

Operant Variability and Voluntary Action

Allen Neuringer and Greg Jensen
Reed College

A behavior-based theory identified 2 characteristics of voluntary acts. The first, extensively explored in operant-conditioning experiments, is that voluntary responses produce the reinforcers that control them. This bidirectional relationship—in which reinforcer depends on response and response on reinforcer—demonstrates the functional nature of the voluntary act. The present article focuses on the second characteristic: a similar bidirectional relationship between reinforcement and the predictability/unpredictability of voluntary acts. Support for the theory comes from 2 areas of research. The first shows that levels of behavioral variability—from highly predictable to randomlike—are directly influenced by reinforcers. Put another way, variability is an operant dimension, analogous to response rate and force. The second source of support comes from psychophysical experiments in which human participants judged the degree to which “choices” by virtual actors on a computer screen appeared to be voluntary. The choices were intermittently reinforced according to concurrently operating schedules. The actors’ behaviors appeared to most closely approximate voluntary human choices when response distributions matched reinforcer distributions (an indication of functionality) and when levels of variability, from repetitive to random, changed with reinforcement contingencies. Thus, voluntary acts are characterized by reinforcement-controlled functionality and unpredictability.

Keywords: voluntary action, theories of free will, reinforced variability, concurrent schedules of reinforcement, choice

Debates have raged for thousands of years concerning voluntary action: its characteristics, explanation, and, indeed, its very existence. One of the more difficult issues has been how to reconcile the apparent freedom of an individual’s actions with a lawful universe. Many writers on this topic have claimed that we are able to initiate actions to do—within reasonable limits—what we want to do, but scientists and advocates of a scientific worldview have also maintained that we are physical/biological beings whose behaviors are subject to the laws of nature. How to reconcile these two? In this article we attempt to identify the types of observable behaviors that lead to attributions of volition. Our claim is that voluntary behavior is indicated by functional changes in reinforcement-controlled behaviors including, importantly, levels of behavioral variability. We provide two main lines of evidence. The first comes from experimental studies showing that levels of response variability, ranging from stereotyped repetitions to randomlike variations, are controlled by reinforcement contingencies. Stated differently, variability is an operant dimension of behavior. The second is that human participants attribute volition to agents whose behaviors,

including the variability (or predictability) of the behaviors, change adaptively in response to environmental contingencies. Together, the evidence indicates that operant variability is an essential characteristic of voluntary action.

The predictability of behavior—or more to the point, its potential unpredictability—has been debated at least as far back as the offerings of the ancient Greek philosophers Democritus and Epicurus. Democritus proposed that the universe was composed of *atoms*, the motions of which were determined by interactions with other atoms. He described a causally ordered universe, one in which the laws of nature could be deciphered, thereby providing his countrymen with the hope of understanding and influencing events. Epicurus had little disagreement with Democritus’s depiction of a natural world but argued that something must be added to account for the novelty and creativity seen throughout nature and for the initiative and volition seen in human behavior. Epicurus’s solution was to posit occasional random *swerves* of the atoms. An assumption of a causally determined universe, one that included humans as well as everything else, on the one hand, and the possibility of uncaused and unpredictable human voluntary actions, on the other, became a focus of debates that continue to this day. Some participants in the debates are described as determinists, who reject a special status for voluntary behaviors; others as compatibilists, who maintain that voluntary behaviors are consistent with universal determinism; and still others as incompatibilists, some of whom argue that voluntary actions differ in kind from physically determined events. We will return to some of these views at the end of the article, but here it suffices to note the ongoing difficulty and confusion surrounding the topic. According to a contemporary philosopher who is deeply involved in this issue, free will provides “perhaps the most voluminously debated

Allen Neuringer and Greg Jensen, Department of Psychology, Reed College.

This research was supported by National Institutes of Health Grant MH068259. We thank Paul Piff and Jacob Rothstein for their contributions to the empirical studies and theory development. We also thank Peter Killen and John Wixted for their constructive critiques.

Correspondence concerning this article should be addressed to either Allen Neuringer, Department of Psychology, Reed College, Portland, Oregon 97202, or Greg Jensen, who is now at Columbia University, 406 Schermerhorn Hall, 1190 Amsterdam Avenue MC 5501, New York, NY 10027. E-mail: allen.neuringer@reed.edu or belarius@gmail.com

of all philosophical problems” (Kane, 2002, p. 3). Another wrote, “The persistence of the free will problem in philosophy seems to me something of a scandal. After all these centuries of writing about free will, it does not seem to me that we have made very much progress” (Searle, 2007, p. 37). Both of these quotes refer to *free will*, and we will use that term interchangeably with *voluntary behavior*, although the latter is taken by some to be only a behavioral manifestation, or subset, of the former and by others to be explained by the former.

Part of the difficulty in explaining volition comes from an apparent incompatibility between two often-noted characteristics. On the one hand, voluntary acts are thought to be intentional, purposeful, goal-directed, rational, or adaptive. We will refer to this set of related characteristics as indicating the “functional” nature of those behaviors. On the other hand, they are simultaneously thought to be internally motivated and controlled and therefore unpredictable, at least some of the time. This unpredictability separates voluntary acts from other functional behaviors (e.g., reflexes) and separates them (according to some) from Newtonian cause–effect explanations of other physical phenomena. Proposed explanations of the unpredictability run the gamut from an autonomous soul or mind to quantum-mechanical random events, but they are all ultimately motivated by the presumed inability of a knowledgeable (and even supremely knowledgeable) observer to anticipate the particulars of a voluntary act.

Thus, the key problem derives from the difficulty in combining the unpredictable (perhaps even the unpredictable in principle) with the functional. Functionality implies potential predictability. If an observer knows what an individual is striving for, or attempting to accomplish, then together with knowledge of the individual’s past experiences and current circumstances, at least somewhat accurate predictions can be made concerning the individual’s future actions. The essential point here is that, to the extent that behavior is related in orderly fashion to environmental variables, as is generally required for functionality, predictions are possible. But, again, voluntary acts are often characterized by their manifest or potential unpredictability.

The Operant Model

Many attempts have been made to reconcile these two attributes of volition, but most theories leave unanswered a basic question: What is voluntary behavior? That is, what does volition look like, and what are its distinguishing characteristics? Missing from discussions are attempts to describe—and by so doing, to define—voluntary acts in terms of empirically observable and testable characteristics. Writers have often assumed that the phenomenon is intuitively obvious, but these appeals to the obvious are a serious problem in the field, especially when different authors have divergent intuitions.

One of the few theories to be based on empirical evidence is the Skinnerian model of operant behavior (Skinner, 1938). Skinner (1974) and other behavioral researchers have distinguished between Pavlovian responses (or stimulus-elicited reflexes) and operant responses (often characterized as emitted voluntary responses). Pavlovian reflexes consist of a conditioned stimulus that, because of associations with an unconditioned stimulus, comes to elicit a conditioned reflex. Operant responses are created by a three-term contingency: S^D - R^C - S^R . When a response (R^C) pro-

duces a reinforcing consequence (S^R) in the presence of a discriminative stimulus (S^D), the response may come to serve as an operant.

Skinner (1938) described two qualities of the operant that are important for our discussion. First, operants are adaptive in producing reinforcing effects. Unlike reflexes (which are closely linked to unconditioned and conditioned stimuli), operant responses also involve a continually changing feedback-loop relationship between organism and its environment. Thus, operants manifest a “looseness” that enables rapid behavioral changes in response to changing circumstances. The operant is often described as purposeful or goal-directed to distinguish it from stimulus-elicited behaviors.

Second, the individual R^C is a member of a class C of instances, a generic class comprised of functionally similar (although not necessarily physically similar) actions (Skinner, 1935/1961). An example may help to explain this point. Jack desires to have his wine glass refilled. He might reach across the dinner table for the wine bottle, ask his dinner companion to pass the wine, or beckon to the waiter. Each of these acts, although differing in physical details, is a member of the same operant class because each serves the same functional relationship between S^D (the empty wine glass) and the S^R goal (wine in the glass). Some responses may be more functional than other members of the class, and cues may indicate which of these responses is most likely. For example, if the waiter is nearby and the wine bottle is distant on the table, then beckoning to the waiter becomes more likely. But in other cases, the behaviors appear to be selected with more equal probabilities, and prediction of the instance becomes difficult.

As just suggested, members of a particular class of behaviors may be divided into subclasses, and even here variability may characterize aspects of the response. For example, if “ask for wine” is the activated subclass, then the exact moment of a verbal request, or the particular words used, or the rhythm or loudness may all be difficult to predict. In a common laboratory model of operant behavior, a rat presses a lever to gain food pellets and the characteristics of the press (one paw vs. both paws, with short or long latency, with high or low force, etc.) are sometimes predictable but are oftentimes not. Thus, according to a Skinnerian model, functionally equivalent instances emerge unpredictably from within a class or subclass, as though generated by a stochastic process (Skinner, 1938; see also Moxley, 1997). To state this differently, there is variance within the operant, manifested as the emergence of instances from a set comprised of functionally related but often physically dissimilar behaviors. Sensitivity to contingent reinforcers causes the operant to appear purposeful and goal-directed and provides a reason (i.e., explanation) for its occurrence. The within-class variability produces an appearance of independence and internal guidance, and Skinner (1938, p. 20) related this variability to the “emitted” and volitional nature of the operant.

An extensive body of research has powerfully demonstrated the predictable ways in which reinforcing events influence operant behaviors and, equally, ways in which operants function to produce these events. Reinforcers, presented contingent upon responses, enable strong predictions in the study of shaping, schedule-maintained responding, and discriminative-stimulus control. Such functionality has been well documented in behavior analysis, behavioral therapies, neuroscience, and education, among

other areas of study. However, the contribution of behavioral variability to explanations of volition is less clearly defined. The key questions are these: What governs differences (or changes) in variability, and how do they contribute to an explanation of voluntary acts? Skinner hypothesized that variability is present in all operants, providing the substrate for shaping new responses, but he also indicated that continued experience with reinforcement tends to define and narrow the response class (Skinner, 1974). A related hypothesis is that variability is a sign of observer ignorance and that with knowledge the domain of variability shrinks. But neither a constant level of variability nor its continual decrease will explain the voluntary nature of voluntary behavior. As argued by others, constant variability suggests a level of uncontrollability that is quite the opposite of the “control from within” posited for volitional acts. And if variability inevitably decreases with reinforcement experience, all operant responding would ultimately be predictable, in ways similar to Pavlovian reflexes, and, again, this is inconsistent with our perceptions of volitional acts. Another alternative is that variability is induced by intermittent withholding of reinforcement, as under variable interval or ratio schedules, or withholding it altogether, as during periods of extinction, or, more generally, by conditions that induce low expectations for reinforcement (Balsam, Deich, Ohyama, & Stokes, 1998; Gharib, Gade, & Roberts, 2004). But here, too, common views of volition are not consistent with unpredictability’s being relegated solely to disappointing situations or to those in which efforts are unsuccessful.

We will provide experimental evidence to show that these phenomena and hypotheses do not suffice to account for the variability of operant behavior, nor, we submit, do they explain volition. Variability *sometimes* decreases with training, it is *sometimes* explained as knowledge accumulates, and it *sometimes* increases when reinforcers are withheld. But in many cases, levels of behavioral variability, from highest to lowest, are sensitive to how the variability intersects with the environment; and these levels can change, moment to moment, depending upon their consequences. This includes extraordinarily high levels of variability, indistinguishable from random processes, when it is functional to behave unpredictably. The controlled nature of operant variability—under contingent control by environmental consequences—provides the basis for a theory of volition.

Operant Variability and Voluntary Action (OVVA) Theory

We describe a theory of volition that differs from Skinnerian theory mainly in terms of the character and contribution of behavioral variability. The operant variability and voluntary action (OVVA) theory proposes that variability itself is an operant dimension—and that operant variability is an essential characteristic of voluntary behavior.

All operant responses are comprised of many dimensions (topography, location, force, speed, etc.). “Reinforcement of an operant” can best be understood as reinforcement contingent upon one or more (but generally not all) dimensions of a response. Experimental and analytic precision can be gained if we describe operants in terms of those dimensions of a response that are influenced by reinforcement (i.e., “operant dimensions”). An example comes from the study of lever pressing by rats, in which reinforcement is contingent upon a minimum *force* that is exerted

in a particular *direction* at a particular *location*, all of these being response dimensions that can be controlled by contingent reinforcers. The operant nature of variability implies that *levels of variability* are influenced or controlled by discriminative stimuli and reinforcing consequences in the same way as these other dimensions are. Thus, in some cases, repetitive or systematic (and therefore easily predictable) responding may be reinforced, whereas in other cases behaving unpredictably may gain reinforcement. We will refer to this capacity of variability to be controlled by reinforcement contingencies as the operant nature of variability, or shorthand, operant variability. As will be seen, reinforcer contingencies exert simultaneous control over multiple operant dimensions (e.g., topography and variability).

A few words about terminology: *Variability* connotes an uncertainty or entropy continuum, the two end points of which are repetitive responses (characterized by stereotypy) and random responses (characterized by equiprobability). Location along the continuum (i.e., the level of variability) can be described in terms of statistical evaluations of relative frequencies (e.g., of instances, dyads, and triads) and other statistical tests. *Predictability* refers to a response by an observer and depends on level of variability. Thus, for example, whereas stereotypic responses are easily predicted, random responses are no more predictable than a chance guess and therefore are referred to as unpredictable.

