hint, for example, that the two schedules differ with respect to repetition of a choice.

To help contrast the Concurrent and Variability schedules, Figure 4 plots the log-ratio of each conditional probability in the Concurrent schedule to its corresponding probability in the Variability schedule for each of the four conditions. Positive numbers indicate choices that are preferred in the Concurrent schedule, whereas negative numbers indicate choices that are preferred in the Variability schedule. Here, the hint provided in Figure 3 becomes easier to see: Repetition (shaded in gray) is consistently avoided in the Concurrent schedule, relative to the Variability schedule. It also appears that, at least some of the time, diagonal moves across the center of the chamber are favored in the Concurrent schedule, relative to the Variability schedule. Although the session-level proportions in Figure 2 suggested very similar distributions of behavior, Figure 4 makes clear that the structure of behavior is different under these two

Figure 3. Probabilities for each of the five response alternatives conditional on each of six possible prior locations the subject could be in the operant chamber. White circles denote means of probabilities under the Concurrent schedule, whereas black diamonds denote means under the Variability schedule. Boxes represent the 95% confidence interval for the mean. All previously reinforced trials are pooled into the “after reinforcer” category (on the right), as subjects are presumed to make their next choice while physically adjacent to the pellet tray. All other conditional choices (e.g. “after CK”) take place after non-reinforced choices.

schedules.

Matching Responses to Reinforcement

The immediate question that arises, from an operant conditioning perspective, is whether this difference in behavior is a function of differential reinforcement. Subjects might avoid repetition in the Concurrent schedule simply because repetition is much less likely to be reinforced. To evaluate this claim, care must be taken because (as described above), obtained reinforcement is a potentially endogenous (and thus compromised) predictor in regression models. One way to evaluate the role of reinforcement without incurring the risk of bias in the regression estimates is to make use of instrumental variable estimation. This, in turn, requires specifying instrumental variables.

For this analysis, the “target distribution” (listed in table 2) for each session was used as the instrumental variable Z_i for each choice. Using equation 5, this was used to estimate the conditional probability of reinforcement R_{ij} expected for each choice alternative. Those estimates were then in turn used to estimate

Figure 4. Log-ratio of conditional probabilities in the Concurrent schedule relative to conditional probability in the Variability schedule. Positive values indicate higher probabilities in the Concurrent schedule, whereas negative values indicate higher probabilities in the Variability schedule.

Figure 5. Conditional matching analysis for choices as a function of reinforcement, made starting from each of six locations in the chamber. Parameters were estimated using IVE, with target distribution as an instrumental variable, and predict proportions of responding relative to estimated obtained reinforcers (Equations 5 and 6). The Concurrent schedule is shown with white circles and the Variability schedule is shown with black diamonds. **(Top Left)** Schematic diagram of the five choices under consideration in each of the six analyses. For example, in the first column, the arrow connecting LK to LL corresponds to choosing the left lever after an unreinforced choice of the left key. **(Bottom Left)** Mean bias, in log units ($LR(k_{i|j})$) for the five choices in each of the six schedules, showing general preference for each choice once rates of reinforcement have been taken into account. Values are in log units, so a bias of 0.0 corresponds to no bias. Boxes represent the 95% credible interval. **(Right)** Sensitivity (s) for each subject under both schedules.

the conditional probabilities of behavior. This approach was used to implement a conditional matching model, modeling the conditional probabilities of choice as a function of reinforcement. These models are provided in the electronic supplement.

Note that the “target distribution” under the Concurrent schedule is not strictly commensurate with that under the Variability schedule. Consequently, the slope d and bias terms $c_{i|j}$ that are used in obtaining the estimates of obtained reinforcement cannot be compared to one another, because they measure fundamentally different units. However, estimates of obtained reinforcement $R_{i|j}$ represent the same quantities for both schedules, and thus, the slope s and the bias terms $k_{i|j}$ that are used to predict response proportions can be compared to one another.

Figure 5 (Left) plots the mean bias across all subjects for each of the 30 conditional probabilities, along with a schematic diagram of which transitions are which in the chamber. This reveals that, contrary to what might be suggested by Figure 3 and 4, subjects consistently favor repetition, *once differential reinforcement is accounted for*. For each response option, the nearest choice is the most favored, whereas choices that

are further away are generally not preferred. Although the strength of these biases appears higher in the Variability schedule than in the Concurrent schedule, the overall pattern among the alternatives is similar.

Figure 5 (Right) plots the sensitivity (s) to obtained reinforcement for each of the subjects in the Concurrent and Variability schedules. In general, sensitivity was higher in the Concurrent schedule, with only a few exceptions. Despite this difference, it is also notable that the sensitivities in general fell within a fairly narrow range of values. In this respect, although a difference was evident, it should still be emphasized that choice as a function of reinforcement was similar in both conditions.

Although the model portrayed in Figure 5 is a sufficient model (in that it describes behavior in a fairly comprehensive manner), it is also exceedingly complex. To describe one subject's behavior under both schedules requires a staggering 102 free parameters: 24 free bias terms $c_{i|j}$, another 24 terms $k_{i|j}$, and 2 slopes d and s , yielding 50 parameters for the Concurrent schedule alone. Double that to include the Variability schedule, and add two error terms σ_R and σ_B to describe the residual error for the regressions. The conditional matching model accounts for an average of 87.1% (SE=0.4%) of the variance of responding, but this isn't surprising given the complexity of the model.

Examining Figure 5 carefully, we can see clues as to how we might simplify the model. In general, once differential reinforcement is accounted for, subjects prefer to remain close to their starting position, and are especially unlikely to cross the center of the chamber. These relationships generally appear symmetric: For example, the bias against going from RK to RL in the Concurrent schedule ($k_{RL|RK} \approx -0.229$) is very close to the bias against going from RL to RK ($k_{RK|RL} \approx -0.320$). The bias is stronger for the Variability schedule, but is again symmetric ($k_{RL|RK} \approx -0.575$ vs. $k_{RK|RL} \approx -0.594$). Because bias grows with distance, and appears symmetric, we can model each bias term as the sum of an operandum press and a series of movements in the chamber. In addition to five components associated with the operands, nine additional components will refer to a 3-by-3 subdivision of the space within the chamber, denoted by their cardinal directions with respect to Figure 1. For example, the upper-left corner of the chamber (where RK is located) will be denoted by NW, for "Northwest." The center of the chamber is denoted by C. These components will subsequently be denoted by m_k , so m_{RK} is the component associated with pressing the Right Key, and m_{NW} is the component associated with being located in the upper-right portion of the chamber.

Table 4 lists the components contributing to each bias term. A "1" indicates that component contributed to the bias associated with the corresponding transition. For example, the bias term associated with the transitions $RK \rightarrow RL$ is $k_{RL|RK}$, and based on the spaces occupied by the subject during that transition, its log-ratio value is given as follows:

$$LR(k_{RL|RK}) = m_{NW} + m_C + m_{SE} + m_{RL} \quad (10)$$

The transition $RK \rightarrow RK$ remains in the NW region of the chamber. Consequently, its bias term $k_{RK|RK}$ depends only on one location, and is constructed as follows:

$$LR(k_{RK|RK}) = m_{NW} + m_{RK} \quad (11)$$

The general rule for constructing a transition bias is to include one term for the operandum selected, and one term for each region of the chamber the subject spends time in. For repetitions, this includes the cell they began in (as in Eq. 11). The full implementation of this spatial model is included in the electronic supplement.

