

This is the peer reviewed version of the following article: Catania, A. Charles; Tactical contingencies in the experimental analysis of reinforcement and operant classes; Journal of the Experimental Analysis of Behavior (2020); <https://onlinelibrary.wiley.com/doi/abs/10.1002/jeab.648> ,which has been published in final form at <https://doi.org/10.1002/jeab.648>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Use of Self-Archived Versions. Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Tactical Contingencies in the Experimental Analysis of Reinforcement and Operant Classes

A. Charles Catania

University of Maryland, Baltimore County (UMBC)

Abstract

In his "Tactics of Scientific Research" (1960), his work on avoidance, his discovery of equivalence classes and his cautions on applications of coercion, Murray Sidman created high standards for behavior analytic research. I illustrate his influence in the context of three examples he might have characterized as pilot studies. Each examined trial N+1 response probabilities depending on whether trial N responding had produced a reinforcer. Differentially reinforced interresponse times, keys pecked in arbitrary matching, and two-key response sequences provided no robust evidence that reinforcing some response property on trial N raises the probability of responding with that property on trial N+1. These negative findings shed light on the nature of operant classes and on the relation of reinforcers to the responses that produced them. Through selection, reinforcers create operant classes and engender variations of the responses within those classes; operant classes are held together by common contingencies. Sidman extended our understanding of operant classes by expanding them to include equivalence relations.

161 words

Key words: Murray Sidman, reinforcement, selection, contingencies, operant classes, tactics

Address correspondence to:

A. Charles Catania, 10545 Rivulet Row, Columbia, MD 21044 USA

Email: catania@umbc.edu

Author Footnote

Address correspondence to the author via email: catania@umbc.edu. Study 1 was supported by a 1961-1962 NSF Grant to B. F. Skinner (Catania, 2002); Study 2 was supported by the UMBC Department of Psychology and used equipment originally acquired through grants. In Studies 2 and 3, which could not have been completed without the able assistance of several UMBC undergraduates, Eliot Shimoff was not only a collaborator but also a prime mover in developing the computer programs that ran them. He once told me that Sidman's *Tactics*, which he read before he began working with Nat Schoenfeld, was a crucial component of his move from psychophysics to behavior analysis. All research reported here conformed to the ethical guidelines for animal care that were in place at the time it was conducted.

Before Murray Sidman's "Tactics of Scientific Research" (Sidman, 1960), his research was already highly visible, especially for the procedures that came to be known as Sidman avoidance (Sidman, 1953). In the 1954-1955 academic year, when I was a sophomore, I took Fred Keller's introductory course in the undergraduate psychology curriculum at Columbia (Keller & Schoenfeld, 1949). Sidman had received his doctoral degree with Nat Schoenfeld at Columbia in 1952, so his findings were too recent to have been cited in our textbook (Keller & Schoenfeld, 1950), but they were often referenced in the courses within that curriculum. In his classes, Nat Schoenfeld used Sidman's avoidance data to illustrate the significance of parametric studies and to emphasize such research desiderata as using the widest feasible range of parameter values to maximize the yield from an experiment. Sidman had quickly become established as one of the younger founders of our discipline, along with Charlie Ferster, Joe Brady, Peter Dews, Og Lindsley and a handful of others.

In retrospect, the time between Sidman's Ph.D. and the beginnings of my commitment to behavior analysis seems startlingly brief. By the time I became a graduate student I was a charter subscriber to the *Journal of the Experimental Analysis of Behavior*. Its first volume included a tour de force case study on the analysis of conditioned suppression (Sidman, 1958). The paper demonstrated what could be done with systematic replication and guaranteed a ready reception for the 1960 book. My recollection is that the book became available after I had completed my doctoral requirements and had taken over the Harvard pigeon lab from Lew Gollub (Catania, 2002). My colleagues and I eagerly awaited our copies. Most of us began reading it from cover to cover as soon as we could get our hands on it.