Although the terms *stochastic* and *random* refer to the same phenomenon in most usages, for purposes of exposition in the present article, we will use *stochastic* to indicate the general case of probabilistic emission and *random* to indicate cases where probabilities are equal. That is, both *random* and *stochastic* imply *statistical independence*, but *stochastic* is silent on *proportional distribution*. As an illustrative example, imagine an urn filled with 1,000 colored balls. The urn is shaken, and one ball is blindly selected. After selection, the ball’s color is noted, returned to the urn, and the selection process is repeated. If there are equal numbers of blue and red balls, then the output will be random, and prediction will be no better than chance (i.e., the probability of a correct prediction would be .50). As the distribution becomes more asymmetric, prediction becomes increasingly accurate (e.g., if the urn contained 900 red balls and 100 blue, then prediction accuracy would rise to .90 if one always predicted red). Thus, stochastic outputs are more or less predictable depending upon the relative frequencies of the items, the two colors. It is also the case that the greater the number of *different* item classes (e.g., different colors), the less predictable any given instance. If the urn contained equal numbers of 20 different colors, for example, then the chance level of prediction would be .05 (rather than .50 in the two-color case). (Discussion of these concepts can be found in Gigerenzer et al., 1989.) For purposes of explaining our position, we posit that individual operant responses are emitted in a fashion similar to the selection of these colored balls (i.e., stochastically). Of most importance for a theory of volition, the constituents of the set (the colors, in this example), the size of the set (the number of colors), and the probability or relative frequency distribution of the within-set instances are all variables, controlled by reinforcement contingencies. All of these contribute to the level of variability and therefore to predictability.

Discussions of free will and volition often focus on unpredictability because, as indicated previously, it is taken to be a sign of independence and an indicator of internal—and therefore autono-

mous—control. However, also as indicated, random responding alone cannot explain volition. We do not attribute volition to truly random events, such as the throw of dice or emission of atomic particles (Dennett, 2003; Popper & Eccles, 1977), and truly random responding would often be maladaptive. Thus, the requirement to behave unpredictably (so as not to be enslaved by environmental determinism) seems to be inconsistent with the requirement to behave predictably (in order to attain a goal). A contemporary philosopher put the problem in these somewhat different terms:

The problem of free will in its broadest outline is this. Free will seems to be incompatible with both determinism and indeterminism. Free will seems, therefore, to be impossible. But free will also seems to exist. The impossible therefore seems to exist. (Van Inwagen, 2000, p. 11)

Although randomness in general cannot solve the problem of volition, under some circumstances randomlike behavior may provide the functionality that is required for the resolution. For example, an excellent tennis player will often hit a ball to a place that the opponent cannot reach, and therefore many return shots are predictable, given the locations and characteristics of both players. But sometimes the ball will be returned in unpredictable fashion, in a way that neither opponent nor observer can anticipate. Tennis is an example of a class of competitive game situations, described by Nash (1951) and others, in which random responding is *sometimes* the best path to success (Maynard Smith, 1982; Moreh, 1994). In a historical example, submarine captains in World War II sometimes used dice to determine the angle of turn of their ships to avoid enemy detection. Thus, as indicated by these examples, random behavior is sometimes emitted in pursuit of, and sometimes results in, desired or reinforcing outcomes (i.e., the randomness is functional).

Note that to be functional, responding must be unpredictable within a given (specifiable, or predictable) set and within a particular context. When we refer to random responding, we generally do not mean “do anything.” The unpredictable tennis-ball return falls within the confines of the court, the submarine’s trajectories include only those that minimize risk to the ship, and so on. Thus, most examples of unpredictable behaviors are predictable in some ways, and the ways generally involve a set of responses that are, at the given moment, functional—what Skinner referred to as a generic class.

In rare cases, the set of appropriate responses may be almost limitless, such as “call out any number,” or “utter a sound that can’t be predicted,” or “behave in a way that is completely unpredictable.” To make a point when discussing free will, a person might reasonably emit *any possible* behavior, and that occurrence will (we hypothesize) be exceedingly difficult to predict although being functional within the context. Thus, although we indicated that randomness (in general) does not solve the problem of free will, functional randomness contributes importantly to the solution.

There is another major problem with the contention that all voluntary behaviors are random. Voluntary behaviors are often quite predictable and, indeed, exist across the range of unpredictable to predictable, or unexpected to habitual. For example, when the traffic light turns red, the driver is likely to step on the brake. Most people use toothbrushes to brush their teeth in the morning.

When you are asked for your name, you generally answer veridically, and so on. But even in cases of predictable behaviors, if voluntary, these responses can be—and sometimes are—emitted in more or less unpredictable fashion. The red light can cause speeding up, slowing down, or cursing. The toothbrush can be ignored, or used with little or no toothpaste, or tossed in the air; and the name offered might be a first name, or last, or fabricated on the spot so as to fool the questioner—for example, when in a game situation. In brief, voluntary responses have the potential to move along a variability continuum, from highly predictable to unpredictable. We claim that a characteristic of all voluntary behaviors is real or potential *variations in levels of variability*.

Operant variability contributes to explanations of volition by combining its two most difficult-to-reconcile aspects. Voluntary responses are goal-directed and functional, and the same holds for operant behaviors. The functionality of the voluntary operant is given by the appropriateness of its emission, given a particular discriminative stimulus, and by the fact that the response leads to a reinforcing result. In some cases, we say that the voluntary response—and the operant—are “intended to be functional” because they are governed by *previous experiences* and because in a variable or uncertain environment, what was once functional may no longer be so.

Because variability is an operant dimension, the ability of a knowledgeable observer to predict behavior will depend upon the particular circumstances (i.e., the discriminative stimuli and reinforcement contingencies). This characterizes volition as well. The suggestion here is that operant and voluntary responses are sometimes predictable and sometimes unpredictable and that the levels of response variability (and therefore observer predictability) are functionally related to environmental demands. Implied is that organisms have the capacity to behave in a way that is functionally predictable or unpredictable, depending upon circumstances.

We suggest, therefore, that although previous attempts to explain volition have been deemed unsatisfactory by most, the term identifies a behavioral characteristic of signal importance. To draw a distant analogy, for thousands of years, the explanation of why the sun moves across the sky was in error, but descriptions of the phenomenon—sun apparently moving—were quite accurate. Just as the sun only appeared to move, so, too, voluntary responses appeared (to many) to emanate from an autonomous “self” residing within the skin, a rational and free-wheeling prime mover. We submit as an alternative that operant variability—and therefore the functional nature of behavioral predictability/unpredictability—describes an aspect of real-world behavior and that it goes far toward explaining what is perceived to be voluntary.

Such perceptions serve as a second source of support for the OVVA theory. We will provide evidence showing that human participants in fact identify voluntary actions by the behavioral characteristics described in the OVVA. Thus, we propose that real-world variations in levels of variability—variations that are related to environmental consequences—lead human observers to judge that the behaviors are voluntarily emitted.

The OVVA theory can be summarized as follows:

1. Voluntary acts produce (or are directed at producing) reinforcing consequences. In more common terms, a voluntary act is intentional and goal-directed.

2. Reinforcers and discriminative stimuli influence the *form* (or type, or class) of the voluntary act, including its topography, rate, force, etc.
3. Reinforcers and discriminative stimuli also influence the *variability* (or predictability, uncertainty, entropy, or stochasticity) of the voluntary act, with levels ranging from patterned or repetitive (and therefore easily predicted, as is the case for most habits) to randomlike (and therefore unpredictable).
4. Observations of a range of exemplars of (1) and (2) and (3) lead to attributions of volition.

Because the OVVA theory is in large part based on the operant nature of variability, we next review relevant evidence in that area of research.

Operant Variability

Reinforcement of Variability

This section reviews operant variability research that provides the foundation for the OVVA theory. Much (but not all) of the research has been done with animal models. For example, Page and Neuringer (1985) rewarded pigeons for pecking Left (L) and Right (R) response keys when a sequence (or pattern) of eight responses, emitted in a given trial, differed from the sequences in each of the preceding five trials. Thus, for example, if the pigeon had just responded LLLLR-RRR, and that sequence had not been present in any of the preceding five trials, then a food reinforcer was provided. If the current sequence repeated any one of the previous five, a short timeout followed. This schedule is referred to as a *Lag 5 variability contingency*, and it resulted in the birds' generating sequences of left and right responses that were highly variable and therefore that produced reinforcers on most trials. Many other studies reported similar results (see Neuringer, 2002, for review).

Before concluding that variability can be directly reinforced, the authors had to consider alternative interpretations. One possibility was that the variability was produced by some aspect of the situation other than the "if vary, then reinforce" contingency, such as a natural tendency to vary in such situations, a noisy environment, or the intermittency of food. As indicated previously, other studies have shown that infrequent reinforcement elicits behavioral variations, perhaps due to frustration or disappointment (Balsam et al., 1998). Was that responsible for the high level of variability, or was it directly reinforced? This question is especially important because some have argued that reinforcement necessarily decreases variability and therefore leads to increasingly stereotyped behaviors. If that were true, then variability could not be reinforced, at least over the long term, a result that would support the claim that reinforcement is incompatible with creativity, problem solving, or other cases in which variable or novel behaviors are characteristically observed. To control for these and other possible explanations, in another phase of the experiment we provided the pigeons with exactly the same reinforcers, at exactly the same intermittencies, but now independently of variability. Some eight-response trials continued to terminate with reinforcers and others with timeouts, but these were now based on, or yoked to, the

previous variability phase (hereafter referred to as Var). Thus, if an individual bird had been reinforced during the *n*th trial in a given Var session, then the analogous trial under the yoke condition would also end with a reinforcer, whether or not the lag contingency had been met. The goal, of course, was to test whether it was necessary to reinforce variability (in the present situation) for responding to vary, and the result was that the birds varied under the Var conditions but tended to repeat pecks on one or the other key in the yoke condition. That is, in the Page and Neuringer (1985) experiment, as well as many performed since then, variability was significantly higher when directly reinforced than when not—even if all other aspects of the situation were kept constant. These results support a claim that variability is an operant dimension, controlled by contingent reinforcers.

The generality of reinforcement-of-variability effects is shown by the many different procedures used to reinforce variability and the many species studied, including pigeons, song birds, rats, fish, monkeys, and humans (Neuringer, 2002). Animal model research has provided the opportunity for long-term evaluations across parameter values, but the results have generally been validated and extended with human participants. For example, Stokes, Lai, Holtz, Rigsbee, and Cherrick (2008) reinforced college-age students for responses that created a path through a triangular array in which a total of 128 paths were possible. When reinforcement was based on the current path differing from the previous five (Lag 5), levels of variability were higher than in the control phase, as in the Page and Neuringer (1985) experiment described before. In another example by Maes (2003), undergraduates chose from among three keys of a computer keyboard, with each trial consisting of three responses (thus providing 27 unique sequences). Reinforcement of low-frequency patterns caused sequences to vary more than when reinforcers were not provided (extinction) or provided independently of sequence variability (analogous to the yoke procedure described earlier). A third example comes from Miller and Neuringer (2000), who compared adolescents diagnosed with autism to adult and child controls under two conditions of a computer game: reinforcement of variability and reinforcement independent of variability. As in the studies previously described, response variability was higher when reinforced than when not. Although the participants with autism generally responded less variably than did both groups of controls, of considerable interest was the finding that their variability could be increased through direct reinforcement. Lee, McComas, and Jawor (2002) extended this research to reinforce verbal responses that varied appropriately from one instance to the next, a potentially important finding because lack of functional variation is a salient characteristic of autism. Thus, the research with animals and people shows that variability is higher when explicitly reinforced than when not and that the effects have considerable generality as well as potential application to real-world problems.

But voluntary behavior is more nuanced than vary or not. Rather, voluntary actions show much finer gradations of predictability. The semantic content of our everyday verbalizations, for example, are more or less predictable, depending on context. So, too, our choices (e.g., of what to eat, read, or how to recreate). For the OVVA theory to adequately explain voluntary behaviors, reinforcers must be shown to control *levels* of operant variability and to do so with precision. The next three subsections describe such precision of control.

Levels of Variability

As part of their experiment, Page and Neuringer (1985) varied the lag requirement (i.e., the number of preceding trials from which the pigeon's current sequence had to differ for reinforcers to be delivered). As lag values increased from 1 to 50, levels of variability increased and, at the higher lag values, approached the levels of a random model. Additional evidence for control by reinforcers over specific levels of variability was seen in an experiment by Grunow and Neuringer (2002) in which rats generated three-response sequences across three operanda—two levers, left (L) and right (R), and one response key (K). Each of four groups was reinforced for a different level of sequence variability. A computer tracked, in real time, the relative frequencies of the 27 possible sequences—LRK, LLR, KRL, and so on. For the high-variability group, the current sequence was reinforced only if it had previously occurred with a relative frequency (or probability) of no more than .037, that value referred to as the “threshold” criterion. (Reinforcement of low relative-frequency sequences is a commonly employed alternative to lag contingencies as a way to generate high response variability; see Denney & Neuringer, 1998). The value of .037 was chosen because a true random generator would distribute its responses approximately equally across three operanda with that relative frequency of each three-response pattern. Thus, the high-variability group was reinforced for approximating the distribution expected from a random model. The other three groups were required to satisfy successively less demanding thresholds: .055, .074, and .37, respectively, the last criterion being 10 times more permissive than the random model. Figure 1 shows that the high-variability group responded most variably, the low-variability group responded least variably, and

the other two groups fell at intermediate levels. Other studies have shown the same precise control over levels of variability in pigeons (Machado, 1997) and in people (Jensen, Miller, & Neuringer, 2006).

The Grunow and Neuringer (2002) study also demonstrated that reinforcement contingencies and frequencies interact to influence variability—as is the case for other operant dimensions—and do so in an informative way. In the initial phase of the experiment, the rats were reinforced each time they satisfied their respective threshold criteria (continuous reinforcement). In a second phase, satisfying these same variability contingencies led to reinforcement no more than once per minute (Variable Interval 1 min, or VI 1 min), and then in a third phase to no more than once per 5 min (VI 5 min). Thus, although the variability contingencies remained unchanged throughout the experiment—requiring high variability for one group, permitting low variability for another, and intermediate levels for two—frequencies of reinforcement were systematically lowered. The result was that, with decreasing reinforcements, levels of response variability decreased in the high variability group (.037), increased in the low group (.37), and remained approximately unchanged in the intermediate groups (.055, .074), a statistically significant interaction shown in Figure 1. A possible parallel to real-world behaviors is this: When low variability is reinforced, perhaps as in assembly-line work, little variability is observed. When high variability is reinforced, as in artistic or creative activities, the resulting product can't be predicted. What happens in each of these cases when reinforcers become less available? Variability of the assembly-line worker's output would be predicted to increase, whereas variability of the creative artist's would be expected to decrease. But, importantly, if the animal-model research is predictive, the artist would continue to behave much more variably overall than would the assembly-line worker (as was the case for the high- vs. low-variability groups of rats) even when reinforcers are sparse: The variability contingencies exert primary control. This preponderance of control by reinforcement contingencies, as opposed to reinforcement frequencies, is observed in other commonly studied dimensions of operant behavior as well, including response rates.