Figure 6 (Left) plots the means of the bias component estimates for both conditions. The five operandum components are positive and convey the relative preferability of each operandum type. The

Bias ($LR(k_{i j})$)	Additive Spatial Bias Components													
	m_{LK}	m_{CK}	m_{RK}	m_{LL}	m_{RL}	m_{NW}	m_N	m_{NE}	m_W	m_C	m_E	m_{SW}	m_S	m_{SE}
LK $\rightarrow$ LK	1	0	0	0	0	0	0	0	0	0	0	1	0	0
LK $\rightarrow$ CK	0	1	0	0	0	0	0	0	1	0	0	1	0	0
LK $\rightarrow$ RK	0	0	1	0	0	1	0	0	1	0	0	1	0	0
LK $\rightarrow$ LL	0	0	0	1	0	0	0	1	0	1	0	1	0	0
LK $\rightarrow$ RL	0	0	0	0	1	0	0	0	0	0	0	1	1	1
CK $\rightarrow$ LK	1	0	0	0	0	0	0	0	1	0	0	1	0	0
CK $\rightarrow$ CK	0	1	0	0	0	0	0	0	1	0	0	0	0	0
CK $\rightarrow$ RK	0	0	1	0	0	1	0	0	1	0	0	0	0	0
CK $\rightarrow$ LL	0	0	0	1	0	0	0	1	1	1	0	0	0	0
CK $\rightarrow$ RL	0	0	0	0	1	0	0	0	1	1	0	0	0	1
RK $\rightarrow$ LK	1	0	0	0	0	1	0	0	1	0	0	1	0	0
RK $\rightarrow$ CK	0	1	0	0	0	1	0	0	1	0	0	0	0	0
RK $\rightarrow$ RK	0	0	1	0	0	1	0	0	0	0	0	0	0	0
RK $\rightarrow$ LL	0	0	0	1	0	1	1	1	0	0	0	0	0	0
RK $\rightarrow$ RL	0	0	0	0	1	1	0	0	0	1	0	0	0	1
LL $\rightarrow$ LK	1	0	0	0	0	0	0	1	0	1	0	1	0	0
LL $\rightarrow$ CK	0	1	0	0	0	0	0	1	1	1	0	0	0	0
LL $\rightarrow$ RK	0	0	1	0	0	1	1	1	0	0	0	0	0	0
LL $\rightarrow$ LL	0	0	0	1	0	0	0	1	0	0	0	0	0	0
LL $\rightarrow$ RL	0	0	0	0	1	0	0	1	0	0	1	0	0	1
RL $\rightarrow$ LK	1	0	0	0	0	0	0	0	0	0	0	1	1	1
RL $\rightarrow$ CK	0	1	0	0	0	0	0	0	1	1	0	0	0	1
RL $\rightarrow$ RK	0	0	1	0	0	1	0	0	0	1	0	0	0	1
RL $\rightarrow$ LL	0	0	0	1	0	0	0	1	0	0	1	0	0	1
RL $\rightarrow$ RL	0	0	0	0	1	0	0	0	0	0	0	0	0	1
Rein $\rightarrow$ LK	1	0	0	0	0	0	0	0	0	1	1	1	0	0
Rein $\rightarrow$ CK	0	1	0	0	0	0	0	0	1	1	1	0	0	0
Rein $\rightarrow$ RK	0	0	1	0	0	1	0	0	0	1	1	0	0	0
Rein $\rightarrow$ LL	0	0	0	1	0	0	0	1	0	0	1	0	0	0
Rein $\rightarrow$ RL	0	0	0	0	1	0	0	0	0	0	1	0	0	1

Table 4. Which additive components to include in order to build each bias term $LR(k_{i|j})$ for the spatial matching model.

Figure 6. The spatial matching analysis, which collapses the 30 conditional bias terms plotted in Figure 5 into five “operandum components” and nine “movement components,” combined as demonstrated by Equations 9 and 10. Sensitivity parameters are effectively identical to those plotted in Figure 5, and so are not shown. **(Left)** Estimates for the mean bias components for the Concurrent (white circles) and Variability (black diamond) schedules. Boxes represent the 95% credible interval. **(Right)** Schematic visualization of the operand chamber, including the means for each component of the bias. Note that the boundaries between the nine regions of the chamber are only approximate, and should not be taken as strict criteria for which region is which.

nine movement components are negative and convey the cost of moving through space. The greater the gap between the operandum components and the movement components, the more the model will avoid moving. Thus, this figure confirms that even when differential reinforcement is accounted for, subjects switch more under the Concurrent schedule (because the cost of movement is small relative to the gains of choosing an operandum), whereas subject switch less under the Variability schedule (because each step through the chamber incurs a greater cumulative cost). The cost of moving through the center is the most extreme cost for both schedules, but is much costlier in the Variability case than it is in the Concurrent case.

Figure 6 (right) situates the means of the operandum and movement components within a scale diagram of the chamber. Drawn in this way, it is clear that, in the Concurrent condition, movement along all the walls of the chamber incur a roughly equal cost (since all movement components along the edge have mean costs between -0.247 and -0.381. Movements near the levers have similarly low costs in the Variability condition, but movement is costlier in the rest of the chamber, especially in the transitions between the levers and the keys. Movement through the center of the chamber is especially costly in both cases.

Introducing the construct of “spatial costs” greatly simplified the model, cutting the number of free parameters roughly in half, to a total of 62 when both Concurrent and Variability are considered. However, there is still a lot of complexity in the model. To simplify the model further, suppose that all moves along an edge have the same cost. Furthermore, suppose that all keys are equally favored with one another, as are both levers. Under this approach, a bias term is constructed by taking the sum of the edge and center moves entailed by a transition. Table 5 lists these counts for a simplified model of bias consisting of just four terms: m_{Key} and m_{Lever} (to account for operandum preference), as well as m_{Edge} and m_{Center} (to account for movement costs). These terms are multiplied by the relevant counts. For example, to construct $k_{\text{RL|RK}}$:

	Simplified Additive Spatial Bias Components			
Bias ($LR(k_{i j})$)	m_{Key}	m_{Lever}	m_{Edge}	m_{Center}
LK $\rightarrow$ LK	1	0	1	0
LK $\rightarrow$ CK	1	0	2	0
LK $\rightarrow$ RK	1	0	3	0
LK $\rightarrow$ LL	0	1	2	1
LK $\rightarrow$ RL	0	1	3	0
CK $\rightarrow$ LK	1	0	2	0
CK $\rightarrow$ CK	1	0	1	0
CK $\rightarrow$ RK	1	0	2	0
CK $\rightarrow$ LL	0	1	2	1
CK $\rightarrow$ RL	0	1	2	1
RK $\rightarrow$ LK	1	0	3	0
RK $\rightarrow$ CK	1	0	2	0
RK $\rightarrow$ RK	1	0	1	0
RK $\rightarrow$ LL	0	1	3	0
RK $\rightarrow$ RL	0	1	2	1
LL $\rightarrow$ LK	1	0	2	1
LL $\rightarrow$ CK	1	0	2	1
LL $\rightarrow$ RK	1	0	3	0
LL $\rightarrow$ LL	0	1	1	0
LL $\rightarrow$ RL	0	1	3	0
RL $\rightarrow$ LK	1	0	3	0
RL $\rightarrow$ CK	1	0	2	1
RL $\rightarrow$ RK	1	0	2	1
RL $\rightarrow$ LL	0	1	3	0
RL $\rightarrow$ RL	0	1	1	0
Rein $\rightarrow$ LK	1	0	2	1
Rein $\rightarrow$ CK	1	0	2	1
Rein $\rightarrow$ RK	1	0	2	1
Rein $\rightarrow$ LL	0	1	2	0
Rein $\rightarrow$ RL	0	1	2	0

Table 5. Which additive components to include, and their respective weight, in order to build each bias term $LR(k_{i|j})$ for the simplified spatial matching model.

Figure 7. The simplified spatial matching analysis, which further collapses the bias components plotted in Figure 6 into four terms, combined as demonstrated by Equations 11 and 12. Sensitivity parameters are effectively identical to those plotted in Figure 5, and so are not shown. **(Left)** Estimates for the mean bias components for the Concurrent (white circles) and Variability (black diamond) schedules. Boxes represent the 95% credible interval. **(Right)** Schematic visualization of the operant chamber, including the means for each component of the bias.

$$LR(k_{RL|RK}) = 2 \cdot m_{\text{Edge}} + m_{\text{Center}} + m_{\text{Lever}} \quad (12)$$

Since this transition consists of two edge moves (SE and NW) the m_{Edge} term is added twice. Similarly, because the repetition of RK consists only of spending time in the NW region of the chamber (and thus one “edge” component), the equation to construct $k_{RK|RK}$ is as follows:

$$LR(k_{RK|RK}) = m_{\text{Edge}} + m_{\text{Key}} \quad (13)$$

Figure 7 (Left) plots the means for the four bias components in the simplified spatial model in each condition. As before, although subjects tended to prefer to stay put in general, the cost of movement was higher for the Variability schedule than it was for the Concurrent schedule. Movement through the center of the chamber was more costly than movement along the edge of the chamber, consistent with Figure 6. Figure 7 (Right) depicts these values within the context of the operant chamber. Overall, the simplified model is indeed much simpler: To explain behavior in both conditions and perform IVE, the simplified spatial model makes use of only 22 free parameters.