In those days, the publication of behavior analytic papers in standard psychology journals was limited by expectations of hypothesis testing and by the experimental designs and statistical analyses regarded as acceptable. "The Behavior of Organisms" (Skinner, 1938) and "Schedules of Reinforcement" (Ferster & Skinner, 1957) had implicitly opened the way to a less constrained experimental style and Skinner had provided some informal guidelines (Skinner, 1956), but Sidman's book made the criteria for our research strategies explicit.

I take this opportunity to celebrate Sidman's extensive contributions to our discipline by recounting three interrelated experimental studies with pigeons inspired in part by the example Sidman had set in his book and, in retrospect, perhaps also relevant to his later work on equivalence classes (Sidman, 1994). They are probably best called exploratory or pilot studies. It is impossible at this temporal remove to provide all of the usual procedural details, such as feeding protocols, dimensions of experimental chambers and the sequencing and sessions of conditions, but for the present purposes it should be sufficient to describe the experiments with a broad brush in the context of the varying research standards of those times. It's unlikely that anyone will wish to replicate them in detail.

Sidman's book gave pilot studies a chapter but they didn't merit a prominent place in his taxonomy of research tactics. Of studies designated as pilot studies in advance, he wrote "They are supposed to provide the investigator with an estimate of the probable success or failure of a subsequent well-designed experiment. But the defining feature of this type of pilot study is its lack of control over certain variables"; with typical candor, he later continued, "A sloppy experiment is a poor experiment, and can never be justified by labeling it 'pilot'" (p. 218).

He went on, however, to discuss exploratory procedures arranged to assess the adequacy and appropriateness of a research design. These are designated as pilot only retrospectively,

after they have been refined in the course of addressing the technical inadequacies they revealed. Sidman provided examples in his own work and in the work of others, concluding that a pilot study "is not a necessary first step in experimental design. It is an unplanned consequence of experimentation that is carried out without sufficient knowledge. . . . But it is also because of our relative state of ignorance that pilot studies may often be turned to good use" (p. 333).

Here I describe three exploratory studies that I would not have undertaken had I recognized my relative state of ignorance with regard to reinforcement schedules. Each was concerned with whether the effects of a reinforcer could be seen in the responding that immediately followed the response that was reinforced, each was tacked on to research that had been initiated for other reasons, and each made do with limited resources under time constraints. Study 1 examined just a single dimension of reinforced responding: interresponse-time or IRT distributions after reinforcement of relatively short and relatively long IRTs during variable-interval reinforcement. The next two studies each examined two different dimensions of the responding that was reinforced. Study 2 involved discrete response classes; it used early acquisition sessions of arbitrary matching with four sample stimuli on four comparison keys to see whether the probabilities of stimulus or location pecked on trial N+1 depended on whether or not pecking was reinforced on trial N. Study 3 involved not discrete classes but rather two response continua, one of time and one of number; two-key contingencies were arranged in which reinforced key-B pecks were based on number of pecks on key-A during stimulus 1 and on time since the first key-A peck during stimulus 2. The data analysis examined whether the distributions of number and time on trial N+1 were affected by whether or not key-B pecking was reinforced in trial N and by the dimension upon which the reinforcer depended during each stimulus. These studies paralleled my transition from thinking about responses as fundamental

building blocks of operant classes to looking instead at how classes are created when different response properties (forces, durations, topographies and so on) are concurrently but independently differentiated. In carrying me through that transition, I hope these pilot studies were "turned to good use." We reinforce not responses but rather their properties, and a fundamental issue is how the effects of differential reinforcement become focused on some properties and not others. For example, it is usually easier to shape response forces or topographies than to shape IRTs, but we still have much to learn about why that is so (e.g., Catania et al., 2015; Killeen & Sitomer, 2003).

Study 1: The Differential Reinforcement of Interresponse Times (IRTs)

By the end of the Spring 1962 semester, Skinner was finishing out the NSF Grant that maintained his pigeon lab at Harvard (Catania, 2002). Those of us who had been running our research in that lab were wrapping up ongoing projects and trying to get the most out of the days remaining to us before we moved on. George Reynolds and I had completed the work we eventually reported in a JEAB monograph (Catania & Reynolds, 1968). We were especially focused on variable-interval (VI) schedules, which over a wide range of parameter values routinely produced lower response rates than variable-ratio (VR) schedules. Doug Anger (1956) provided one way to account for this rate difference with his insight that the way these schedules work made longer IRTs more likely than shorter ones to be followed by a reinforcer in VI than in VR.