An extraordinary example of precise control was shown in an experiment by Ross and Neuringer (2002) in which college students were rewarded for simultaneously varying along two dimensions of response while repeating along a third dimension. The task was to draw rectangles on a computer screen, with one point being awarded following some rectangles but not others. No information was given to the participants about the fact that three dimensions of the rectangle were being evaluated, nor were they told about the reinforcement of variations and repetitions. The students were randomly divided into three groups, and all were instructed to simply gain points. One group was reinforced for rectangles whose sizes (defined by the area of the rectangles) were approximately the same (within a certain “delta” window), trial after trial, but whose locations (defined by the centroids of the rectangles) and shapes (defined by the ratios of height to width) both varied. A second group obtained points for drawing rectangles in approximately the same location on the screen while varying size and shape, and a third group was required to draw rectangles of similar shapes while varying sizes and locations. All participants learned to meet their respective three-part contingencies, varying and repeating as required (see Figure 2). Thus, binary feedback—




Figure 1. U value, an index of behavioral variability, as a function of reinforcement frequencies (CRF, or continuous reinforcement, given every time variability contingencies were met; VI 1, or Variable Interval 1 min, reinforcement for meeting variability contingencies no more than once per minute, on average; VI 5, reinforcement no more than once every 5 min). Each line represents a different group: .037 = very high variability required for reinforcement; .37 = very low variability required; and .055 and .074 = intermediate levels required. (Adapted from Grunow & Neuringer, 2002)




Figure 2. U values for each of three dimensions—area, shape, and location—as a function of whether participants were required to repeat the areas (left set of bars), shapes (middle set of bars), or locations (right set of bars) of their rectangles. In each condition, repetitions were required for one dimension and variations were required for the other two dimensions. Error bars indicate standard errors. (Adapted from Ross & Neuringer, 2002)

reinforcement presented or not—sufficed to independently influence variability and repetitions along three orthogonal dimensions, thereby again demonstrating the precise way in which reinforcers control the variability dimensions of operant responses.

Response Sets

Reinforcers also establish the set of responses from which acceptable variations emerge. That is, when variability is reinforced, the contingencies work simultaneously to define the set of possibilities and levels of variation from within that set. This is relevant because the types or classes of voluntary behaviors often change, moment to moment, depending upon contexts and contingencies. Evidence for response-set specification—together with reinforcement of variability—was seen in an experiment by Mook and Neuringer (1994). Rats were first rewarded for variable four-response sequences across L and R levers (lag schedule), following which only sequences that began with two right responses (RR) could gain reinforcement. Thus, now only RRLL, RRLR, RRRL, and RRRR patterns were effective. In the initial phase, all 16 possible sequences were emitted, whereas in the second phase, most sequences began with two right responses (RR). Thus, the reinforcement contingency generated behaviors that satisfied the “appropriate set” definition while simultaneously producing a required level of variability within that set. Additional evidence was seen in an experiment (Neuringer, Kornell, & Olufs, 2001) in which rats were reinforced in a chamber containing five operanda: left lever (LL), right lever (RL), left key (LK), center key (CK), and right key (RK). In an initial phase, reinforcers were contingent upon variations across only three of the operanda: LL, RL, and CK. The rats quickly learned to respond variably across these three “hot” operanda and generally to ignore the “off” operanda. Later, when reinforcers were withheld during a period of extinction, responses were distributed across all five operanda. Thus, extinc-

tion caused a broadening of the previously reinforcement-defined response set.

Discriminative-Stimulus Control

Voluntary behaviors are context-dependent, and research shows this to be the case as well for operant variability. For example, Page and Neuringer (1985, Experiment 6) reinforced pigeons for repeating a single sequence (LRRL) in the presence of blue key lights and for variable sequences in the presence of red. Blue and red alternated after every 10 reinforcements. The birds learned to repeat in the presence of blue and to vary in the presence of red, and when the stimulus relationships were reversed, the birds learned to reverse and vary in blue and repeat in red. In another experiment, rats learned to emit variable four-response sequences across L and R levers in the presence of one set of lights and tones, and to repeat a single pattern (LLRR) in the presence of a different set (Cohen, Neuringer, & Rhodes, 1990). In a more stringent test (Denney & Neuringer, 1998), rats were reinforced for varying under one stimulus whereas reinforcement was provided independently of variability under an alternating yoke stimulus. Recall that under yoke conditions, the rats were free to vary or not whereas reinforcement frequency was identical to that in Var. Nevertheless, the cues came to exert strong stimulus control, and when variability was required, the animals varied, but when variability was permitted (although not required), response sequences became much more repetitive and predictable. These results indicate that, consistent with the OVVA theory, an individual may behave in a habitual and predictable manner in one context, whereas in a different context, perhaps occurring only a few moments later, the same individual will respond unpredictably or in novel ways. The results further indicate that to maximize behavioral variations, one may need to explicitly reinforce variability rather than, as in a laissez-faire environment, to simply permit individuals the “free-

dom” to vary. This point is especially relevant to situations in which variations are adaptive, such as in solving problems, participating in creative activities, and attempting to break unwanted habits including addictions.

Responding Randomly

Behaviors that can't be predicted (e.g., possibly because they are randomly generated) hold a special place in many discussions of volition, in part because they imply (for some) an autonomous, within-person controller. A much-debated issue is whether indeterminate processes underlie (apparently) unpredictable behaviors. The question we next consider is whether operant responses can approximate stochastic models. If so, the possibility of “truly unpredictable” voluntary responding would be supported.

Much research has been directed at whether human participants can generate random response sequences, either verbally or on a computer keyboard (see Brugger, 1997, for review). The results consistently showed that when participants were simply asked to respond randomly—feedback was not provided—the resulting responses could readily be distinguished from random responses through any of a number of statistical evaluations. There was one notable exception in which feedback *was* provided following each set of 100 responses: Students received statistical feedback indicating how closely their response distributions matched those expected from a random model (Neuringer, 1986). At first, feedback was based on one statistical test and, after the students learned to satisfy that evaluation, then on two different tests, then three, and so on until they were receiving feedback simultaneously based on 10 statistical tests of randomness. Over many sessions and tens of thousands of responses, the students learned to distribute their responses in a manner that was statistically indistinguishable from a random model, at least according to the 10 tests used. In a self-experiment, the senior author learned to satisfy as many as 30 statistical tests, with many of these the same as used to evaluate computer-based random-number generators (Roberts & Neuringer, 1998). Note that in all experiments on human random generation, the responses were quite variable and individual responses were difficult to predict. It is only with respect to the strict criteria of a random distribution that feedback training was required: People are quite capable of highly variable responding without training. The ability to behave “truly randomly,” however, appears to be a skill, much as is the case for other operant behaviors, and one that depends upon the particulars of experience with consequences (see also Budescu & Rapoport, 1994).

Nonhuman animals also can learn to respond in a way that approximates a truly random output. In an impressive demonstration, Blough (1966) reinforced pigeons for pecking a response key with interpeck intervals that would parallel emissions of atomic particles. Blough established 16 interpeck intervals and reinforced a peck only if it fell within an interval currently containing the fewest peck instances. The durations or widths of these intervals were constructed such that if pecks were emitted randomly, they would fall equally often into the 16 intervals, as would be expected from the atomic emitter. Although the birds showed a nonrandom tendency to emit double pecks (two pecks with a short time between them), they otherwise approximated the “random peck” requirement. Blough's results showed that animals can be reinforced for approximating a random distribution in time, and his

work was followed by similar findings regarding response sequences (Machado, 1989).

Any claim that an organism can learn to respond randomly should, however, be considered with great care. It is one thing to show that response distributions parallel those from a random source and quite another to conclude that a random process is responsible for generating those responses (see Nickerson, 2002). One problem in testing the random-process hypothesis is that it is impossible to “prove” randomness: No matter how many tests of randomness are passed, there may always exist another test that would distinguish between the response sequence and that expected from a random source. Another problem is that nonrandom processes can produce randomlike outputs, such as iterating the digits of π (Bailey & Crandall, 2001).

More generally, chaos theory describes ways in which random-appearing outputs are generated by deterministic processes that are highly sensitive to “initial conditions.” That is, chaotic outputs (including behaviors) are controlled by prior events and rely on a kind of “memory” that utilizes prior responses to compute current ones. If an observer knew the generating function and history of prior responses, then the next response could be accurately predicted. Neuringer and Voss (1993) demonstrated that human participants could in fact learn to respond in a way that parallels one chaotic function, namely the logistic difference function (however, see Ward & West, 1994, for an alternative interpretation). Thus, the question is not whether highly variable responses *can* be chaotic: Many studies have suggested an affirmative answer (Robertson & Combs, 1995). Rather, it is whether processes that yield predictable outputs are responsible for the variable behaviors generated when reinforcers are contingent upon highest levels of variability. Put in general terms, are instances of an operantly varying sequence caused by identifiable prior events, including prior responses, or do they occur randomly and independently of those events? This issue is especially important for those who argue in favor of indeterministic underpinnings of free will. We will describe a number of tests, each of which provided results that are consistent with an indeterministic, “random-generator” model.

Page and Neuringer (1985), as part of the pigeon experiment described earlier, tested random generation by manipulating trial length. In one phase of the experiment, four-response sequences were reinforced under a Lag 3 variability contingency; in a second phase, the number of responses per trial was increased to six, with the lag remaining the same; and in a third phase, to eight, again with the same lag contingency. The question was whether the birds used memory of past sequences to respond in the current trial (e.g., used previous trial patterns as discriminative cues to emit a “different” pattern in the current trial) or responded randomly across the two keys (and therefore independently of previous trials). A memory hypothesis predicted that increasing trial length might reduce success rates (as indicated by the frequency of reinforcement), because memory for 24 previous responses (8 responses per trial under Lag 3) is more demanding than for 18 (6 responses per trial), and that, in turn, is more demanding than for 12 (4 responses per trial). The random-generator hypothesis predicted the opposite result: Increases in trial length would be expected to increase the success rate because degrees of freedom were greater under 8 responses per trial than 6, and greater under 6 than 4. (If a true coin is tossed once, then the chances of repeating the previous toss is .5; if the coin is tossed four times, the chances of repeating the

previous four tosses is .0625.) The results were clearly consistent with the random-generator hypothesis: As responses per trial increased, the pigeons were reinforced more and more frequently, exactly as is the case for a simulated random model.

But these results could have been produced by “noise” in the system (whether stemming from extraneous cues or “mistakes” made by the participants) rather than emanating from an endogenous random source. A noisy environment, for example, would be expected to produce higher variability—and therefore more reinforcement—under the 8-responses-per-trial condition than under the 6- or 4-responses-per-trial condition. Jensen et al. (2006) therefore studied the same question but maintained a constant number of responses per trial. Pigeons were reinforced for variable responding across 2, 4, or 8 different response keys in three phases of an experiment. As the number of response operanda increased, trial lengths were decreased—from 6 to 3 to 2—so as to maintain a constant number of different possible patterns across the phases (i.e., 64 possible response sequences, because $64 = 2^6 = 4^3 = 8^2$), effectively keeping the influence of statistical noise constant. The results were again consistent with a random-generator hypothesis.

The random-generator hypothesis was tested in a third way by interfering directly with the possible influences of previous responses. This was accomplished by increasing the time between individual responses in each sequence, thereby slowing responding and presumably interfering with memory for prior occurrences (Neuringer, 1991). Two conditions were compared, one in which we hypothesized that memory would be involved and the other a reinforcement-of-variability condition. We reasoned that if response interference were similar in the two cases, it would favor a memory-based process to account for responding under the variability contingencies and against the random-generator hypothesis. In particular, one group of rats was reinforced for repeating a single pattern (LLRR). Once the pattern was well learned, an interresponse timeout period (IRT) was introduced between consecutive responses, the duration ranging from 0.1 s to 20 s across different phases of the experiment. We hypothesized that the increasing IRTs would adversely affect emissions of the LLRR

sequence, presumably because each individual response served as a cue for the next response. That was indeed the result: Figure 3 shows that as timeout durations increased, errors increased and reinforcement rates fell. Thus, interfering with memory for prior responses degraded fixed-pattern, LLRR responding.

A second group of rats was reinforced for varying sequences of four responses under lag contingencies. Again, if their variable responding were memory based, then the same effects should be found as with the LLRR group, namely that performances would increasingly suffer as IRTs increased. In fact, as the timeout duration increased across the same range, performances by the variability group actually improved (i.e., the rats were reinforced more frequently; see Figure 3). Some have suggested that absence of memory for prior responses is necessary for random responding (Weiss, 1965), implying that memory actually interferes with random generation. In any event, the results were clearly inconsistent with a memory hypothesis to account for operant variability.