Given three different matching analyses, with three different levels of complexity, it is natural to want to compare the efficacy of each model in a way that controls for its complexity. While a linear regression models’ ability to explain the variance of its outcome is measured by R^2 , it offers no protection against overfitting by a complex model. The Bayesian Information Criterion (BIC) controls for the goodness of fit by penalizing models as a function of their number of free parameters q (Schwartz, 1978). However, its calculation is slightly more complicated when using IVE, because the model being fit consists of multiple

	Conditional Matching Model	Spatial Matching Model	Simplified Spatial Matching Model
Mean R^2 ($LR(R_{i j})$) (SD)	0.871 (0.004)	0.651 (0.007)	0.609 (0.007)
Mean R^2 ($LR(B_{i j})$) (SD)	0.859 (0.035)	0.660 (0.008)	0.599 (0.009)
Mean BIC (SD)	-326.56 (13.41)	-190.29 (14.64)	-347.94 (14.74)

Table 6. Goodness of fit and information criteria for model comparison.

simultaneous equations: In this case, one predicting B and one predicting R . Consequently, the BIC for simultaneous linear equations was computer as follows, per Gorobets (2005):

$$\text{BIC} = n \cdot \log(\det(\Sigma)) + \log(n) \left(\frac{m(m+1)}{2} + \sum_{i=1}^m q_i \right) \quad (14)$$

$$\det(\Sigma) = \text{VAR}(\epsilon_B) \cdot \text{VAR}(\epsilon_R) - \text{COV}(\epsilon_B, \epsilon_R)^2$$

Here, n is the sample size, m is the number of simultaneous equations, and each q_i is the number of parameters used by each of the equations. Σ is a covariance matrix describing how the residuals of the m simultaneous models covary with one another, and $\det(\Sigma)$ is the determinant of that matrix.

The model with the highest R^2 does the best job retrodicting the observed data. Thus, in order to obtain the most detailed description of what subjects did, one should favor the model that explains the largest share of the variance. However, being able to describe the current data does not necessarily mean that a model will do well in predicting behavior under future circumstances. The model with the *lowest* BIC score is the model for which the evidence is strongest, having taken its complexity into account.

Table 6 compares the mean R^2 for both the predictions of R and B for each model, as well as its overall BIC. From this, although it is clear the conditional matching model does an excellent job explaining the variance in both the obtained reinforcers and in the behavior, it only comes in second at the level of model comparison. The transition to the spatial model results in a substantial reduction in the variance explained, and despite its lower complexity, it is judged the worst of the three models. Surprisingly, despite it being a further simplification of the spatial model, the simplified spatial model is judged to provide the best overall fit, because its dramatic reduction in complexity offsets its lower variance explained.

At the outset of this process, the analysis of choice among five response alternatives appeared bewilderingly complex, particularly because the endogeneity of obtained reinforcers compromised their usefulness as a predictor. However, it turns out that a sensible description of the behavior can be achieved by modeling the costs of movement through the chamber directly. According to this analysis, the Concurrent schedule induces subjects to move more freely through the chamber than does the Variability schedule, even after having corrected for differential reinforcement.

Variability of Behavior

One of the clear takeaways of the matching analyses above is that subjects were less prone to moving around the chamber under the Variability schedule than they were under the Concurrent schedule. In a trivial sense, therefore, behavior under the Variability schedule should be easier to predict, since there are many ways to move around the chamber, but only one way to stay put.

Figure 8. Entropy of response sequences under the Concurrent schedule (white circles) and the Variability schedule (black diamonds), based on Equation 8. These values presume that the subject's current location in the chamber is known, and represent the overall unpredictability of the subject's next response.

Figure 8 confirms this suspicion. The entropies of each subject's observed conditional probabilities of behavior were calculated according to Equation 9, and those entropies were higher under the Concurrent schedule for most subjects most of the time. From the perspective of "predicting the next response," behavior under the Variability schedule was indeed easier to predict, and in that sense was less variable.

This is not, however, the only way of looking at whether behavior is unpredictable. Insofar as both schedules of reinforcement have something in common, it is that subjects must approximate particular distributions of behavior over time to maintain a high rate of reinforcement. If they are doing so randomly, then there will be stretches of time in which certain operanda won't be selected simply by happenstance. However, if subjects are *systematically* trying to maintain a particular distribution, then as more time passes without an operandum being selected, the greater the pressure will be to begin selecting it again. In other words, if behavior is random, then the chance of making a choice should not depend on how much time has passed since that choice was last made. This premise is at the core of Barba's "balance hypothesis" of operant variability.

Formally, behavior is random if $p(\text{choice}|t) = p(\text{choice})$, given some number of trials t since that choice was last made. If, on the other hand, subjects are maintaining their current distribution of responses by steering back toward items they haven't visited recently, then $p(\text{choice}|t) > p(\text{choice})$ as t becomes large. By the same logic, when the log-ratio of these two probabilities are examined, then random behaviors are expected to yield a value of $\log\left(\frac{p(\text{choice}|t)}{p(\text{choice})}\right) = 0.0$, whereas if Barba's "balance hypothesis" is correct, then $\log\left(\frac{p(\text{choice}|t)}{p(\text{choice})}\right) > 0.0$. By this metric, we can evaluate whether behavior depends on events that occurred many trials before the current response.

Figure 9 plots the log-ratio of the probability of making each choice given t trials since its last selection, relative to the overall rate at which each choice was made. Here, again, we see the pronounced proclivity for repetition (i.e. when $t = 1$) in the Variability condition, and the corresponding avoidance of repetition in the Concurrent condition. What we also see, however, is that the conditional probability of making a choice tends to grow as time passes, and in the Concurrent case tends to consistently grow larger than the baseline rate. Figure 10 helps to make this effect clear by presenting the mean log-ratios for each operandum across the four schedules. In the Concurrent schedule, subjects tend to steer back toward alternatives they have avoided if they have been away from them too long. In the Variability schedule,

Figure 9. Log-ratios comparing the probability of making a choice after t trials of not having made that choice, relative to the probability of making that choice overall. White circles represent means for the Concurrent condition, whereas black diamonds represent means for the Variability condition. Positive values correspond to conditional probabilities that exceed baseline levels, whereas negative values correspond to conditional probabilities that fall below baseline levels. The dark shaded region corresponds to the 80% credible interval of the mean, and light shaded regions represent the 99% credible interval.

however, subjects tend instead to steer back to the global probability. Knowing what a subject did 7 or 8 trials previously will usually help you predict what the subject will do next under the Concurrent schedule, but not what they will do under the Variability schedule.

In this sense, although the Variability schedule is meaningfully less variable with respect to short-term response-by-response transitions (as measured by the entropy of its conditional probabilities), it is *more* variable with respect to resembling a random walk in the long term. Barba’s balance hypothesis is indeed a compelling account of how response distributions are maintained under Concurrent schedules, but falls short in explaining behavior under Variability schedules.

DISCUSSION

In this study, subjects chose from among five response alternatives under two different schedules of reinforcement: The “Concurrent” schedule (that delivered reward probabilistically) and the “Variability” schedule (that rewarded behavior when it roughly approximated a particular stochastic distribution, as

Figure 10. Averages of the log-ratios of probabilities of making each choice conditional on t trials since a choice was last made, relative to overall probabilities of making that choice.

evaluated by frequency thresholds). Although motivated by different objectives, both schedules presented subjects with uncertainty about the future location of reinforcers. Furthermore, care was taken to make the schedules as similar as possible in every way besides their internal logic. Consequently, a comparison of behavior in these two schedules provides an opportunity to ask questions about how behavior responds to feedback in general terms, rather than merely within the scope of a single experimental paradigm.