But there was a paradox. On the one hand, response rate changes occurred quickly with changes in type of schedule during VI and VR responding; on the other, differential-reinforcement of low-rate or DRL schedules, which directly reinforced long IRTs, were

relatively ineffective in reducing response rates. How could the former happen if the latter was the case? Maybe it would be helpful to see whether the effects of differential reinforcement on IRTs could be observed within ongoing VI performance? If IRT reinforcement is effective, shouldn't IRTs be shorter briefly after a short IRT has been reinforced and longer briefly after a long IRT has been reinforced? The chamber and some pigeons that George and I had used were now free, so I decided to give it a try.

Those were the days of electromechanical equipment (Catania, 1972; Keister, Ritchie, & Washburn, 1951), so data collection required electronic timers, the pulses from which advanced stepping switches that distributed pecks to banks of counters corresponding to the cells of IRT distributions. Each peck reset the appropriate stepping switch to its initial position. A separate circuit was arranged to count responses during reset separately, because reset, though rapid, took long enough that occasional responses couldn't be routed to appropriate counters or other control circuits.

Four distributions were recorded during fixed time windows (usually 30 s or 60 s) both before and after a reinforced short IRT and both before and after a reinforced long IRT; no responses were reinforced during those times. But even with 30-s windows the minimum inter-reinforcer time would be 60 s, so half the scheduled reinforcers became available without an IRT contingency and the other half were equally divided between short-IRT (e.g., <.5 s) and long-IRT (e.g., 2-3 s) contingencies. Without this arrangement, the VI schedule would have included no short inter-reinforcer intervals and would therefore have been more like a fixed-interval (FI) than a VI schedule. As it was, the distribution of inter-reinforcer intervals was skewed toward longer intervals.

Another complication was that times between the end of an arranged before period and the delivery of a reinforcer could vary depending on how long it took after the reinforcer setup for the pigeon to produce an eligible IRT; responding during those times was not included in the IRT distributions. An informal criterion was to exclude sessions in which the mean of those times was greater than 10 s. At the end of each session, data were written down by hand from the banks of counters, and various checks of totals versus summated counts within distributions were used to assess the accuracy of the system. In lieu of starting out with naive pigeons, some of those from the VI monograph study were used to test the equipment. A sample data sheet is shown in Figure 1.

<<Insert Figure 1 about here.>>

Two pigeons with minimal experimental histories were added later, on the assumption that effects of reinforcers might be more prominent early in VI acquisition. One had served in a classroom demonstration for Skinner's undergraduate course (NatSci 114); sessions for the other began immediately after key pecking had been shaped. Table 1 presents data from an initial session from the latter, Bird 201. After reinforcement of IRTs in the 0 to 0.5-s class or in the 2 to 3-s class, the relative frequencies of IRTs in those respective classes were slightly higher than before reinforcement (bold underlined entries in the table). Other instances also showed up in data from other birds. But increases were about as likely as decreases. I'd like to report that I recognized these differences as small enough to qualify as negative results. My notes from the time, however, show that I regarded them as real though small effects. I assumed the phenomenon was there for capture and saw the problem as an experimental one of amplifying small effects by refining details such as the duration of the windows for sampling the IRT distributions rather than as one involving my assumptions about how behavior worked.

<<Insert Table 1 about here.>>

There was also a problem of data analysis. The appropriate measure of the likelihood of an IRT is not its relative frequency but rather its probability given that there has been an opportunity for it to occur. The appropriate statistic, IRTs/op (Catania & Reynolds, 1968), is calculated by dividing IRTs within a given class by the sum of IRTs in that and longer classes. But my time in the lab was almost over, so I packed away the data and didn't resurrect them for a more thorough treatment until long after. My later quantitative analysis for sample data from two pigeons is illustrated in Table 2. These two sessions (top frame for each bird) appear to show increases, though only in the longer IRT class (1 to 2 s for Bird 278, 2 to 3 s for Bird 279). For both birds and for short as well as long IRT classes, however, with relative frequencies recalculated as IRTs/op (bottom frame for each bird), reinforcement within a given IRT class was followed in each instance by a lower probability of responding within that class.