In a related study, alcohol was administered to rats that had learned to respond variably under one stimulus condition and repeat LLRR sequences under a second set of stimuli (Cohen et al., 1990). The two stimuli alternated throughout each session under what is referred to as a “multiple schedule.” Figure 4 shows that with increasing alcohol dose, rats’ ability to repeat an LLRR sequence was seriously impaired, whereas their ability to meet a lag variability criterion was unaffected (see also Doughty & Lattal, 2001). These effects were all the more impressive because, within a single session, the “drunk” rats were quite unsuccessful when the conditions indicated “repeat” but were highly proficient when they indicated “vary.” Thus, both interposed time delays and alcohol, two ways to affect memory for prior responses, degraded performances of fixed-pattern sequences and either improved or left unaffected operant variability. With reference to volition, these results suggest that drunken individuals may behave in unusual or variable ways—alcohol does not inhibit variability—but that they will not adequately vary *levels* of variation to meet current demands. Here, and throughout this article, it is not variability per se,




Figure 3. Percentages of correct (or reinforced) trials as functions of interposed interresponse intervals (timeouts) in one group of rats reinforced for varying response sequences (left panel) and another reinforced for repeating a single LLRR (L = left response, R = right response) sequence (right panel). Lines connect group arithmetic means, and vertical lines show standard errors. (Adapted from Neuringer, 1991)




Figure 4. Percentages of reinforced, or correct, sequences as a function of ethanol dosage for each of five rats (BG, BR, RG, B, and P). The same rats were reinforced for varying sequences under one stimulus condition (left panel) and for repeating LLRR (L = left response, R = right response) under another set of stimuli (right panel). The lines connect averages of the five subjects. (Cohen et al., 1990)

but rather *functional variations in levels of variability* that identifies volitional behavior.

The just-reviewed studies are consistent with the claim that operant responses are emitted randomly in situations where randomness is reinforced. Therefore the evidence supports, by implication, the possibility of truly unpredictable and adaptive voluntary behaviors.

(Un)predictability Within Response Classes

The experimental evidence discussed to this point has been based on sequences of responses or, as in Blough's (1966) experiment, on interresponse times. We next review experiments showing that reinforcers also control the predictability/unpredictability of qualitative instances that are members of a class or category of appropriate responses. In one case, the class consists of different physical movements; in another, the construction of different objects in space; in a third, different members of a linguistic category; and in a fourth, different locations on visual drawings. These experiments add to the generality of control by reinforcers over response variability/predictability.

Novel Behaviors

Pryor, Haag, and O'Reilly (1969) rewarded porpoises for novel flips, swims, turns, and the like, and eventually some of the porpoises generated behaviors that qualitatively differed from any that had previously been observed in that species. In an extension to humans, Goetz and Baer (1973) rewarded preschool children for block constructions that differed in form. Across training, the children generated increasingly varying forms, including ones never before emitted by these children. When reinforcement was later made contingent upon repeating a single form, the children satisfied the repetition contingency as well, thus indicating that operant reinforcement was responsible for the novelty. Similar results were obtained with drawings. Thus, low probability, or

novel, drawings and constructions were engendered by contingent reinforcers (see Goetz, 1989).

Learning to emit topographically varying responses may have some generally beneficial effects. For example, Arnesen (2000) studied whether reinforcing rats for novel interactions with objects in one context would facilitate problem solving in a different situation. In a training stage, the rats were reinforced with food pellets for interacting variably with objects in an experimental chamber. Each session (or few sessions in some cases), the rats experienced a new object. In one case, for example, a tin soup can was the object, and reinforcement followed the first interaction (e.g., a touch), but then a different topographic response was required (e.g., poking the nose into the open can or rolling the can). The individual rats engaged in different types of interactions and in different orders, but all were reinforced for novelty of such interactions. Ten different objects were experienced during a series of reinforcement-of-novel-interactions training sessions, following which each rat was separately tested in a new environment. For this test, the rat was placed in a large enclosure containing 30 objects that were different from those in the training sessions and ones that the rat had never before experienced. Hidden in each object was a small piece of food. The problem, therefore, confronting the hungry rats was to discover the food. The number of test objects explored and number of food items discovered were recorded. Two control groups received exactly the same test but they had different prior experiences. One control had spent exactly the same time as the experimental animals with the 10 training objects but without explicit reinforcement for novel interactions. Thus, these rats had ample opportunities to interact with the objects but were not explicitly reinforced for doing so and, in fact, did not demonstrate the same high level of interaction as did the experimental group. The other control group had simply been handled for the duration of the training periods. Differences during the test sessions were large and statistically significant: The experimental animals explored many more of the objects and found

many more food items than did either of the control groups. Many of the control rats hovered close to the walls of the test room, an indication of fear. Thus, reinforcement of novel interactions facilitated exploration and discovery in a different situation. These results suggest that at least under some circumstances, when an individual learns to respond variably in one environment, that competency may generalize to other related situations. Of course, as with all cases of generalization, parameters are important (Goetz, 1989).

Variations Within Linguistic Categories

Human verbalizations demonstrate an immense range of variation. Verbal responses to requests for instances of a particular category can vary from predictable (analogous to stereotyped response patterns or repeated interresponse times) to highly improbable and unpredictable (analogous to randomlike sequences or latencies), depending upon contexts and contingencies. Previous research showed that when individuals are asked to generate an instance of a category (e.g., “flowers”), responses vary, both across individuals and across time in the same individual (Barsalou, 1987). Such responses can readily be organized along a common (or typical or predictable) to uncommon (or atypical or unpredictable) continuum (e.g., Rosch, 1978; Voorspoels, Vanpaemel, & Storms, 2008). Context, prior experiences, and emotional states influence the degree to which a response is a common versus less common exemplar (Barsalou, 1987). Do reinforcement contingencies exert similar, and possibly more precise, influence? We used the research on graded category structures to test whether individuals could generate instances of categories at *specific levels of predictability*. Despite the large body of research on within-category structures, we know of no study in which people are asked explicitly for instances at particular locations along a predictable-to-unpredictable (or common-to-uncommon) continuum. The study attempted to explore operant predictability in what is generally thought to be a uniquely human cognitive domain and will be described in some detail because the data have not previously been published.

Undergraduate participants were instructed to imagine that many individuals had been asked to name a single instance of a category, such as animals or fruits. Our participants ($n = 89$) then estimated three instances of the category: one instance that was “very likely” (VL) to have been provided by members of the imagined group, another that was “only somewhat likely” (OSL), and a third that was “very unlikely” (VU). For example, if a person were asked to imagine that a number of his or her colleagues and friends had been asked to write on slips of paper the name of a sport, he or she could then provide what would seem to be the most common response given; a somewhat common response; and an unlikely, rarely emitted response. In our instructions we presented, as an example, the category A Mode of Transportation and suggested that “car” might be an appropriate VL response, “rowboat” an appropriate OSL response, and “hot air balloon” an appropriate VU response. Although we were interested in the participants’ ability to generate instances along the predictability continuum, the task was phrased in terms of estimating the responses of a group of other people in order to avoid individual differences based on participants’ thinking of themselves as quirky or unusual, or trying to outwit the experimenters, and so forth.

Participants were asked to provide VL, OSL, and VU responses for each of the following categories:

- An Animal
- A Fruit
- A Thing to Eat on a Diet
- A Thing Someone Might Do During Lecture
- A Thing That Might Fall on Your Head

Results. Following the exclusion of fewer than 5% of the responses—due to not being members of the requested category (e.g., a brick is not a diet food) or being nonsense words—and the correction of spelling errors, the responses were ranked in descending order of frequency across the 89 participants, and the frequency of each response was then plotted as a function of that response’s ordered rank. This was done separately for VL, OSL, and VU estimates for each of the five categories. The results were well described by the following function:

$$p(i) = c \cdot \left(\frac{1}{i}\right)^d, \text{ where } \sum_{i=1}^n p(i) = 1. \quad (1)$$

Equation 1 shows that the relative frequency, or proportion $p(i)$ of total responses for a given item (e.g., “dog”) was a function of its rank i among n items. Two free parameters were used to fit the data: c , a scaling constant, and d , the rate at which frequencies decreased as rank increased. In this situation, c is not of interest, because it serves merely to scale the function; the main variable of interest is d , which provided an index of the diversity (or variability) of the answers. The larger the value of d , the more homogeneous the responses (and thus the more predictable); correspondingly, the smaller the value, the more variable. This equation closely resembles Zipf’s law, first observed in philology (Zipf, 1935); it describes the distribution of word frequencies in a variety of contexts.

Figure 5 shows the Equation 1 fits for the responses generated when “Animals” was the category. The figure shows that responses for VL likelihood (left panel) were far more predictable and more often repeated (i.e., more sharply skewed) than were responses for the OSL likelihood (center panel), which were in turn more predictable than responses for the VU likelihood. Table 1 provides the d parameters for the three likelihood estimates in each of the five linguistic categories and shows a similar pattern of increasing variability (or uncertainty) as participants named increasingly unlikely members of each category. That is, d was consistently larger for VL than for OSL, and in four of five cases larger for OSL than for VU.

Table 2 provides the top-ranked responses for two categories: the Fruit category (which had the highest d values) and the Thing That Might Fall on Your Head category (which had the lowest). In both cases, despite the broadness of the category and their qualitatively different characteristics, participants agreed to a notable extent. Additionally, it was not merely the range of the responses (as captured by the d parameter) that changed with likelihood: The top-ranked members also changed across the three likelihood subcategories. The same patterns of data were observed for the other categories.

These results show broad agreement among participants as to which instances are appropriate to each of the VL, OSL, and VU

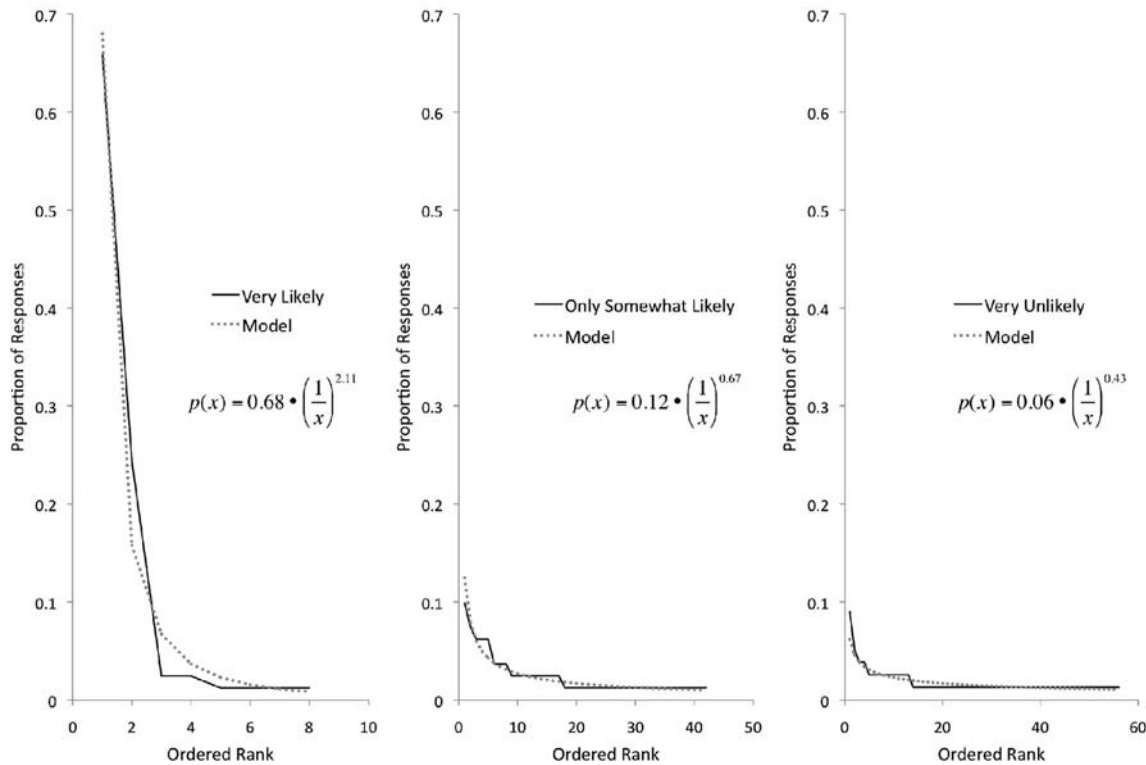


Figure 5. Proportion of total responses as a function of rank when participants were asked to provide very likely (left panel), only somewhat likely (middle panel), and very unlikely (right panel) examples of animals. The dashed lines show best fitting Zipf functions.

subclasses. This agreement indicated that the responses functionally satisfied the likelihood criterion. In addition, variability of responses generally increased across VL, OSL, and VU subcategories. Thus, not only did the types of responses change, but so too did the size of the activated sets across the three likelihoods. A third finding was that the same patterns of increasing variability were obtained whether the categories referred to well-defined objects (such as fruits) or ad hoc groupings (such as an object that might fall on one's head). Thus, the likelihood, or predictability, of verbal responses came under the same type of discriminative control as demonstrated for response sequences and latencies. This facility to vary levels of predictability is consistent with the volitional control that is a salient characteristic of verbal behavior.

Table 1
Decay Parameter (d) as a Function of Likelihood and Category

Category	Very likely	Only somewhat likely	Very unlikely
An Animal	2.107	0.667	0.432
A Fruit	2.670	0.901	0.861
A Thing to Eat on a Diet	1.247	0.723	0.758
A Thing Someone Might Do During Lecture	1.705	0.747	0.607
A Thing That Might Fall on Your Head	1.182	0.915	0.640

Variations Within Visual Categories

In a second procedure we tested for controlled variations in a visual category task in which the same participants were asked to

Table 2
Five Most Frequent Responses (With Percentage) to the Categories A Fruit and A Thing That Might Fall on Your Head, for Each Likelihood

Likelihood and rank	A Fruit	A Thing That Might Fall on Your Head
Very likely		
1	Apple (.91)	Rain (.61)
2	Orange (.04)	Leaf (.06)
3	Banana (.02)	Apple (.05)
4	Strawberry (.01)	Book (.05)
5	Starfruit (.01)	Rock (.04)
Only somewhat likely		
1	Tomato (.10)	Leaf (.19)
2	Mango (.10)	Bird droppings (.14)
3	Banana (.09)	Book (.10)
4	Kiwi (.09)	Snow (.06)
5	Pear (.08)	Piano (.05)
Very unlikely		
1	Starfruit (.17)	Piano (.08)
2	Tomato (.12)	Anvil (.05)
3	Guava (.09)	Airplane (.05)
4	Kumquat (.08)	Brick (.04)
5	Pomegranate (.06)	Apple (.04)

select a position on a pictorial image that would be very likely, only somewhat likely, or very unlikely to be chosen by a large group of individuals who were simply asked to click once on the image. Five different pictures were shown, and each participant provided three likelihoods for each of the pictures. The five pictures were as follows:

- A Blank Field
- An Abstract Shape
- A Face
- A Block of Text
- A UFO Hovering Over a Landscape

Figure 6 (left column) displays two of the images: A Blank Field (top row) and A UFO Hovering Over a Landscape (bottom row). The remaining columns display the pattern of clicks made by all participants superimposed over a light-gray version of the image, in which the left-middle column corresponds to the VL likelihood, the right-middle corresponds to the OSL, and the right to the VU.