Firstly, subjects were clearly responsive to the changing schedules, shifting their behavior to favor alternatives that yielded many reinforcers and avoiding those that yielded few (Figure 2). However, it was also clear that movement through the cage profoundly influenced the preferences animals had for their next choice. This was first evident in the overwhelming preference for levers on trials that immediately followed the delivery of a reinforcer to a pellet tray that was much closer to the levers than it was to the keys (Figure 3). This raised the further question: Even if subjects were reasonably good at approximating the target distributions on average, where their trial-by-trial choices impacted heavily by movement through the chamber, and how did those impacts differ between the two reinforcement schedules?

Figure 4 helps to clarify one major way in which the Concurrent and Variability schedules differed: Selecting the same operandum several times in a row was much less likely to occur under Concurrent schedules than under Variability schedules. By contrast, subjects appeared more willing to make long trips between operandum (e.g. LK→LL) under the Concurrent schedule than they were under the Variability schedule. Broadly speaking, the Concurrent schedule appeared to be more favorable toward subjects moving through the chamber than was the Variability schedule.

The influence of reinforcement was measured using the conditional matching model, which made use of instrumental variable estimation (IVE) to obtain an unbiased estimate. This analysis revealed that subjects generally sought to minimize their travel time (Figure 5), with repeats being most favored, *relative to the rate of reinforcement*. In other words, although subjects generally preferred to repeat in both schedules, repetition was differentially reinforced in the two schedules, resulting in different apparent preferences in Figure 3. Meanwhile, the bias terms for the Variability schedule were more exaggerated (i.e. further from zero), and the sensitivity to reinforcement was lower. Collectively, the patterns among the biases suggested that a simpler model could be achieved. As shown in Figure 6, the bias terms could be decomposed into the relative contributions of operandum selection and movement through the chamber by dividing the chamber into nine zones and estimating the “cost” of spending time in each zone over the

course of a response's execution. This model was further simplified to explain the bias among transitions entirely in terms of four factors: Whether responses were made to a key or a lever, and how many edge or center zones the subject had moved to in order to reach that operandum (Figure 7). This last, simplified model proved to be best supported by the available evidence, according to the BIC (Table 6).

A comparison of the variability of behavior under both schedules yielded mixed results. Consistent with the suggestion of the spatial matching analysis, subjects tended to move around the operant chamber less under the Variability schedule, and thus displayed behavior that was easier to predict on a trial-by-trial basis (Figure 8). However, when looking at behavior over the longer term of tens of responses, behavior under the Concurrent schedule showed signs of “balancing” itself over time (with a response becoming more likely the longer the interval of time since its last selection; Figures 9 and 10). This balancing over the long term was not evident in the Variability schedule.

Follow the Reinforcers

Differences emerged between the Concurrent and Variability schedules, but behavior was nevertheless broadly similar under both schedules. Subjects consistently took both travel time and relative rates of reinforcement into account when deciding on their next choice. The result was behavior that responded efficiently to changing schedules. Although there were apparent differences in behavior under the two schedules, these need not be explained in high-level terms, such as subjects “recognizing” those differences and “strategizing” accordingly. A large proportion of the variance in behavior can be accounted for as a function of proportion of reinforcers received and distance traveled to obtain them.

Although the spatial matching models explained less of the variance in behavior, their general success helps to explain how subjects approach the complexity of a five-operandum task. Keeping track of the reinforcers earned by performing every possible response transition constitutes a very considerable memory load. Even though the spatial analysis has the flavor of a traditional operant analysis of behavior, it is important to recognize that the parameters that emerge effectively constitute a map of the operant chamber (or at least, of the features of the chamber that are relevant to the schedules of reinforcement). Inasmuch as this model does a good job describing the behavior of subjects, it also motivates research into how such tasks are encoded by the brain. Although behavior analysts often contrast cognitive proposals (such as “cognitive maps”) with “simpler” learning processes, analyses like the simplified spatial matching model (Figure 7) present an example of how encoding the spatial characteristics of the task directly into a learning model both yields a map of the chamber *and* is simpler than a naïve operant model.

Despite the similarity of behavior in the two schedules, subjects were more prone to move around the chamber in the Concurrent schedule, and less willing to do so in the Variability condition, even after accounting for differential reinforcement. This suggests that subjects were sensitive to some other factor intrinsic to the schedules that was not captured by probabilities conditional only on the subject's last location in the chamber. Figures 9 and 10 give a clue as to what this might be: The time since an operandum was last visited.

In general, concurrent schedules of reinforcement routinely lead to frequent switching behavior, and for good reason: The conditional probability of reinforcement at a particular operandum depends, in many cases, on how long it has been since that operandum was visited, and *not* on how long it has been since it yielded a reward. Our schedule, for example, was trial-driven, and this makes the probability of reinforcement easy to calculate (Jensen and Neuringer, 2008):

$$p(\text{Rein}|X, t = n) = 1 - (1 - p(\text{Rein}|X, t = 1))^n \quad (15)$$

Here, t is the number of trials since the last selection of an operandum X . When $t = 1$, the probability

is equal to the probability of food setup, as listed in Table 2. For example, $p(\text{Rein}|\text{LK}, t = 1) = 0.022$ in configuration C1. With such a low probability of reinforcement, it makes sense that subjects would avoid it. However, after 20 trials, $p(\text{Rein}|\text{LK}, t = 20) = 1 - (1 - 0.022)^{20} = 0.359$, a 16-fold increase. This probability depends strictly on the time since the last visit to an operandum, and not on whether any given response was reinforced. Therefore, a subject with perfect memory and knowledge of the schedule will almost always switch between operandum, deterministically favoring whichever has accrued the highest cumulative probability of reinforcement. The low cost of movement for subjects in this schedule, as well as the lookback analysis in Figure 9 and 10, suggest that although subjects are not deterministic in their responding, they are nevertheless sensitive to the time spent away from an operandum.

That concurrent schedules actively steer behavior by gradual changes in conditional reinforcement has long been appreciated. Killeen (2015) describes concurrent VI schedules as “complex feedback machines designed to moderate the allocation of behavior among alternatives” (p. 79), and goes on to describe those dynamics in great detail:

There is negative feedback inherent in most concurrent schedules — the longer the subject spends on one alternative, the more likely the other alternative is to pay off. Even though one payoff may be of less utility to the subject, that lesser utility at some point is balanced by its greater immediacy or probability. Concurrent schedules thus constitute dynamic delay/probability discounting experiments in miniature. (p. 81)

Consequently, it should come as no surprise that Concurrent schedules appear to yield evidence for the balance hypothesis: They *actively encourage* such long-term steering of behavior back toward some target distribution, and insofar as subjects don’t match the target distribution precisely, this can be chalked up to a combination of their insensitivity to ratios of reinforcement and the behavioral biases induced by differing operandum and their spatial configuration.

Unlike the Concurrent schedule, however, behavior under the Variability schedule is not compelled by reinforcement to display constant switching. Indeed, doing so would reduce the rate of reinforcement, since excessive switching would lead to an overrepresentation of switch-heavy sequences, pushing their frequency above the target threshold and preventing those sequences from being reinforced. In fact, when variability schedules combine sequence-based thresholding with a moving window, it can even lead to higher rates of reinforcement for repetition than for other sequences. This is because repetitive sequences (“AAA”) can be visited more rapidly than nonrepetitive sequences (“ABC”); the former can be sampled again with a single selection of “A,” whereas the latter requires three responses (another “ABC”) to be selected again. The optimal foraging strategy can, in such cases, consist of staying put longer (Wilson, 1986, as described by Neuringer, 2002). The higher cost of moving under the Variability schedule suggests that subjects were sensitive to this aspect of the task.

When longer-term conditional probabilities were examined ($t > 10$ in Figure 9 and 10), subjects did not show signs of the balance hypothesis. Even for rich response alternatives (LL in condition V1, RL in condition V2, LK in condition V3), subjects were no more likely to visit that alternative after being away for 10 responses than they were after being away for 20 responses. Formally, $p(\text{choice}|t) \approx p(\text{choice})$ under the Variability schedule for large values of t , a feature expected of behavior that is random in the long term, even if it is more predictable when modeled only one response at a time.