<<Insert Table 2 about here.>>

Within the first IRT class, IRT frequency and IRTs/op should be identical. The data in Table 2 differ slightly, however, because of rounding conventions that changed over time. For example, to avoid the systematic biases of always rounding down or always rounding up one convention was always to round to the nearest even number. Our data are typically too variable for such niceties to matter. But I don't recall which conventions were in place when I did the various calculations, so I've let the differences in Table 2 stand. Most computer programs these days simply lop off the extra digits when rounding. A more general point, however, follows from these observations. In the early days of computer programming one concern was that such built-in routines took the control of such details out of our hands and, as I have argued elsewhere (Catania, 2011), we must remain vigilant that differences in subtle details of computer programs

that arrange schedule contingencies can have substantial effects on behavior, and such details are sometimes deeply embedded within programs.

I eventually tabulated data from twenty sessions distributed among six pigeons. As measured by IRTs/op in pigeons with both minimal and extensive experimental histories and equally for both short-IRT and long-IRT contingencies, in about 70% of the available comparisons the short-term effect of reinforcing a response within an IRT class was not to raise the probability of an IRT within the class but rather to reduce it. The effects in either direction were typically small, with a mean of less than .03. I came to regard the outcome as a negative finding, perhaps attributable to the atypical temporal arrangements within the baseline VI schedule, but the tendency toward a result in a direction opposite from what I had been looking for nagged at me. I didn't get back to the question until the 1990s.

Study 2: Arbitrary Matching

An opportunity to revisit the question of whether the effect of reinforcing response N could be seen immediately in response N+1 arose when Eliot Shimoff and I set up some arbitrary matching procedures to study equivalence relations. Our plan was to start with a test of transitivity. We wanted to try our hand at creating equivalence classes (Sidman, Cresson, & Willson-Morris, 1974; Sidman, Wynne, Maguire, & Barnes, 1989), we had taken to heart Sidman's case that two stimuli are not enough in matching procedures (Sidman, 1987), and we had also been following the methodological advances that eventually allowed Urcuioli (2008) and Wright (2012) to break out of the constraints of standard matching procedures to demonstrate new varieties of stimulus control with pigeons and other nonhuman organisms.

To start our training of arbitrary matching, we set up a pigeon chamber with five keys, configured as in Figure 2. The top middle key was our sample key, with four comparison stimuli presented in changing positions on the other four keys. Except that there were four comparison keys instead of two, ours was a standard matching procedure: the sample key was lit and a peck on the sample lit the four comparison keys; a peck on the arbitrary matching comparison key produced food and a peck on any other comparison key instead produced a brief time-out.

<<Insert Figure 2 about here.>>

We created three stimulus categories of four stimuli each (colors, forms and textures), as illustrated in Figure 3. Thus, instead of just two equivalence classes we could establish four of them, each consisting of one color, one form and one texture (i.e., the columns of Figure 3). The data presented here were drawn from the first and fifth sessions with *ab* matching and the first session with *bc* matching, where the assignment of colors, forms or textures to *a*, *b* or *c* categories was counterbalanced across birds as shown in Table 3. Each pigeon was experimentally naive, and each was provided with a prior history after the shaping of key pecks that included reinforcing pecking (1) on the sample key lit alone with each potential sample stimulus, (2) on each comparison key lit alone with each potential comparison stimulus, and then (3) on *ab* and *bc* trials in which samples were presented but only the correct comparison stimulus was lit. Full arbitrary-matching trials with the appropriate *ab* stimuli for each pigeon were then initiated with all four comparison stimuli available, followed later by *bc* training.