A visual examination of the patterns of clicks in all cases showed a result similar to that for the linguistic-category task: Participants appeared not only to widen the range of clicks as the likelihood decreased but also shifted which “landmarks,” present or implicit, were focused on. In the case of the UFO, for example,

VL responses were chiefly focused on the UFO itself and a beam of light descending from it. By contrast, the OSL responses shifted away from the beam and toward other landmarks (such as the house and the clouds), whereas the VU responses avoided the UFO entirely and (to the extent that they clustered at all) tended to cluster around background landmarks such as the mountains.

An objective measure of dispersion was obtained by subdividing each picture into a 9-by-9 grid. Responses to each of the 81 cells were then ranked in the same way as responses in the linguistic task. We again used Equation 1 to estimate the relationship between proportion of clicks to a given cell (analogous to frequency of verbal instances) and its ordered rank. As in the linguistic task, values for d were lower in the VU condition than in the OSL, and likewise lower in the OSL than in the VL (these values are shown in Table 3). In other words, responses became less predictable as the requested likelihood decreased.

Linguistic and visual tasks therefore demonstrate that human participants have a ready facility to move along the predictability–unpredictability continuum and do so in an appropriate, or functional, manner. The types of responses and their predictability were both influenced by requests for different levels of predictability. Thus, as in the response-sequence tasks described earlier, levels of within-category predictability changed as a function of discriminative stimuli—in this case, instructions.

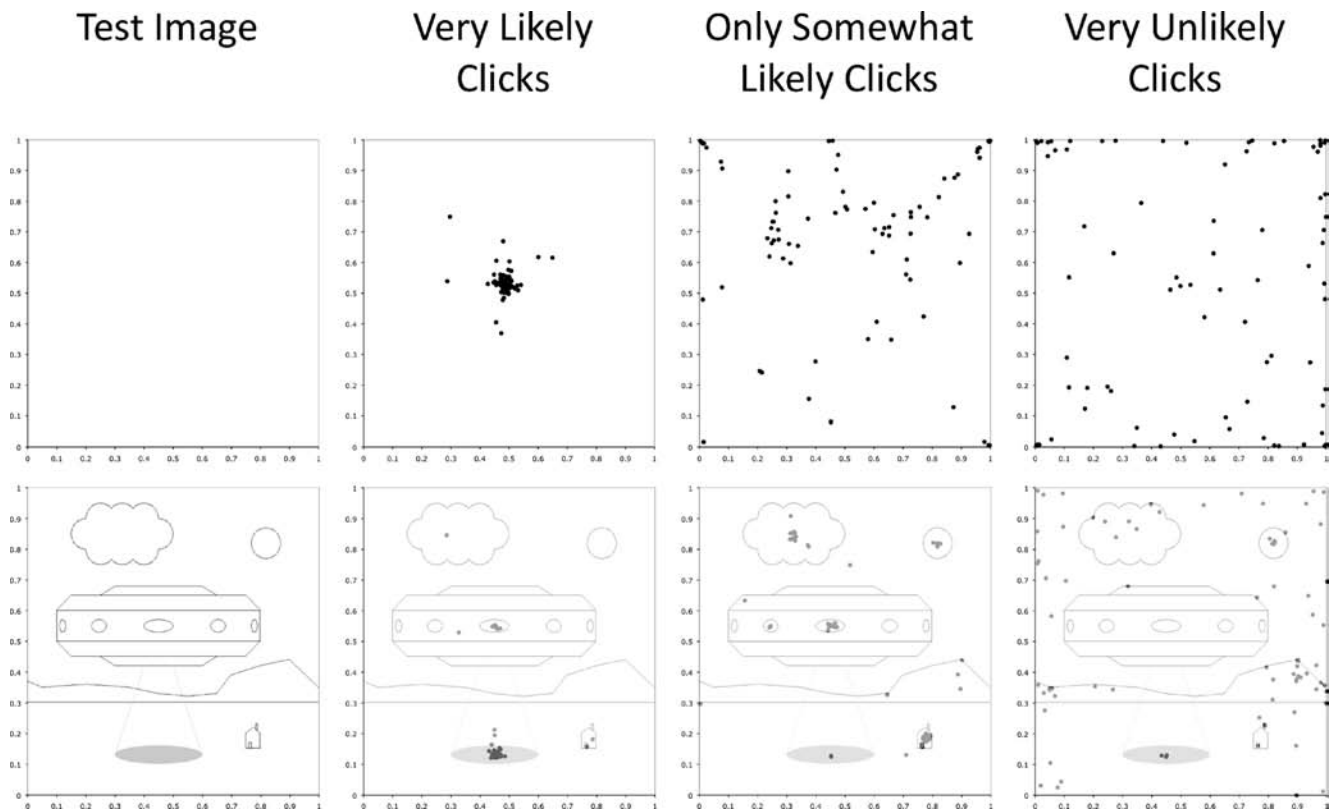


Figure 6. Responses in the visual task to the blank field (upper row) and the UFO picture (lower row). The left-most column shows the displays as presented, the next column shows the very likely responses, the next the only somewhat likely responses, and the right-most column the very unlikely responses. Each point represents a click on the computer image by one of the 89 participants.

This document is copyrighted by the American Psychological Association or one of its allied publishers. This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Table 3
Decay Parameter (*d*) as a Function of Likelihood and Image

Image	Very likely	Only somewhat likely	Very unlikely
A Blank Field	2.0166	0.8398	0.6950
An Abstract Shape	1.9745	1.0423	0.5615
A Face	1.6403	1.1965	0.6226
A Block of Text	1.2462	0.7809	0.6943
A UFO Hovering Over a Landscape	2.1288	1.2770	0.7544

Choice

Another behavioral competency closely associated with volition is choice. As will be seen, under some circumstances an individual choice may be easily predicted, but under other conditions predictions will be no better than chance. That is, choice allocations vary across the levels-of-predictability continuum, much as with the other cases previously discussed. Many interpretations of choice have been offered, including the necessity of prior conscious deliberations or the presence of a central supervisor or controller (in the mind or brain). A different way of conceptualizing choice is provided by analyses of behavior–environment relationships under concurrent schedules of reinforcement (Davison & McCarthy, 1988; Williams, 1988). The literature on “concurrent choice” provides evidence consistent with the OVVA theory and also suggests a way to test the theory. We will first describe the basic procedures and results from concurrent-choice studies and then describe an experimental test of the OVVA theory.

In experiments employing concurrent-reinforcement schedules, animals or human participants are generally provided with two options. Frequencies of reinforcers often differ for the two, and these are systematically changed across phases of the experiment. The main question concerns the relationship between distribution of choices—or relative frequencies of responses for the two options—and distribution of reinforcers. An example is provided by a concurrent VI 1 min: VI 3 min schedule, in which VI stands for “variable interval.” Under this schedule, a reinforcer becomes available, or “sets up,” unpredictably on the average of once every 1 min for responses on the left, and a reinforcer sets up independently on the average of once every 3 min, again unpredictably, for choices of the option on the right. Once a reinforcer has set up, it remains available until collected, much as with mail delivered to a home mailbox or money deposited in a (secure) bank account. In the behavioral studies in question, the exact time at which a given reinforcer becomes available is uncertain, as is the location of the available reinforcer (i.e., whether a reinforcer has set up for a left choice, a right choice, or neither). Two general results are most relevant to the OVVA theory. First, ratios of left to right choices change as a function of ratios of obtained reinforcers, a relationship commonly described as a power function and referred to as the *generalized matching law* (Baum, 1974), as seen in the following equation:

$$\frac{C_X}{C_Y} = \left(\frac{k_X}{k_Y} \right) \cdot \left(\frac{R_X}{R_Y} \right)^s \quad (2)$$

In Equation 2, *C* refers to choices of *X* and *Y* alternatives, *k* to bias toward or away from *X* relative to *Y*, *R* to reinforcers, and *s* to

the sensitivity of choice ratios to reinforcement ratios. To the extent that the generalized matching law provides an accurate description (and there is much support for it), predictions are possible at a “molar” level of choice allocation (i.e., overall ratios of choices can accurately be described as a function of obtained reinforcer ratios; Davison & McCarthy, 1988). Note that Equation 2 is used to describe a wealth of actual choice data (i.e., the function is used to fit the data obtained over hundreds of studies). As will be explained in the Psychophysical Test of Perceived Volition section, the *s* exponent, often referred to as the *sensitivity* parameter, indicates the extent to which response ratios change as a function of reinforcer ratios, and it will play an important role in our test of the OVVA.

A second relevant result is that individual choices are difficult to predict: They often appear to be emitted stochastically (Glimcher, 2003, 2005; Jensen & Neuringer, 2008; Nevin, 1969; see also Silberberg, Hamilton, Zirrax, & Casey, 1978, for an alternative view). In the previous example given, an observer might accurately predict that the left option will be chosen three times more frequently than the right but be unable to accurately predict any given choice. A recent example of such stochasticity was observed when pigeons were reinforced concurrently for choices among three keys. Figure 7 shows that run lengths—defined as the average number of choices on one key prior to switching to a different key—approximated those expected from a stochastic process (see Jensen & Neuringer, 2008, for a discussion of these findings, including the small divergence of data points from the theoretical curve).

The concurrent-reinforcement studies discussed to this point provided discrete-choice options (such as left key vs. right key). A different approach defined a choice in terms of strings of responses, classified according to some criterion. In one such experiment a computer evaluated four-response sequences emitted by




Figure 7. Logarithms (base 10) of mean run lengths by pigeons on each of three response keys as a function of the proportion of responses to that key. The drawn line is the expected function if responses were stochastic. (Adapted from Jensen & Neuringer, 2008)

pigeons as meeting either a “vary” criterion or a “repeat” criterion (Neuringer, 1992). Vary and repeat choices were reinforced in a way similar to the concurrent schedule just discussed, and the results were also similar: The ratios of vary to repeat emissions were described by the generalized matching function. Figure 8 shows that as reinforcement frequencies increased for varying, responses increasingly varied. Thus, the choices “to vary” or “to repeat” were governed by the same reinforcement relationships as discrete-response choices.

Much has been written about possible reasons why animals and people sometimes match choice proportions to reinforcer proportions (Williams, 1988). One consequence of such “strict matching” (in which choice ratios exactly equal reinforcer ratios, a relationship that will be discussed later) is that the effort expended per reinforcer is equalized across choice alternatives, and therefore there is a balancing of effort (output) to reward (input). A second consequence is that matching often leads to maximizing the overall frequency of reinforcement.

Less has been written about reasons for stochastic emission, but possibly important are avoidance of local maxima (where repetitive responding might cause an organism to forgo more highly reinforcing possibilities); exploration of the physical or problem space; or, when in a social context, competition with other organisms for limited rewards or protection against attack or predation. For present purposes, it is sufficient to note that the concurrent literature shows that choice proportions are power functions of reinforcer proportions and that choices are emitted stochastically when rewards are uncertain. These results can be summarized by the phrase “stochastic matching of choices to reinforcers,” a phenomenon that will be used to test the OVVA theory in the psychophysical procedures to be described next.

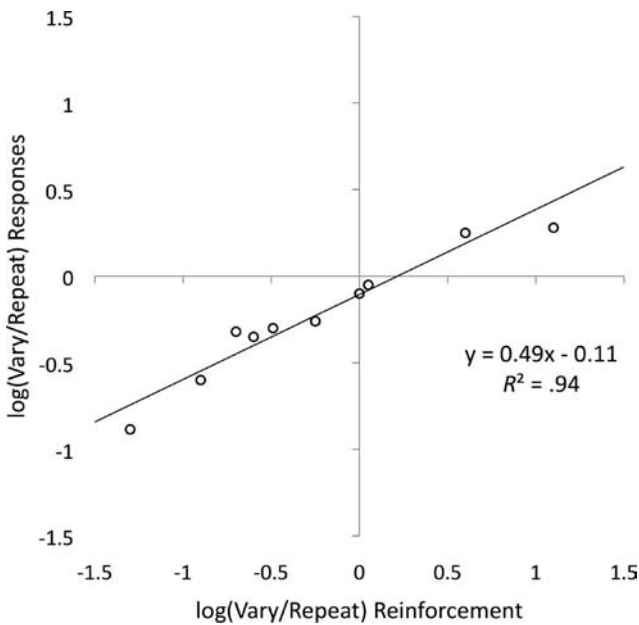


Figure 8. Logarithms (base 10) of the ratios of variable sequences (Vary) to repeated sequences (Repeat) as a function of the logarithms of the ratios of Vary to Repeat reinforcements. Averages of the six pigeon subjects' performances are shown. (Adapted from Neuringer, 1992)

Psychophysical Test of Perceived Volition

An implication of the OVVA theory is that a behavior will appear to be voluntary when its predictability varies—possibly across a wide range—in ways that seem functional (meaningful, purposeful, adaptive, or reinforced). A combination of (un)predictability and functionality will indicate volition. The concurrent-reinforcement literature provided an experimental procedure that enabled us to test this claim. Human participants observed computer screens on which several different virtual actors made thousands of choices (Neuringer, Jensen, & Piff, 2007). The choices were represented as dots moving around the screens so as to minimize other cues (e.g., whether the actor resembled a human). The virtual actors were said to be choosing from among three alternative gambles, similar to slot machine gambles, each represented by the dot's movement in one of three directions. Whenever a choice was successful, reinforcement was indicated by a change in the dot's color. Participants were asked to judge how well the dot represented the voluntary choices made by a human player.