Don’t Follow the Obtained Reinforcers

It will come as a surprise to some readers that “obtained reinforcers” (whether it be the sum of the reinforcers collected as a consequence of choosing Operandum X over the course of a session, or a “rate” calculated by dividing that empirical sum by the session’s duration) cannot be used as a valid predictor in

regression models. Such empirical estimates of the reinforcement rate, whether at session levels or with respect to interreinforcer-level analysis, have long been the norm in applications of generalized matching (see, for example, the extensive discussion of regression methods for matching analyses in Davison and McCarthy (1988)), and continue to be used regularly (as reviewed by Cowie and Davison, 2016).

Despite this, it is self-evident that responses and reinforcement have a reciprocal causal relationship: A subject cannot be reinforced for a choice it has never made. Because reinforcers depend on responses, and responses are presumed to be influenced by reinforcers, a mutual dependency exists, and as such the empirically observed “obtained reinforcers” cannot be considered independent of responses and do not qualify as an IV for a regression analysis. Caron (2015) puts it succinctly: “Regardless of the effect of the reinforcer, fitting the [generalized matching law] to behavioral data should produce spurious correlations (i.e. systematically positive non-null correlations regardless of sampling errors)” (p. 228). Nor is this concern limited to generalized matching: It applies to any regression analysis in which reinforcers are labeled as an “IV.” Contrary to the requirement that regression predictor be “exogenous,” obtained reinforcers are an example of an “endogenous predictor” (Antonakis et al., 2014). This may, unfortunately, explain why the variance explained by generalized matching regressions is often very high: An excess of confidence driven by a spurious correlation. This difficulty is not unknown to behavior analysts. It has been identified in the past as a “feedback function” (Rachlin, 1978) and as a “coupling coefficient” (Killeen, 1994). Nevertheless, the severity of endogeneity’s distortion of parameter estimates in regression models remains largely unappreciated.

Although endogeneity essentially guarantees that a model will be overconfident and that its parameter estimates will be biased, the character and severity of the problem varies from case to case. In some schedules (such as concurrent VI schedules with very long intervals), the influence of responding on the distribution of reinforcement is generally very small, and only introduces a small bias. In other schedules (such as the Variability schedule in this paper), the influence is enormous, rendering useless the direct estimation of s by using Equation 5 as a regression model. Fortunately, the severity of the endogeneity problem can be detected with “specification tests,” developed for use in econometrics (Hausman, 1978). If there is a concern that a predictor based on obtained reinforcers may be compromised, then one should run the corresponding specification test before undertaking ordinary regression. If the test fails to reject the null, it does not guarantee that the predictor is exogenous, but at least suggests that the effect size of the endogenous confound is small.

Provided a reliable instrumental variable is available, running a specification test is straightforward. Davidson and MacKinnon (1993) describe how to run the Durbin-Wu-Hausman (DWH) test as an augmented regression. Consider Equations 5 and 6, described earlier:

$$LR(B_{i|j}) = LR(k_{i|j}) + s \cdot LR(R_{i|j}) + \epsilon_B \quad (5)$$

$$LR(R_{i|j}) = LR(c_{i|j}) + d \cdot LR(Z_{i|j}) + \epsilon_R \quad (6)$$

Rather than assume that obtained reinforcers $R_{i|j}$ are an endogenous predictor of responses $B_{i|j}$, we can use our data, in combination with an instrumental variable $Z_{i|j}$, to test whether endogeneity is present. To do this, we must first fit the parameters of Equation 6, in order to compute the residuals ϵ_R . Since $Z_{i|j}$ is known to be exogenous (i.e. to be a predictor free of confounds), this will give us an unbiased estimate of the residuals, which we can calculate directly from our parameter estimates:

$$\epsilon_R = LR(R_{i|j}) - (LR(c_{i|j}) + d \cdot LR(Z_{i|j})) \quad (16)$$

Durbin-Wu-Hausman Specification Tests

Figure 11. Critical statistics for the Durbin-Wu-Hausman specification test, used to detect whether obtained reinforcers are an endogenous predictor. The test was performed for each subject in each schedule separately. Scores from the Concurrent schedule are shown as white circles, while scores from the Variability schedule are shown as black diamonds. The dashed lines indicate the two-tailed significance criterion for $\alpha = 0.05$.

To perform the DWH test, we add the residuals from Equation 16 to the terms in Equation 5, giving the following model):

$$LR(B_{i|j}) = LR(k_{i|j}) + s \cdot LR(R_{i|j}) + \beta_{\epsilon} \cdot \epsilon_R + \epsilon_B \quad (17)$$

If the parameter β_{ϵ} in Equation 17 is significantly different than zero, then the predictor $R_{i|j}$ should be considered endogenous with respect to $B_{i|j}$, and consequently the estimate of s given by fitting Equation 5 directly will be biased. Furthermore, dividing β_{ϵ} by the standard error of its estimate (denoted by $\sigma_{\beta_{\epsilon}}$) gives a standard test statistic for how strong the effect is. If the resulting score is positive, it suggests that the obtained reinforcers will tend to be obtained at a faster rate if subjects visit that alternative more often. If, on the other hand, the resulting score is negative, it suggests that visiting an alternative too often will instead slow the rate of reinforcement.

Figure 11 plots the results of 40 DWH specification tests: One for each subject for each of the schedules, fitted using the conditional matching model. For 19 out of 20 cases, obtained reinforcers were detectably endogenous under the Variability schedule, showing large positive effects. This helps to explain why proportions of obtained reinforcement were so much more similar to the response proportions than they were to the target distribution in Figure 2. Under the Concurrent schedule, however, 9 subjects failed the test, and eight of those did so with negative coefficients. This is indicative of the “feedback” characteristic identified by Killeen as being intrinsic to concurrent schedules: If a subject responds too often to an operandum in a concurrent schedule, it will depress the relative proportion of reinforcement earned from that alternative. These DWH tests make clear that, in these data, obtained reinforcers should not be used as a predictor in either condition. However, in other studies (e.g. concurrent VI schedules with long intervals), it may be that obtained reinforcers would only be mildly endogenous, and as such the bias incurred would be relatively mild.

Even if an analyst recognizes that obtained reinforcers are compromised as predictors, there may still

be resistance to using IVE to bypass the problem. An alternative approach is to stop using obtained reinforcement (estimated or otherwise) altogether. Jensen (2014b), for example, simply uses the programmed reinforcer setup probabilities directly to predict behavior, rather than using them as instruments as the current manuscript does. Although this approach gives a good description of behavior, it has theoretical limitations. For example, if had such an analysis been performed in the present study, then the slope of a regression using the Concurrent schedule parameters (Table 2 top) could not be directly compared to that of a regression using the Variability schedule target distribution (Table 2 bottom). Although both types of schedule parameter are independent variables, they are also specific to their respective schedules. As such, the units associated with their slopes would be different, and any comparison between them would be an apples-to-oranges comparison. Furthermore, such a model might be predictive, but could not itself describe how subjects learned. Since subjects do not have direct knowledge of the experiment's settings, it is not plausible that the learning process could depend on them directly.

This is the real advantage of using IVE: The sensitivity parameters depicted in Figure 5 may be directly compared, because they are values in the same units (log responses per estimated log obtained reinforcers). Furthermore, whereas the inner workings of schedules of reinforcement are hidden from view, obtained reinforcers are the main form of feedback that subjects receive. Another advantage is that IVE is easy to adopt: No fancy statistical software is required to implement IVE. As the electronic supplement demonstrates, IVE analyses and the corresponding specification tests may be performed using linear regressions in Excel.

Residuals: The Undiscovered Country

“How variable was behavior?” Framed in this way, the question is deceptive, because it implies an objective standard whereby behavior may be more or less variable. In practice, however, the question is always framed in terms of some theoretical expectation. For example, if behavior is always strictly random, with equal probabilities for all response alternatives, then at one level it is correct to describe the behavior as being “maximally unpredictable” (or, equivalently, as displaying “maximum entropy”). However, at another level, if this behavior is invariant across all schedules of reinforcement, then the parameters describing the behavior are highly predictable (e.g. $p(LK)$ is the same in every scenario). If an analyst expects to see high entropy behavior, and the observes low-entropy behavior, the behavior is surprising, even though the responses it generates are predictable.