<<Insert Figure 3 about here.>>

<<Insert Table 3 about here.>>

With only a single chamber available for this experiment, we ran sessions every day including weekends, and to get all six birds through in a reasonable time we arranged intertrial

intervals of just 1 s. Each trial ended with a single comparison peck, and we included a correction procedure in which trials without a match were repeated until a peck on the correct comparison key produced the reinforcer (food). Predetermined irregular trial sequences were designed to counterbalance samples, distributions of comparisons across keys, and trial N to trial N+1 sequences of stimuli and locations. With these contingencies, accuracies over the first one or two matching sessions were typically lower than the chance level of 25%, probably because of the correction procedure, but they rose quickly. Based on data from the first five sessions of *ab* training immediately followed by five sessions of *bc* training (see Table 3), accuracies were typically highest with colors as samples but the stimuli in each modality were discriminable, though most levels at this stage of the research were far short of what would be appropriate for the transitivity tests for which we were planning.

By this time, recording data by hand from electromechanical counters was ancient history. Now we had computers and printers. We didn't have the capacity to save individual trials for later analyses, but we had enough to record data matrices of trial N versus trial N+1 pecks for both comparison stimuli and comparison locations separated out by whether or not the peck on trial N had produced a reinforcer. Figure 4 shows a sample data printout for Pigeon 76 on its first session of *ab* training.

<<Insert Figure 4 about here.>>

For our purposes, the relevant parts of the printout are the four 4x4 arrays under the heading "Autocorrelation matrices." The top two are based on comparison-stimulus responses and the bottom two on comparison-location responses across trials, with rows for trial N and columns for trial N+1. The left column of matrices shows all responding in which the peck on trial N was not reinforced; the right column shows all responding in which it was. The

diagonals, then, provide the totals for which the stimulus or the key pecked on trial N+1 was the same as the one pecked on trial N. Consider, for example, the upper left matrix, showing the stimulus pecked on trial N+1 given that the peck on trial N was not reinforced. The sum of the diagonal (59, 49, 56, 74) is 238, and the sum of the entire matrix is 477; thus, the probability of a peck on the same stimulus on trial N+1 given that the trial-N peck was not reinforced is $238/477$, or .50. Similarly, for the upper right matrix, based on trials in which the trial-N peck was reinforced, the sum of the diagonal is 12 and the sum of the entire matrix is 60; this probability is $12/60 = .20$. In other words, in this session the probability of a trial N+1 peck on the same stimulus as pecked on trial N was lower if the trial-N peck was reinforced than if it was not.

We had six birds and saw all of them through further sessions of *ab* and *bc* training. Our plans had been to continue with transitivity tests and other procedures. But difficulties with lab funding, staffing and the increasing obsolescence of our equipment limited our options. Perhaps we were also too strict with our criteria. We had originally planned to start transitivity tests with any bird only if its accuracy levels exceeded 90% in both *ab* and *bc* performances. Although we did get some results that seemed promising with a couple of pigeons, we had to close down that part of the lab before we could continue further. But that's another story. Here we are concerned with the N versus N+1 data during the five sessions each of *ab* and *bc* acquisition for these six pigeons, as summarized in Table 4.

<<Insert Table 4 about here.>>

The data are presented both as percentages and as number of trials, mainly because the latter show that the general findings were maintained across different levels of accuracy and different proportions of nonreinforced and reinforced responding. Whether we look at trial N to trial N+1 pecks on the same stimulus or on the same key, or early during training versus later

during training (Day 1 versus Day 5), this table provides no evidence with regard either to stimuli or to locations that reinforcing a Trial-N peck raises the probability that the same peck will occur on Trial N+1. Of eighteen cases involving comparison stimuli, none produced a higher probability after a trial that ended with a reinforcer than after a trial that didn't. Of eighteen involving location, only three produced higher probabilities (26% versus 32% for Bird 76; 30% versus 37% for Bird 80; 34% versus 37% for Bird 79) with one case in which they were essentially equal (25% versus 25% for Bird 81); all these occurred in the fifth session of *ab* training, though we had anticipated that the phenomenon, if observed, would be most likely early in acquisition. Furthermore, in several cases showing more repetitions after trial N without a reinforcer than trial N with a reinforcer the differences were substantial (e.g., the first entry for Bird 76, 50% versus 20%).