In the first of a number of experiments, participants observed six different actors and estimated, on a scale of 0 to 100, the degree to which the responses represented the voluntary choices of a real, human player. The actors differed only in how they allocated their choices, these being controlled by algorithmic iterations of the generalized matching power function shown in Equation 2 extended to a three-alternative choice situation (Jensen & Neuringer, 2008). In turn, the algorithms differed in only one aspect, namely the value of the s exponent that controlled the sensitivity of response allocations to reinforcer distributions. (Bias across the three options was set to be equal and therefore did not play a role in the outputs.) We will describe the way in which reinforcers were programmed and then the strategies used by the different actors. None of the following details were provided to the participants, who were told only that the movement of the dots on the screen may represent (more or less—and they were to judge that) the voluntary choices of human players.

Reinforcers were programmed (set up) probabilistically, and once set up they remained available until collected, as described in the previous section on concurrent choice. The set-up probabilities changed across phases of the experiment, and they generally differed across the three choice alternatives (see Neuringer et al., 2007, for details). Participants observed six different actors, each choosing under six different sets of reinforcement probabilities, these referred to as different “games.” After observing all actors, the participants judged the degree to which each actor's responses appeared to be those of a human player who was making voluntary choices.

As indicated, the actors differed only in the values of the sensitivity exponents employed by the algorithms that governed their choices. The effects of these different exponent values were as follows: For one actor, s value equaled 1.0, and its choice ratios “strictly matched” obtained reinforcer ratios. Assume, for example, that this actor had gained a total of 100 reinforcers at some point in a session: 50 reinforcers for option X, 30 for Y, and 20 for Z. The probability of the next X choice would therefore equal .5 (50/100), a Y choice .3 (30/100), and a Z choice .2 (20/100). The $s = 1.0$ actor therefore distributed its choices probabilistically in exact proportion to its received reinforcers.

Another actor had an s value of 0.4, the consequence of which was that it tended to respond with probabilities that were more equal across the three alternatives throughout the six experimental phases than indicated by the reinforcement ratios. In the example given, the $s = 0.4$ actor would choose X with a probability of .399 (rather than .5 for the exact matcher), choose Y with a probability of .325 (rather than .3), and choose Z with a probability of .276 (rather than .2). In general, algorithms with s values < 1.0 are referred to as *undermatchers*: They distribute choices across the available options more equally than does the exact matcher. The opposite was the case for actors with s values > 1.0 , whose preferences were more extreme than indicated by the reinforcer ratios and were referred to as *overmatchers*.

In the first experiment, participants were informed in advance that all of the actors' choices were generated by computer algorithms, and they were asked to rate the algorithms in terms of volitional appearance. In another experiment, participants were told that some actors' choices were based on computer algorithms, that others' choices actually depicted voluntary choices made by real humans, and that their task was to "identify the humans," a kind of Turing test. Here is a portion of the instructions read by participants in the latter case:

There is a common intuition that people will know voluntary behavior when they see it . . . In this study, we wish to establish a baseline measure of how well viewers (such as yourself) can discriminate a voluntary human player from an artificial one. To this end, we brought in a group of individuals who played a simple gambling-type game. Their performance was used to determine their wage . . . We also designed a set of computer algorithms that were also able to play the game, and used them to generate a set of nonhuman game play records . . . (You should) attempt to decide whether the choices represented across the 6 games represent voluntary choices made by a real human player or not. (Neuringer et al., 2007, p. 26)

Across all experiments, effects were large, consistent, and statistically significant: The strict matcher ($s = 1.0$) was judged to

best represent volitional choices. When observing actors they knew to be algorithms, participants gave the strict matcher the highest ratings, and when asked to identify which actor represented the voluntary choices of a real human player, they again identified the strict matcher a higher percentage of the time than they did the under- or overmatchers. Figure 9 shows ratings from both of these experiments—different participants were involved in the two experiments—and shows the high degree of concordance in the judgments. Sensitivities closer to 1.0 were rated higher than those farther from 1.0, suggesting a continuum of more or less apparent volition.

A noteworthy aspect of these results is that maximum unpredictability did not generate highest judgments of volition: Undermatchers ($s < 1.0$), actors who tended to choose among the alternatives equally and therefore most unpredictably, were judged not to represent voluntary human choices, despite their choices being less predictable than those of the strict matcher. Responding randomly under all circumstances does not represent voluntary behavior. Similarly, low judgments were made for the overmatching actors ($s > 1.0$), who tended to fixate on one alternative and whose choices tended to be relatively predictable in most games.

These results are consistent with the OVVA theory: From the perspective of our experimental participants, the $s = 1.0$ strict matcher sometimes responded unpredictably (when reinforcers were equally allocated across the three alternatives), other times highly predictably (when most reinforcers were obtained from one alternative), and at yet other times at intermediate levels. In each case, however, performance seemed to be related to the reinforcement distribution generated in a particular game environment, an indicator of functional changes in behavior. Thus, functional variations in levels of predictability indicated volitional choice. However, other interpretations are possible, and these were tested with a series of control procedures, again with different groups of participants.

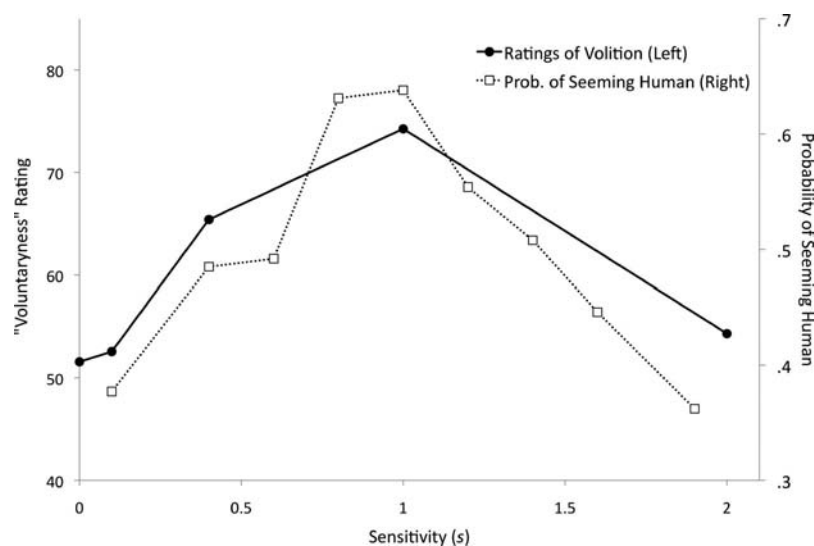


Figure 9. Judgments of how closely actors' responses approximated voluntary human choices (on left y-axis) and probabilities of identifying actors as a voluntarily choosing human player (on right y-axis) as functions of the actors' s -value exponents. Prob. = probability. (Adapted from Neuringer et al., 2007)

One alternative explanation derives from a comparison of the frequencies of reinforcers obtained by the different actors. In some game contexts, the overmatchers were most frequently reinforced (contexts in which almost all of the reinforcers were programmed for one of the choice options). In other games, the undermatchers received the most reinforcers (contexts in which reinforcers were equally distributed across options). And in yet other games, the matcher excelled. Averaged across all of the games, however, the strict matcher's performance somewhat outpaced the other actors', and therefore a possible interpretation was that the actor who was most frequently reinforced overall was seen as best representing voluntary choice.

Although the differences in the overall frequencies of obtained reinforcers were in fact relatively small, to test the possibility that they were responsible for the volitional judgments, we had participants compare a strict matcher (same as in the experiment described previously) with two other actors who appeared to have some degree of "knowledge" concerning the availability of reinforcers. This knowledge resulted in higher reinforcement frequencies (reinforcers per choice) than those obtained by the strict matcher. One of the actors was "perfect omniscient": Whenever a reinforcer became available, that actor knew where to collect it and did so. A second comparison actor, "imperfect omniscient," appeared to know the presence and location of reinforcers on 70% of its choices. Both perfect (100%) and imperfect (70%) omniscient actors were therefore reinforced more frequently than was the strict matcher described earlier. The results, shown in Figure 10 (left panel), were that the strict matcher (where $s = 1.0$ and choices were stochastically generated throughout) was judged to better represent a voluntary chooser than was either of the other two and that the 70% omniscient matcher was judged to approximate voluntary choices better than did the 100% omniscient actor. Comments from participants suggest that these results emerged

because the omniscient actors' responses appeared to be controlled—or determined—in a way inconsistent with choices in a gambling-type situation in which reinforcers were uncertain. Imagine, for example, someone at a roulette wheel who always won: Whenever he chose to place a bet on a number, that number would be the next to appear. A lucky player? Not if the luck continued across hundreds of trials and therefore the choices appeared to be controlled by something other than the free (and uncertain) choices made by most players at this gambling game. So, too, it appeared that our participants judged the strict stochastic matcher as better representing voluntary choice, despite the omniscient actors' being reinforced with much higher probabilities than was the matcher. High frequency of reinforcement alone did not generate high judgments of volition.

A second control procedure tested the importance of functionality to judgments of volition. Recall that the OVVA theory posits that volition is indicated by functional changes in levels of (un)predictability. In this second control, reinforcers were hidden, and therefore it was not possible for the participants to assess functionality. All that could be seen were patterns of choices, these differing across actors and across games. The participants were told that the actors were reinforced "behind the scenes" and therefore that the participants would not be able to see when a reinforcer was provided or for which choice. The result, under these hidden-reinforcement conditions, was that the participants based judgments of volition exclusively on levels of predictability and therefore judged the undermatchers to best represent a voluntary chooser (a result related to the animacy research described later in the Animacy subsection of the Related Areas section). The implication is that when observers cannot judge whether behaviors are functional, unpredictability indicates volition. However, when observers can judge functionality, they favor actors whose stochas-

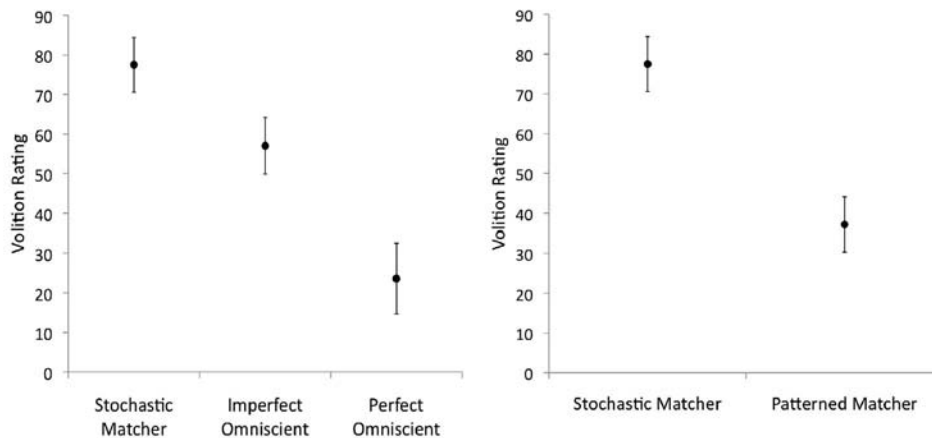


Figure 10. Judgments of human volition as a function of the algorithms used to program the actors' choices. Left panel: The stochastic matcher algorithm allocated responses according to the generalized matching function (i.e., matched response proportions to reinforcer proportions). The perfect omniscient actor responded as if it could identify the availability and location of each reinforcer, and therefore reinforcer rates were exceedingly high. The imperfect omniscient matcher combined aspects of the stochastic matcher and the perfect omniscient algorithms. Right panel: The patterned matcher matched response proportions to reinforcer proportions in a patterned and predictable manner. Note that the stochastic matcher judgments are the same in left and right panels. Error bars show standard errors. (Adapted from Neuringer et al., 2007)

ticity varies in a functional manner according to context, as was the case for the strict matcher.

A third possible explanation derived from the relationship between choice distributions and choice predictability. The OVVA theory asserts that variations in variability (or predictability) will indicate volition when the variations are seen to be functionally related to environmental conditions. In the present studies, the strict matcher ($s = 1.0$) demonstrated the widest range of predictability but also the widest range of response distribution, and these were correlated. Thus, for example, when the strict matcher distributed its choices equally across the three options (given equal reinforcers for the three), its responses were relatively unpredictable; and when—under a different set of reinforcement conditions (i.e., a different game)—it chose one alternative predominantly, predictability of those choices became relatively easy. Response distributions and response predictability were correlated. It was therefore possible that strict matching alone (predictable or otherwise; stochastic or not) could explain the results. The third control procedure therefore examined whether participants' judgments were based exclusively on matching or, as predicted by the OVVA theory, on the combination of matching and stochasticity (i.e., stochastic matching). As a test, we introduced another actor who was also a strict matcher ($s = 1.0$) but who responded in a highly predictable fashion across all reinforcer distributions. The choices of this "patterned matcher" were seen to occur in repeated blocks. If, for example, obtained reinforcers were distributed 5:2:1, then the patterned matcher would emit the sequence XXXXXYYZXXXXXXYYZ . . . (i.e., it would repeatedly choose one gamble five times, then the second gamble twice, and the third a single time, with this pattern repeated over and over until the obtained distribution of reinforcers changed). This patterned matcher was reinforced exactly as often as the stochastic matcher was, so that both matching and functionality were identical for the two. The only difference was in the predictability of responses. Because the patterned matcher followed a rule, its behavior was consistently predictable, no matter the reinforcer ratios. Because the strict matcher's choices were probabilistically generated, predictions of its choices depended on the distributions of reinforcements—easy when reinforcers were preponderantly from one option and no better than chance when reinforcers were equal across the three options. Figure 10 (right panel) shows that participants rated the stochastic matcher as better representing voluntary human choice, a significant effect, suggesting that the combination of functionality (resulting from matching) and different levels of unpredictability (as produced by stochastic emission) was important, as predicted by the OVVA.