Keeping this in mind, whether Figure 8 or Figure 10 constitutes a better map of subjects' unpredictability is a matter of which level of behavior the analyst wants to emphasize. On the one hand, Figure 8 sets a strict criterion for what would make behavior stochastic: The predictability of the next response, given only a subject's current location in the chamber. By this standard, the Concurrent schedule yielded behavior that was harder to predict than behavior under the Variability schedule. However, Figure 10 examined long-term behavior, and found evidence of the balance hypothesis under the Concurrent schedule: When a response had not been made for a long time, there was a predictable increase in its subsequent rate of selection. The variability schedule did not show this pattern. Thus, the choices made under the Variability schedule appeared to be more stochastic when viewed over the long term, at least in this respect.

Reducing behavioral uncertainty by identifying valid predictors is an inevitable consequence of the expectation game that is built into analyses of behavioral variability. Its flipside, however, is that behavioral structure we don't think to look for is often mistaken for variability. For example, if a subject were to travel clockwise through the chamber, pressing each operandum in sequence, then their behavior would display high entropy in the base rates (since equiprobable distributions have higher entropy than peaked distributions), but that same behavior would display low entropy when examined in terms of conditional probabilities. Similarly, when this paper was originally in preparation, response dyads were examined

whether or not a reinforcer was delivered in the interim, which concealed a major source of behavioral structure from view. This source of behavioral structure is now evident in the left-most column of Figure 3: If you know that a reinforcer has just been delivered, you can predict the next response with very high accuracy in most schedules.

As wisely chosen predictors explain more and more of the variance in behavior, it is tempting to conclude that any remaining behavioral variability either conceals further structure (to be revealed by further predictors). Under this view, the residuals in a description of a behavior should be understood as experimenter ignorance. For example, Machado (1997) describes a series of experiments in which subjects were rewarded for switching between two operanda, and did so with high levels of apparent stochasticity (despite there being no constraint forcing them to switch unpredictably). Although Machado observes variability that he cannot account for, he writes that “we cannot rule out the hypothesis that [schedule-behavior dynamics] underlay the differences in the degree of sequence variability across the three experiments” (p. 22). Barba (2012) is even more forceful in claiming that behavioral variability must be explainable: “Even if a contingency increases levels of behavioral variability but in fact differentially reinforces some sequences on the basis of properties other than their variability, the increase in variability levels, whatever it is, does not allow one to measure a differentiation process. In this case, *the increase in variability represents a secondary effect of differential reinforcement*” (p. 222, italics mine). According to this view, even when a schedule apparently increases the variability of a behavior, any dimension of the behavior that can be shown (or, indeed, imagined) to receive differential reinforcement should be given explanatory priority.

There is another way of conceptualizing the residual variability of behavior, however. Rather than treating the residuals as “error” and thus something to be conquered by additional predictors, the variability of behavior (or, indeed, of any measure of scientific interest) can be viewed as a property of those data. Even if that variability is ultimately explained by additional predictors, any comprehensive description of the behavior must include a description of its variability at the level relevant to the analyst’s question.

For example, it has long been understood that, in predator-prey pursuit dynamics, prey gain an advantage by making their escape behavior unpredictable but not entirely random (since being *entirely* unpredictable would include counterproductive behaviors like leaping into the predator’s jaws). These so-called “protean behaviors” (Driver and Humphries, 1988) are generally assumed to be instinctual, and thus depend on some hard-wired mechanism. And indeed, careful analysis of these behaviors sometimes reveals underlying structure. When Domenici et al. (2008) studied the escape behaviors of cockroaches, they found that although the direction of a cockroach’s escape is highly unpredictable within a $60^\circ - 90^\circ$ arc, the precise vector of escape was not selected uniformly within that span. Instead, each subject chose from one of three or four “preferred angles,” spaced $15^\circ - 20^\circ$ apart, being much more likely to move along one of those vectors than those that lay between. Objectively, going from “uniform along a 90° arc” to “mostly concentrated along three or four vectors” constitutes a dramatic reduction in unpredictability, so in one sense, the variability of this behavior is revealed to be less than what was previously believed. On the other hand, however, what matters to the survival of the cockroach is the unpredictability of the behavior *from the perspective of a predator*. If the radial structure of the escape behavior is too subtle for predators to exploit, then it is sensible to describe the adaptive randomness of the behavior, even while it is also important to acknowledge its underlying structure in pursuit of its underlying mechanism.

Just as “any action that closes the circuit” is a typical description of a response given in operant conditioning, anything that results in the cockroach evading a predator deserves to be called escape behavior. In both the operant response and the escape behavior, then, we are faced not with singular atomic responses, but *distributions* of possible responses that belong to the same functional class. The entropy of that distribution is a fundamental descriptor of a class of behaviors, as much a dimension of behavior as

characteristics like rate, duration, force, or topography. Indeed, an account of learning that tries to deal with singular atomic behaviors, rather than with distributions, falls prey to *reductio ad absurdum*, since no two singular instances of a behavior are precisely identical.

In general, it is sensible to ask, “How variable is behavior at a particular level of analysis?” while simultaneously asking, “How does that variability arise?” Having a satisfactory answer to the second question that explains some share of the variance does not change the answer to the first question. It is accurate to say that the Concurrent schedule in this study was more variable on a trial-by-trial basis than the Variability schedule, and it is also accurate to say that this arose in part because the feedback function built into the concurrent schedule drove behavior toward the target distribution via differential reinforcement, just as predicted by Barba’s balance hypothesis.

So: Can response variability be reinforced? How theorists answer this question depends critically on how they define “the response” in an operant context. Purists like Barba (2014) advocate an atomic approach to defining the response, such that feedback that increases variability across whole classes of responding must instead be understood in terms of many separate learning processes, which only *appear* to reinforce variability as a secondary effect. The present study, on the other hand, argues for a more flexible and powerful approach, whereby any given response is recognized to possess multiple dimensions (which can potentially be shaped independently of one another) and to belong to a hierarchy of overlapping response classes. Whether reinforcement can influence any given dimension of behavior, or of broad classes of behavior, is a theoretically valid and ultimately empirical question. Furthermore, even if a deterministic mechanism is uncovered that explains how a series of reinforcers led to an increase in behavioral variability, it would still be valid to say that variability had been reinforced. There is much more to be gained from a “Yes, and?” approach to explaining behavioral variability than from “But, actually?” contrarianism.

That said, studying the mechanisms that underlie behavioral variability remains an interesting and worthwhile activity. This may reveal that behavioral variability has non-random underpinnings, but such discoveries won’t obviate the needs to measure that variability and to understand what contexts give rise to it. In this study, Barba’s balance hypothesis has some support in the Concurrent case, likely due to the manner in which concurrent schedules steer behavior by differential reinforcement. Such behavior still appears highly unpredictable, so better understanding this link will make it easier to anticipate what sorts of behaviors will arise due to schedules of reinforcement outside a lab setting.

Behavior analysts should not rule out the possibility that some mechanisms behind behavioral variability may be intrinsically unpredictable. The general consensus in behavioral neuroscience is that decision processes are not only intrinsically variable, but that said variability is deliberately harnessed to ensure that behavior is adaptive (Gold and Shadlen, 2007). For example, Tervo et al. (2014) propose that the anterior cingulate cortex is involved in memory-driven responses. When subjects transition from low- to high-stochasticity behavior, they report that the cingulate is inhibited by inputs from the locus coeruleus. In effect, locus coeruleus inputs could act as a modulator of stochasticity, ramping the degree to which behavior is based on past experience up or down as needed. If new evidence remains consistent with this interpretation, and if locus coeruleus activity is modulated by reinforcement, then even under the strictest Skinnerian interpretation, the reinforcer would have to be described as increasing variability by inhibiting learned reward signals.

Insofar as external behavior is a behavior analyst’s sole venue for theorizing, some critical mechanisms underlying behavior will never be explained. Skinner (1985) writes, “A behavioural account is incomplete, in part because it leaves a great deal to neurology” (p. 297). In the decades since this concession the scope and sophistication of neuroscience has vastly grown, particularly as it relates to the brain activity of awake behaving animals. Insofar as behavior arises from neural activity, and that neural activity is being revealed

in ever-growing detail, the tenability of a behaviorism in which “the response” is a discrete atomic unit is increasingly questionable. If they recognize that the activity of neurons represents a lower rung on the ladder of reductionism, behaviorists must then also accept a measurable response arises from activity distributed throughout a network over time. Rather than being discrete, responses necessarily take a fuzzy form that resist being discretized. Without discretized responses, “pure” Skinnerian accounts can only ever be higher-level approximations, rather than describing fundamental mechanisms. In light of this, declaring theoretical injunctions against certain kinds of conditioning seems premature.

A Closing Thought on Models

A common piece of advice that statisticians give researchers is to perform the best analysis they are able, but not to use analyses they do not understand. The thinking goes that because all statistical procedures require a certain amount of expertise to implement and interpret, it would be better to use a slightly ill-fitting analysis while being aware of its shortcomings than to use a cutting-edge algorithm without being able to recognize when an error has been made.

I raise this point because the analytic methods described in this paper are a compromise, favoring breadth of readership over depth of analysis. Although two central pillars of this paper’s analysis (the conditional matching model and instrumental variable estimation) will be unfamiliar to most readers, their implementation does not require advanced techniques or special software. The analyses described in this paper were performed in Microsoft Excel to demonstrate that no command more advanced than its `LINEST()` least-squares multiple regression function is required to use these tools.

Had I run my own ideal analysis, however, I would have done several things differently. Rather than perform the analysis in Excel, I would have used the R programming language. Rather than fit the regressions using least squares, I would have made use of a framework called Stan to obtain Bayesian parameter estimates. And, finally, rather than fitting the data for each subject separately and reporting on the distribution of obtained parameters, I would have fit a multi-level model that yielded estimates of both subject-level and population-level parameters simultaneously. Every one of these alternatives would have given rise to a stronger analysis, but would have done so at the expense of some share of this study’s potential readership.

The electronic supplement is an Excel file, because Excel is very widely used and is comparatively easy to learn. As a scientific tool, however, spreadsheet software such as Excel has two major shortcomings. The first is that its statistical toolkit is quite limited, as its development is governed chiefly by the demands of Microsoft’s enterprise customers. The second, which is more serious when it comes to communicating science, is that the analysis is forced to coexist among the data in the cells of the spreadsheet. This makes trying to read an Excel analysis, in order to understand how it works, devilishly difficult. Programmers often lament the difficulty of working with “spaghetti code,” which arises when a every part of a program is full of references to other bits of code. To understand such code, the analyst must trace a chain of events that weaves a spaghetti-like nonlinear path back and forth through the program. A statistical analysis in Excel is even worse, because the path woven by the analysis runs in two dimensions through the cells of the spreadsheet. Consequently, complex Excel analyses are difficult to understand and almost impossible to debug. Both of these difficulties contributed, for example, to the scandal surrounding a very influential analysis of national debt and GDP by Reinhart and Rogoff, whose use of Excel concealed both questionable practices and outright coding errors (described in Dafoe, 2014). With a more transparent analytic framework, the erroneous conclusions of the paper might have come under scrutiny much earlier.

The R programming language, by contrast, has a lot to offer the behavior analyst. It is free, open source, and arbitrarily extensible through add-on libraries. Since it is one of the primary platforms for statisticians, essentially every analytic tool worth taking seriously has already been implemented. Its greatest benefit,

however, comes from encoding the analysis as a sequential script, with each line performing a specific analytic step. Not only does this help the analyst understand and double-check their own work, but it also helps their colleagues to reproduce and modify the original analysis. A huge array of free tutorials is available for every kind of user, and the time saved once one is proficient is well worth the initial effort. Weston and Yee (2017) provide a list of resources for psychologists interested in adding this tool to their toolkit.

Making the transition to R opens up a wide range of models and methods. Of particular interest to behavior analysts will be the possibility of obtaining parameter estimates for custom nonlinear models. For many years, analysts trying to fit the parameters of a model to their data have been forced to make a choice between two options: (1) manipulate a proposed model arithmetically until it is sufficiently linear to use ordinary least squares, or (2) use iterative curve-fitting for unavoidably non-linear functions. When OLS is used, statistical inference is possible but can be complicated by the model not quite abiding by linear regression's many assumptions. When curve-fitting is used, statistical inference is generally impossible, as it is difficult to be certain that the estimates are indeed the best-fitting (as opposed to being trapped in some local minimum) and impossible to put error bars on the estimates. Using R gives the analyst ways of escaping this dilemma and fitting highly non-linear models without resorting to curve-fitting.

My preferred framework for doing so is Stan (Carpenter et al., 2017), a framework for specifying and fitting arbitrary models. Stan is a Bayesian framework that uses a Markov chain Monte Carlo algorithm to estimate the distributions of parameters (for more on Bayesian inference generally, see Heino et al. 2018). This means that although parameter fitting is computationally intensive, the distribution of those parameters may take any form. If the parameter distributions are multimodal (perhaps because there are multiple solutions to describing the observed data), Stan can discover this. If parameter estimates are correlated (as they almost always are in non-linear models), Stan can discover this as well. And because Bayesian models operate in the common currency of probability, a wider range of model comparison tools can be exploited by analysts to help determine how well models perform relative to one another.

The single greatest advantage of adopting Stan is that it greatly enhances the possibilities of running multi-level models, or MLMs (Gelman and Hill, 2007). Under an MLM, the analyst concedes that it may be possible to learn something about Subject 1 from the data collected from Subject 2, and vice versa. Subjects, in other words, are probably somewhat similar, and so while a subject's own data is the best source of evidence about that subject's parameters, some additional evidence can be gleaned from the other subjects. The challenge is to discover how much subjects resemble one another, and to calibrate each subject's influence on the other subjects accordingly. MLMs strike this balance, and in doing so, two benefits are gained. The first is that the parameters for each subject tend to be more accurate, since every source of evidence is considered. The second is that the degree of inter-relatedness between subjects takes the form of a population-level distribution of subjects. With such a model, the analyst may not only put error bars on individual estimates, but can also identify whether any given subject is a typical member of the population or is instead an outlier. The mathematical challenge of fitting an MLM are tremendous, so although such models have been known conceptually for decades, their use was not practical for most research problems until very recently. Stan is one of a number of new frameworks that is able to use MCMC efficiently to obtain MLM estimates.

A major reason that behavior analysts should be intrigued by MLMs is because they allow good inferences to be made about behavior with less data. Traditionally, studies of behavior have collected enormous datasets because that is the only way to get precise estimates of a subject's behavior. In the present study, for example, each subject completed 120 sessions, a process that (on a five-day work week) took half a year to complete. When each subject is a data island, each must be studied exhaustively. In an MLM, however, the possibility of being able to incorporate evidence from other subjects greatly reduces

the amount of data needed to obtain an estimate of comparable precision. By using data more efficiently, analysts can run more conditions and more experiments without a corresponding increase in time or cost.

Would bringing these addition tools to bear on the present data change the results? Not that I've seen, largely due to the sheer volume of data (about 2.36 million responses). The chief benefit of an expanded analytic toolkit is not that it yields incrementally better estimates for a given model, but that it makes entirely new classes of model available for consideration. Thanks to the wider availability of quantitative tools and the computational resources to exploit them, the next generation of behavior analysts will be able to ask questions never before considered by the field.

ACKNOWLEDGMENTS

This manuscript owes a tremendous debt to Allen Neuringer, who oversaw the original experiment and collaborated on the manuscript's many subsequent iterations. Thanks also to Armando Machado and Peter Killeen, who provided valuable commentary. Finally, thank you to Albyn Jones for his feedback and advice on statistical matters, and to Gene Olson for his assistance with animal care. This research was supported Research Grant MH068259 from the National Institute of Health.