Choice distributions and choice variability therefore combine to indicate voluntary behavior. Choice distributions alone did not lead responses to be evaluated as highly voluntary; this was shown by the patterned matcher, whose choices could readily be predicted. Thus, matching of response frequencies to reinforcer frequencies does not suffice to produce highest judgments of volition. Unpredictable responding alone (as with the $s < 1.0$ undermatchers) also did not lead to highest judgments. Choices were most voluntary in appearance when probabilities and distributions of stochastic responses changed with distributions of reinforcers. These results are consistent with the OVVA theory claim that operant variability provides evidence for volition.

Operant variability implies that response variability is functional (or intended to be functional) and that the level of variability (which translates into predictability by another organism) changes, depending upon the needs of the moment. The same is true for voluntary behaviors: They are predictable or unpredictable, depending upon circumstances. But it is important to note that even when individual responses are unpredictable, the response class—or set of responses from which individual instances emerge—can often be identified and predicted. For example, under some circumstances, the response to the "How are you?" question can readily be predicted for a given acquaintance. But even when the situation warrants unpredictable responses, as when the question is asked in the context of a "fool the questioner" game, some veridical predictions can be made: that the response will be verbal, that it will contain particular parts of speech, and so on. The functionality of variability implies a *degree of predictability* in the resulting behaviors that is related to the activated class of possibilities from which the response emerges. The class can often be predicted on the basis of knowledge of the behaving organism and current environmental conditions.

The unpredictable end of the continuum is highlighted in many discussions of volition. Indeed, as suggested previously, the size of the active set can be exceedingly large—and functionally so—for if someone were attempting to prove that she is a "free agent," the set of possibilities might consist of all responses in her repertoire (see Scriven, 1965). But we return to the fact that voluntary behaviors can be predictable as well as not. The most important characteristic is *functionality of variability*, or the ability to change levels of predictability in response to environmental demands.

An objection may be raised that the just-described psychophysical research does not constitute evidence for the existence of volition in the outside-the-laboratory real world. After all, a consensus on how Santa Claus looks does not prove Santa's existence. If aspects of a Santa Claus-like figure were systematically modified and individuals were asked to judge how well the figures approximated Santa Claus, there might be consistency in responses at least as high as those in the present study. This result would indicate that people agree as to what Santa looks like, but it would say nothing about whether a real Santa flies in a sleigh above the earth. So, too, it is possible that people agree about the appearance of voluntary behavior, even though volition itself (however it might be defined) is not real.

As a theory, however, the OVVA deliberately distances itself from definitions of volition that are disconnected from behaviors. In so doing, we follow natural-language philosophers. For example, in *The Concept of Mind*, Ryle (1949) argued forcefully that interpreting volition in terms of mental or spiritual events that precede voluntary acts is gratuitous and unhelpful. Voluntary acts can be explained by behaviors and the situations in which they occur. We, too, argue for a behavior-based theory, but unlike the offerings of philosophers, the OVVA theory relies on experimental evidence to identify controlling conditions. The psychophysical experiments presented in this article were based on two behavioral phenomena that have empirical support: a functional relationship between choice and reinforcer probabilities, and control of variability by reinforcement contingencies. The extensive behavioral research on concurrent choices, the many demonstrations of operant variability, and the psychophysical evidence regarding estimations of voluntary-like responding together describe and define a

broad class of real-world behaviors that, we submit, are reasonably and usefully referred to as voluntary.

Related Areas

Volitional Continua

The OVVA theory rejects the claim that voluntary acts emanate from an extraphysical source, as proposed by Descartes (1614/1931) and others. Rather, the capacity to behave voluntarily is a natural quality that evolved under the same types of selection pressures that led to other behavioral competencies (see Dennett, 2003). The evolutionary history of this ability goes back as far as single-celled organisms. For example, *E. coli* bacteria have two basic responses to chemicals in their environment. If the chemical gradient (across time) is positive, the bacterium moves ahead in a straight line (i.e., it behaves predictably). If the gradient is flat, the bacterium tumbles randomly (Macnab & Koshland, 1972; discussed in Staddon, 2001). More complex organisms manifest a similar ability to repeat and vary, in increasingly complex ways. In some mammals and birds, variable, unpredictable flights and movements are responses to threatening situations, such as the presence of a predator (Driver & Humphries, 1988). Other functions of variability include attracting attention, perhaps of potential mates, as shown by increased variability and complexity of songs by male birds when in the presence of females (Catchpole & Slater, 1995; Searcy, 1984). Many of these examples constitute “elicited variability,” in which an entire species reliably displays certain kinds of behavioral variability (or systematicity) when faced with particular environmental challenges. These forms of species-typical, epigenetically generated variable and predictable behaviors, although adaptive under certain circumstances, often cannot adapt to novel situations.

At some point in evolutionary history, organisms evolved the capacity to vary levels of variability as a rapid, adaptive response to changing environmental circumstances (Potts, 1998). That is, behavioral variability became an operant. The chief difference between elicited and operant variability is immediate sensitivity, in the latter case, to changing demands and feedback from the environment. This includes the rapid generation and modification of operant response classes as well as moment-to-moment changes in predictability. At present, we don't know when elicited variability evolved to become operant variability, but research has shown considerable cross-species generality in this operant competency, including among fish, birds, rodents, and humans (Neuringer, 2002).

Both behavioral and psychophysical lines of research described in the present article are consistent with a *continuum of volition*, a position that runs contrary to most views of free will. Judgments of volition were not confined to the exact matcher but increased as a function of approximations to matching. This indication of a voluntary–involuntary continuum parallels, and indeed may depend upon, the operant variability continuum that enables moving from repetitions to randomlike responding. Conceptualizing volition in terms of a levels-of-variability continuum provides access to empirical tests of differences in volitional competency. Species differences, as described previously; age differences; psychological health and pathology; central nervous system injuries; and drugs may all contribute to people's ability to move successfully

along the variability continuum and consequently their ability to behave voluntarily (see Brugger, 1997; Neuringer, 2002; Nickerson, 2002). Explorations of such differences may contribute importantly to understanding volition as well as helping to ameliorate volitional deficiencies.

Animacy

In research on animacy, human participants judge the extent to which moving objects on a computer screen appear to be alive.¹ In some studies, simple objects (e.g., geometric forms or letters of the alphabet) move around the screen with types and amounts of interactions varied. Participants are asked to describe what the forms are doing or to rate appearance along, for example, a scale of 1 (*definitely not alive*) to 7 (*definitely alive*). The results of these studies are consistent with and complement the previously described psychophysical studies of volition. Animacy judgments are found to depend upon the type of movement (e.g., whether readily predictable or not) and the apparent purposefulness or functionality of the movement (e.g., whether one icon [e.g., a “predator”] appears to move in the direction of another [a “prey”]). Unexpected or unpredictable changes in speed or angle of movement of an icon increase judgments of animacy (Tremoulet & Feldman, 2000). Purposeful interactions between the moving icon and its environment, or what we have referred to as functionality, also increase the likelihood of attributing animacy (Opfer, 2002; Tremoulet & Feldman, 2006). Thus, attributes of animacy appear to be related to those of volition, a relationship well worth exploring in future research.

Consciousness

It is commonly claimed that conscious deliberations cause (or are required for) voluntary actions—a form of the “mental events precede and are responsible for voluntary behaviors” position previously discussed. But recent research questions whether consciousness is necessary for (or indeed involved in) the generation of at least some voluntary acts (e.g., Bargh, 2005; Libet, 1985; Wegner, 2002, 2005). Mounting evidence shows that many acts that seem voluntary (both to actor and to observers) actually occur independently of awareness of the causes or initiators of the act. Additionally, many voluntary-seeming behaviors can be predicted from environmental contexts about which the actor is unaware, suggesting that after-the-fact reasons given for these actions may be confabulated post hoc (Nisbett & Wilson, 1977).

These important findings do not rule out the possibility of conscious (or attentional) influences on the generation of voluntary behaviors. As previously described, when participants are asked to generate a random sequence without feedback, their responses, though quite variable, do not pass statistical tests of randomness. If the participants simultaneously engage in a task that competes for attention, responding becomes even more predictable than otherwise (Baddeley, 1966; Evans & Graham, 1980). These studies show that withdrawal of, or competition for, attention interferes with the ability to behave variably. Although attentional studies that involve *reinforced*

¹ We thank an anonymous reviewer for highlighting this area of research.

variability have not been published, observations in our laboratory also indicate a decrease in ability to approximate a random model when there is competition for attention (e.g., Deiss, 1994). One interpretation of these attentional effects is given by Baddeley, Emslie, Kolodny, and Duncan (1998). They propose that generation of randomlike responding requires “executive control,” a suggestion related to some libertarian theories of volition (see Doyle, 2009). Another interpretation is that attention may be involved in the formation of the sets or classes of responses from which variations emerge. For example, Bargh (2005) suggested that consciousness is a “workplace” where components of behavioral acts may be formulated and combined.

In terms of the OVVA theory, attention may contribute to the ongoing construction of generic response classes, or those sets of *potential responses* from which variable instances emerge. Evidence related to this claim is seen in the role of attention in the formation of categories and concepts (Barsalou, 1987). An issue of considerable importance is how response sets are formed and modified, and how within-set probability distributions change as a function of environmental demands (see, e.g., Neuringer et al., 2001). Volition involves stochastic emergence from within continually changing sets of responses, and attention may play an important role in set creation.

Perception

To the extent that behaving in a functionally varied manner contributes to the emission of voluntary behaviors, assessing variability in the behaviors of others may be important for judging the nature of an act—voluntary or not—and predicting (or dealing with) future behaviors. Wasserman, Young, and Cook (2004) reviewed the evidence showing that people, monkeys, and pigeons have an impressive ability to discriminate different degrees of variability in arrays of visual stimuli. Perceptions and judgments of volition—related to perceptions of variability in the ways previously described—may influence how one organism reacts to the actions of another. In human societies, volitional judgments are important in assessing psychopathological states and in deciding upon consequences—punishing or not—for example, when laws of society are broken.

Indeterminacy and Volition

Many philosophers, including some libertarian philosophers, have invoked indeterminate physical influences to explain freely willed decisions and choices. For example, Karl Popper once said:

New ideas have a striking similarity to genetic mutations. Mutations are, it seems, brought about by quantum theoretical indeterminacy. . . . [O]n them there subsequently operates natural selection which eliminates inappropriate mutations. Now we could conceive of a similar process with respect to new ideas and to free-will decisions. . . . That is to say, a range of possibilities is brought about by a probabilistic and quantum mechanically characterized set of proposals, as it were—of possibilities brought forward by the brain. On these there then operates a kind of selective procedure which eliminates those proposals and those possibilities which are not acceptable to the mind. (Popper & Eccles, 1977, p. 540)

A related explanation was offered by Dennett (1978), who, although generally opposed to the use of indeterminism to explain

volition (Dennett, 2003), described a two-stage model of the following kind. When confronted with a situation in which a decision must be made, an individual entertains a set of alternative possibilities that are generated partly by undetermined, random processes. There follows a consideration period resulting in selection of a decision. The selection process is based on causal events, including the individual’s make-up, prior experiences, the current situation, and rational considerations, and therefore is predictable.

Philosophers are not the only ones to invoke indeterminism to explain volition. A number of physicists have offered similar hypotheses. For example, Margenau wrote:

. . . quantum mechanics leaves our body, our brain, at any moment in a state with numerous (because of its complexity we might say *innumerable*) possible futures, each with a predetermined probability. Freedom involves two components: *chance* (existence of a genuine set of alternatives) and *choice*. Quantum mechanics provides the chance, and we shall argue that only the mind can make the choice by selecting . . . among the possible future courses. (Leshan & Margenau, 1982, p. 240)

Other physicists, including Arthur Eddington and Arthur Compton, have invoked physical indeterminism to explain human volition. In each of these cases, indeterministic influences generate possibilities followed by determined (or causally accountable) selections.

An inverted form of this type of two-stage hypothesis was developed by Kane, a philosopher who has written extensively about free will. In cases in which an important decision must be made, and in which long-term considerations have resulted in a number of alternatives remaining in contention, Kane (1985) hypothesized that there is a sensitivity of neurons in the brain to quantum-indeterminate or chaotic influences that leads ultimately to a choice being made. In these cases, the final stage in the decision process—leading to action—is indeterministic selection. Thus, for Kane, possibilities are determined but selections are, at least in some cases, randomly produced. Helpful discussions of these and other related theories can be found in Doyle (2009).

Such theories differ from the OVVA in two important ways. First and foremost, the OVVA is based on empirical findings rather than theoretical conjectures. Second, the just-described theories *separate* indeterminist and determinist processes, as in the variation-and-selection evolutionary model. The OVVA theory describes behavioral output in terms of a *combination* of stochasticity and functionality. For an organism capable of operantly varying its behaviors, the parameters of its stochastic generator process are influenced by its current state; its experiences, including conditioning history; and the environmental context. Also, the instances upon which stochastic emergence is based—the members of a set of possibilities—are selected or determined by experiences and the like. Similarly, the size of the set, the probability distribution, and so forth are influenced in a determined manner (e.g., selected by consequences). Thus, the stochastic process does not proceed independently of deterministic influences. When behavior is voluntary, the process of variability-generation is not blind. It is equally the case that functional, or selective, aspects of voluntary behavior are not independent of stochastic influences. When an organism chooses from among options, it may match responses to reinforcers, but it does so stochastically. More generally, learning to respond adaptively to environmental influences involves stochastic changes. To summarize, according to

the OVVA theory, the stochasticity of voluntary behaviors involves functionally selective (or deterministic) aspects, and the functionality of those behaviors involves stochastic influences. Stated differently, volition merges functionality and stochasticity.