REFERENCES

- Antonakis, J., Bendahan, S., Jacquart, P., and Lavive, R. (2014). Causality and endogeneity: Problems and solutions. In Day, D. V., editor, *The Oxford Handbook of Leadership and Organization*, pages 93–117. Oxford University Press, New York, NY.
- Balsam, P. D., Deich, J. D., Ohyama, T., and Stokes, P. D. (1998). Origins of new behavior. In W. T. O'Donohue, editor, *Learning and Behavior Therapy*, pages 403–420. Allyn & Bacon, Needham Heights, MA.
- Barba, L. S. (2012). Operant variability: A conceptual analysis. *The Behavior Analyst*, 35:213–227.
- Barba, L. S. (2014). Controlling and predicting unpredictable behavior. *The Behavior Analyst*, 38:93–107.
- 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.
- Baum, W. M. (2013). What counts as behavior? the molar multiscale view. *The Behavior Analyst*, 36:283–293.
- Baum, W. M. and Aparicio, C. F. (1999). Optimality and concurrent variable-interval variable-ratio schedules. *Journal of the Experimental Analysis of Behavior*, 71:75–89.
- Boren, J. J., Moerschbaecher, J. M., and Whyte, A. A. (1978). Variability of response location on fixed-ratio and fixed-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 30:63–67.
- Braitenberg, V. (1986). *Vehicles: Experiments in synthetic psychology*. MIT Press, Cambridge, MA.
- Caron, P. (2015). Matching without learning. *Adaptive Behavior*, 23:227–233.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M. A., Guo, J., Li, P., and Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76:1–32.
- Cowie, S. and Davison, M. (2016). Control by reinforcers across time and space: A review of recent choice research. *Journal of the Experimental Analysis of Behavior*, 105:246–269.
- Dafoe, A. (2014). Science deserves better: The imperative to share complete replication files. *PS: Political Science and Politics*, 47:60–66.
- Davidson, R. and MacKinnon, J. G. (1993). *Estimation and Inference in Econometrics*. Oxford University Press, New York, NY.

- Davison, M. and McCarthy, D. (1988). *The Matching Law: A Research Review*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Denny, J. and Neuringer, A. (1998). Behavioral variability is controlled by discriminative stimuli. *Animal Learning & Behavior*, 26:154–162.
- Diebold, F. X. (2004). The Nobel Memorial Prize for Robert F. Engle. *Scandinavian Journal of Economics*, 106:165–185.
- Domenici, P., Booth, D., Blagburn, J. M., and Bacon, J. P. (2008). Cockroaches keep predators guessing by using preferred escape trajectories. *Current Biology*, 18:1792–1796.
- Driver, P. M. and Humphries, D. A. (1988). *Protean Behavior: The Biology of Unpredictability*. Oxford University Press, Oxford, UK.
- Gelman, A. and Hill, J. (2007). *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge University Press, Cambridge, UK.
- Glaze, C. M., Filipowicz, A. L. S., Kable, J. W., Balasubramanian, V., and Gold, J. I. (2018). A bias-variance trade-off governs individual differences in on-line learning in an unpredictable environment. *Nature Human Behavior*, 2:213–224.
- Glimcher, P. W. (2005). Indeterminacy in brain and behavior. *Annual Review of Psychology*, 56:25–56.
- Gold, J. I. and Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30:535–574.
- Gorobets, A. (2005). The optimal prediction simultaneous equations selection. *Economics Bulletin*, 3:1–8.
- Hausman, J. A. (1978). Specification tests in econometrics. *Econometrica*, 46:1251–1271.
- Hendry, D. F. (2004). The Nobel Memorial Prize for Clive W. J. Granger. *Scandinavian Journal of Economics*, 106:187–213.
- James, L. R. and Singh, B. K. (1978). An introduction to the logic, assumptions, and basic analysis procedures of two-stage least squares. *Psychological Bulletin*, 85:1104–1122.
- Jensen, G. (2014a). Beyond dichotomy: Dynamics of choice in compositional space (doctoral dissertation). Technical Report DOI 10.7916/D8NC5ZC4, Columbia University.
- Jensen, G. (2014b). Compositions and their application to the analysis of choice. *Journal of the Experimental Analysis of Behavior*, 102:1–25.
- Jensen, G. and Neuringer, A. (2008). Choice as a function of reinforcer “hold”: From probability learning to concurrent reinforcement. *Journal of Experimental Psychology: Animal Behavior Processes*, 34:437–460.
- Jensen, G. and Neuringer, A. (2009). Barycentric extension of generalized matching. *Journal of the Experimental Analysis of Behavior*, 92:139–159.
- Jensen, G., Ward, R. D., and Balsam, P. D. (2013). Information: Theory, brain, and behavior. *Journal of the Experimental Analysis of Behavior*, 100:408–431.
- Killeen, P. R. (1994). Mathematical principles of reinforcement. *Behavioral and Brain Sciences*, 17:105–172.
- Killeen, P. R. (2015). The logistics of choice. *Journal of the Experimental Analysis of Behavior*, 104:74–92.
- Kono, M. (2017). The effects of fixed-interval schedules on variability of pigeons’ pecking location. *Journal of the Experimental Analysis of Behavior*, 108:290–304.
- Lau, B. and Glimcher, P. W. (2005). Dynamic response-by-response models of matching behavior in rhesus monkeys. *Journal of the Experimental Analysis of Behavior*, 84:555–579.
- Luce, R. D. (1959). *Individual Choice Behavior: A Theoretical Analysis*. Wiley, New York, NY.
- MacDonall, J. S. (2009). The stay/switch model of concurrent choice. *Journal of the Experimental Analysis of Behavior*, 91:21–39.
- Machado, A. (1992). Behavioral variability and frequency-dependent selection. *Journal of the Experimen-*

- tal Analysis of Behavior*, 58:241–263.
- Machado, A. (1997). Increasing the variability of response sequences in pigeons by adjusting the frequency of switching between two keys. *Journal of the Experimental Analysis of Behavior*, 68:1–25.
- Machado, A. and Tonneau, F. (2012). Operant variability: Procedures and processes. *The Behavior Analyst*, 35:249–255.
- Manohar, S. G. and Husain, M. (2013). Attention as foraging for information and value. *Frontiers in Human Neuroscience*, 7:Article 711.
- Neuringer, A. (2002). Operant variability: Evidence, functions, and theory. *Psychonomic Bulletin & Review*, 9:672–705.
- Neuringer, A. and Jensen, G. (2013). Operant variability. In Madden, G. J., Dube, W. V., Hackenberg, T. D., Hanley, G. P., and Lattal, K. A., editors, *The APA Handbook of Behavior Analysis: Vol. 1, Methods and Principles*, pages 513–546. American Psychological Association, Washington, DC.
- Neuringer, A., Kornell, N., and Olufs, M. (2001). Stability and variability in extinction. *Journal of the Experimental Analysis of Behavior*, 27:79–94.
- Odum, A. L., Ward, R. D., Burke, K. A., and Barnes, C. A. (2006). The effects of relayed reinforcement on variability and repetition of response sequences. *Journal of the Experimental Analysis of Behavior*, 86:159–179.
- Racey, D. E., Young, M. E., Garlick, D., Pham, J. N., and Blaisdell, A. P. (2011). Pigeon and human performance in a multi-armed bandit tasks in response to changes in variable interval schedules. *Learning & Behavior*, 39:245–258.
- Rachlin, H. (1978). A molar theory of reinforcement schedules. *Journal of the Experimental Analysis of Behavior*, 30:345–360.
- Schwartz, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, 6:461–464.
- Silberberg, A. and Fantino, E. (1970). Choice, rate of reinforcement, and the changeover delay. *Journal of the Experimental Analysis of Behavior*, 13:187–197.
- Skinner, B. F. (1985). Cognitive science and behaviourism. *British Journal of Psychology*, 76:291–301.
- Stokes, P. D., , Lai, B., Holtz, D., Rigsbee, E., and Cherrick, D. (2008). Effects of practice on variability, effects of variability on transfer. *Journal of Experimental Psychology: Human Perception & Performance*, 34:640–659.
- Stokes, P. D. and Balsam, P. D. (1991). Effects on reinforcing preselected approximations of the topography of the rat's bar press. *Journal of the Experimental Analysis of Behavior*, 55:213–231.
- Tervo, D. G. R., Proskurin, M., Manakov, M., Kabra, M., Vollmer, A., and Branson, K. (2014). Behavioral variability through stochastic choice and its gating by anterior cingulate cortex. *Cell*, 159:21–32.
- Ward, R. D., Kynaston, A. D., Bailey, E. M., and Odum, A. L. (2008). Discriminative control of variability: Effects of successive stimulus reversals. *Behavioural Processes*, 78:17–24.
- Weston, S. J. and Yee, D. (2017). Why you should become a useR: A brief introduction to R. *Observer*, 30:18–19.
- Wilson, N. (1986). Standards of deviation: Tools for reinforcing variability (undergraduate thesis). Technical Report OCLC Number 268872878, Reed College.