Conclusion

The sun's apparent movement across the sky is a real phenomenon. We submit that the same is true for volition: It refers to a real behavioral phenomenon. But as with the sun, valid explanations of the observations have been hard to come by. We asked, "What does voluntary behavior look like to an outside observer?" To answer, we described experiments showing that behavioral variability is influenced by reinforcement contingencies. We then provided evidence that judgments of volition are related to such variability. Levels of variability that adapted most readily to environmental conditions and that spanned a range from repetitions to randomlike responding were most likely to be judged as voluntary. Voluntary behavior is characterized by functionally changing levels of variability. Relating volition to an empirically observable characteristic, such as functional variations in variability, enables testing for effects on volitional competency of developmental disabilities, brain injury, drugs, and other variables. Furthermore, empirical research can resolve the apparent inconsistency—noted by many writers on volition and free will—between "free" and "functional." The explanation is seen in the evolved capacity of animals and people to behave functionally and, when reinforced for so doing, unpredictably.

References

- Arnesen, E. M. (2000). *Reinforcement of object manipulation increases discovery*. Unpublished undergraduate thesis, Reed College.
- Baddeley, A. D. (1966). The capacity for generating information by randomization. *Quarterly Journal of Experimental Psychology*, *18*, 119–129.
- Baddeley, A. D., Emslie, H., Kolodny, J., & Duncan, J. (1998). Random generation and the executive control of working memory. *Quarterly Journal of Experimental Psychology*, *51(A)*, 819–852.
- Bailey, D. H., & Crandall, R. E. (2001). On the random character of fundamental constant expansions. *Experimental Mathematics*, *10*, 175–190.
- Balsam, P. D., Deich, J. D., Ohshima, T., & Stokes, P. D. (1998). Origins of new behavior. In W. O'Donohue (Ed.), *Learning and behavior therapy* (pp. 403–420). Boston, MA: Allyn & Bacon.
- Bargh, J. A. (2005). Bypassing the will: Toward demystifying the nonconscious control of social behavior. In R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious* (pp. 19–36). New York, NY: Oxford University Press.
- Barsalou, L. W. (1987). The instability of graded structures in concepts. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 101–140). New York, NY: Cambridge University Press.
- 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.
- Blough, D. S. (1966). The reinforcement of least-frequent interresponse times. *Journal of the Experimental Analysis of Behavior*, *9*, 581–591.
- Brugger, P. (1997). Variables that influence the generation of random sequences: An update. *Perceptual and Motor Skills*, *84*, 627–661.
- Budescu, D. V., & Rapoport, A. (1994). Subjective randomization in one- and two-person games. *Journal of Behavioral Decision Making*, *7*, 273–302.
- Catchpole, C. K., & Slater, P. J. (1995). *Bird song: Biological themes and variations*. Cambridge, England: Cambridge University Press.
- Cohen, L., Neuringer, A., & Rhodes, D. (1990). Effects of ethanol on reinforced variations and repetitions by rats under a multiple schedule. *Journal of the Experimental Analysis of Behavior*, *54*, 1–12.
- Davison, M., & McCarthy, D. (1988). *The matching law: A research review*. Hillsdale, NJ: Erlbaum.
- Deiss, C. (1994). *On the nature of operant variability: Evaluating dual-task interference between schedules of variability and repetition*. Unpublished undergraduate thesis, Reed College, Portland, Oregon.
- Dennett, D. (1978). *Brainstorms*. Cambridge, MA: MIT Press.
- Dennett, D. (2003). *Freedom evolves*. New York, NY: Viking Adult.
- Denney, J., & Neuringer, A. (1998). Behavioral variability is controlled by discriminative stimuli. *Animal Learning & Behavior*, *26*, 154–162.
- Descartes, R. (1931). *The philosophical works of Descartes* (E. S. Haldane & G. R. T. Ross, trans.). New York, NY: Cambridge University Press. (Original work published 1614)
- Doughty, A. H., & Lattal, K. A. (2001). Resistance to change of operant variation and repetition. *Journal of the Experimental Analysis of Behavior*, *76*, 195–215.
- Doyle, R. (2009). The cogito model. Retrieved from the Information Philosopher website: <http://www.informationphilosopher.com/freedom/cogito/>
- Driver, P. M., & Humphries, D. A. (1988). *Protean behavior: The biology of unpredictability*. Oxford, England: Oxford University Press.
- Evans, F. J., & Graham, C. (1980). Subjective random number generation and attention deployment during acquisition and overlearning of a motor skill. *Bulletin of the Psychonomic Society*, *15*, 391–394.
- Gharib, A., Gade, C., & Roberts, S. (2004). Control of variation by reward probability. *Journal of Experimental Psychology: Animal Behavior Processes*, *30*, 271–282.
- Gigerenzer, G., Swijtink, Z., Porter, T., Daston, L., Beatty, J., & Kruger, L. (1989). *The empire of chance: How probability changed science and everyday life*. Cambridge, England: Cambridge University Press.
- Glimcher, P. W. (2003). *Decisions, uncertainty, and the brain*. Cambridge, MA: MIT Press.
- Glimcher, P. W. (2005). Indeterminacy in brain and behavior. *Annual Review of Psychology*, *56*, 25–56.
- Goetz, E. M. (1989). The teaching of creativity to preschool children. In J. A. Glover, R. R. Ronning, & C. R. Reynolds (Eds.), *Handbook of creativity* (pp. 411–428). New York, NY: Plenum Press.
- Goetz, E. M., & Baer, D. M. (1973). Social control of form diversity and emergence of new forms in children's blockbuilding. *Journal of Applied Behavior Analysis*, *6*, 209–217.
- Grunow, A., & Neuringer, A. (2002). Learning to vary and varying to learn. *Psychonomic Bulletin & Review*, *9*, 250–258.
- Jensen, G., Miller, C., & Neuringer, A. (2006). Truly random operant responding: Results and reasons. In E. A. Wasserman & T. R. Zentall (Eds.), *Comparative cognition: Experimental explorations of animal intelligence* (pp. 459–480). Oxford, England: Oxford University Press.
- Jensen, G., & 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.
- Kane, R. (1985). *Free will and values*. New York, NY: SUNY Press.
- Kane, R. (Ed.). (2002). *The Oxford handbook of free will*. Oxford, England: Oxford University Press.
- Lee, R., McComas, J. J., & Jawor, J. (2002). The effects of differential and lag reinforcement schedules on varied verbal responding by individuals with autism. *Journal of Applied Behavior Analysis*, *35*, 391–402.
- Leshan, L., & Margenau, H. (1982). *Einstein's space and van Gogh's sky: Physical reality and beyond*. New York, NY: Macmillan.

- Libet, B. (1985). Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and Brain Science*, 8, 529–566.
- Machado, A. (1989). Operant conditioning of behavioral variability using a percentile reinforcement schedule. *Journal of the Experimental Analysis of Behavior*, 52, 155–166.
- Machado, A. (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.
- Macnab, R. M., & Koshland, D. E., Jr. (1972). The gradient-sensing mechanism in bacterial chemotaxis. *Proceedings of the National Academy of Sciences, USA*, 69, 2509–2512.
- Maes, J. H. R. (2003). Response stability and variability induced in humans by different feedback contingencies. *Learning & Behavior*, 31, 332–348.
- Maynard Smith, J. (1982). *Evolution and the theory of games*. Cambridge, England: Cambridge University Press.
- Miller, N., & Neuringer, A. (2000). Reinforcing variability in adolescents with autism. *Journal of Applied Behavior Analysis*, 33, 151–165.
- Mook, D. M., & Neuringer, A. (1994). Differential effects of amphetamine on reinforced variations versus repetitions in spontaneously hypersensitive rats. *Physiology & Behavior*, 56, 939–944.
- Moreh, J. (1994). Randomness, game theory, and free will. *Erkenntnis*, 41, 49–64.
- Moxley, R. A. (1997). Skinner: From determinism to random variation. *Behavior and Philosophy*, 25, 3–28.
- Nash, J. (1951). Non-cooperative games. *Annals of Mathematics*, 54, 286–295.
- Neuringer, A. (1986). Can people behave “randomly”? The role of feedback. *Journal of Experimental Psychology: General*, 115, 62–75.
- Neuringer, A. (1991). Operant variability and repetition as a function of interresponse time. *Journal of Experimental Psychology: Animal Behavior Processes*, 17, 3–12.
- Neuringer, A. (1992). Choosing to vary or repeat. *Psychological Science*, 3, 246–250.
- Neuringer, A. (2002). Operant variability: Evidence, functions, and theory. *Psychonomic Bulletin & Review*, 9, 672–705.
- Neuringer, A., Jensen, G., & Piff, P. (2007). Stochastic matching and the voluntary nature of choice. *Journal of the Experimental Analysis of Behavior*, 88, 1–28.
- Neuringer, A., Kornell, N., & Olufs, M. (2001). Stability and variability in extinction. *Journal of Experimental Psychology: Animal Behavior Processes*, 27, 79–94.
- Neuringer, A., & Voss, C. (1993). Approximating chaotic behavior. *Psychological Science*, 4, 113–119.
- Nevin, J. A. (1969). Interval reinforcement of choice behavior in discrete trials. *Journal of the Experimental Analysis of Behavior*, 12, 875–885.
- Nickerson, R. A. (2002). The production and perception of randomness. *Psychological Review*, 109, 330–357.
- Nisbett, R., & Wilson, T. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 84, 231–259.
- Opfer, J. E. (2002). Identifying living and sentient kinds from dynamic information: The case of goal-directed versus aimless autonomous movement in conceptual change. *Cognition*, 86, 97–122.
- Page, S., & Neuringer, A. (1985). Variability is an operant. *Journal of Experimental Psychology: Animal Behavior Processes*, 11, 429–452.
- Popper, K. R., & Eccles, J. C. (1977). *The self and its brain*. New York, NY: Springer-Verlag.
- Potts, R. (1998). Variability selection in hominid evolution. *Evolutionary Anthropology*, 7, 81–96.
- Pryor, K. W., Haag, R., & O'Reilly, J. (1969). The creative porpoise: Training for novel behavior. *Journal of the Experimental Analysis of Behavior*, 12, 653–661.
- Roberts, S., & Neuringer, A. (1998). Self-experimentation. In K. A. Lattal & M. Perone (Eds.), *Handbook of research methods in human operant behavior* (pp. 619–655). New York, NY: Plenum Press.
- Robertson, R., & Combs, A. (Eds.). (1995). *Chaos theory in psychology and the life sciences*. Mahwah, NJ: Erlbaum.
- Rosch, E. H. (1978). Principles of categorization. In E. H. Rosch & B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Hillsdale, NJ: Erlbaum.
- Ross, C., & Neuringer, A. (2002). Reinforcement of variations and repetitions along three independent response dimensions. *Behavioural Processes*, 57, 199–209.
- Ryle, G. (1949). *The concept of mind*. London, England: Hutchinson.
- Scriven, M. (1965). An essential unpredictability in human behavior. In B. Wolman (Ed.), *Scientific psychology* (pp. 411–425). New York, NY: Basic Books.
- Searcy, W. A. (1984). Song repertoire size and female preference in song sparrows. *Behavioral Ecology & Sociobiology*, 14, 281–286.
- Searle, J. (2007). *Freedom and neurobiology: Reflections on free will, language, and political power*. New York, NY: Columbia University Press.
- Silberberg, A., Hamilton, B., Ziriax, J. M., & Casey, J. (1978). The structure of choice. *Journal of Experimental Psychology: Animal Behavior Processes*, 4, 368–398.
- Skinner, B. F. (1938). *The behavior of organisms: An experimental analysis*. New York, NY: Appleton-Century.
- Skinner, B. F. (1961). The generic nature of the concepts of stimulus and response. In *Cumulative record* (Enlarged ed., pp. 347–366). New York, NY: Appleton-Century-Crofts. (Original work published 1935)
- Skinner, B. F. (1974). *About behaviorism*. New York, NY: Knopf.
- Staddon, J. E. R. (2001). *Adaptive dynamics: The theoretical analysis of behavior*. Cambridge, MA: MIT Press.
- Stokes, P. D., Lai, B., Holtz, D., Rigsbee, E., & Cherrick, D. (2008). Effects of practice on variability, effects of variability on transfer. *Journal of Experimental Psychology: Human Perception and Performance*, 34, 640–659.
- Tremoulet, P. D., & Feldman, J. (2000). Perception of animacy from the motion of a single object. *Perception*, 29, 943–951.
- Tremoulet, P. D., & Feldman, J. (2006). The influence of spatial context and the role of intentionality in the interpretation of animacy from motion. *Perception & Psychophysics*, 68, 1047–1058.
- Van Inwagen, P. (2000). Free will remains a mystery. *Philosophical Perspectives*, 14, 1–19.
- Voorspoels, W., Vanpaemel, W. M., & Storms, G. (2008). Exemplars and prototypes in natural language concepts: A typicality-based evaluation. *Psychonomic Bulletin & Review*, 15, 630–637.
- Ward, L. M., & West, R. L. (1994). On chaotic behavior. *Psychological Science*, 5, 232–236.
- Wasserman, E. A., Young, M. E., & Cook, R. G. (2004). Variability discrimination in humans and animals: Implications for adaptive action. *American Psychologist*, 59, 879–890.
- Wegner, D. M. (2002). *The illusion of conscious will*. New York, NY: Worth.
- Wegner, D. M. (2005). Who is the controller of the controlled processes? In R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious* (pp. 19–36). New York, NY: Oxford University Press.
- Weiss, R. L. (1965). “Variables that influence random-generation”: An alternative hypothesis. *Perceptual and Motor Skills*, 20, 307–310.
- Williams, B. A. (1988). Reinforcement, choice, and response strength. In C. R. Atkinson, R. J. Herrnstein, G. Lindzey, & R. D. Luce (Eds.), *Stevens' handbook of experimental psychology* (Vol. 2, pp. 167–244). New York, NY: Wiley.
- Zipf, G. K. (1935). *The psycho-biology of language: An introduction to dynamic philology*. New York, NY: Houghton-Mifflin.

Received July 31, 2009

Revision received January 21, 2010

Accepted January 22, 2010 ■