The conclusion to be drawn from these data seems to be that reinforcing a response makes it less likely rather than more likely. If we're reluctant to accept that conclusion, we must ask what's wrong with the experiment. For example, we could argue that the correction procedure guarantees that early acquisition will include many cases in which a reinforcer follows a peck after a switch from a different key or stimulus on the previous trial. That also might account for the attenuation of the effect after a longer history in the procedure, as errors decrease and corrections occur less often (Day 1 versus Day 5). Or perhaps an even simpler answer is that responses during correction should not count toward the data, and that the differences would disappear if we excluded all but the first response whenever a new match was presented. Unfortunately, (1) the data were not saved in a way that would allow such an analysis, and (2) even if that was possible all of the reinforced N+1 trials following an N trial with correction would also have to be excluded. We might have run acquisition without correction, but without

correction it is difficult to get high levels of accuracy in arbitrary matching. At least this study extended its measures from a single dimension of responding (IRTs) to two dimensions (stimulus pecked and location pecked). And we had not exhausted the alternatives. For example, perhaps reinforcers had raised the probability of some property of responding other than the stimulus and location dimensions that we had chosen to measure.

Study 3: Two-key response sequences

This third study arose as part of an attempt to make sense of some findings in a study of human behavior with UMBC undergraduates (Shimoff & Catania, 1998). A joystick allowed us to shape and measure both the direction and latency of movement, with points arranged as the nominal reinforcers. We obtained high correlations between trial-N+1 responding and reinforced trial-N responding, but they depended on whether the human's verbal behavior described direction or time contingencies and not on whether that dimension was the one upon which consequences depended. Here again were paradoxical findings. The property of responding most likely to be affected by point deliveries depended on verbal behavior. But if points were reinforcers, they should have affected the likelihood of whatever they followed and not just the dimension for which verbal behavior had been established. To explore the properties of these contingencies independently of verbal behavior, we arranged a pigeon analogue. Creating nonhuman analogues forces us to make explicit our assumptions about the variables that contribute to complex behavior.

We introduced three pigeons to a two-key procedure in which the keys irregularly alternated between green on both keys or red on both keys. In green trials, the left key was lit and the first peck on that key also lit the right key; a right peck after at least # left pecks was

reinforced without regard to the time from the last left peck to the right peck; a right peck after fewer than $\#$ left pecks produced a brief intertrial interval. Similarly, in red trials the left key was lit and the first peck also lit the right key; in this condition a right peck was reinforced if at least t s had elapsed since the first left peck; a right peck after fewer than t s produced a brief intertrial interval. In each session, a shaping procedure was arranged within each condition; during green a reinforced peck automatically incremented the required minimum number of left pecks ($\#A-B$), and during red a reinforced peck incremented the minimum required time from the last left peck to a right peck ($tA-B$). The requirements returned to their starting values at the beginning of the next session but were never decremented within a session. The correlations we wanted to measure called for a range of response values along these dimensions, so these procedures were designed to change response values over the course of the session for at least one of the two response dimensions, $\#A-B$ and $tA-B$.

Figure 5 provides a sample from one pigeon of the autocorrelation scatterplots printed out at the end of each session, with N values along the x-axis and $N+1$ values along the y-axis. If $N+1$ data were highly correlated with N data, the data would begin at the origin and fall along the diagonal. As shown in this example, these correlations were ordinarily relatively low and occasionally even negative.

<<Insert Figure 5 about here.>>

Table 5 shows, for each pigeon in each condition (green and red) and for each of the two response dimensions, trial N to trial $N+1$ autocorrelations with and without reinforcers on trial N . The table includes the first session in which sufficient responses occurred for correlations to be calculated, the one that included trial 1000, and the one that included trial 2000. In the differences between post nonreinforcement (*nr*) and post reinforcement (*rf*) correlations these

data provide little evidence that early in training pigeons are more likely to repeat on trial N+1 what was followed by a reinforcer on trial N, though correlations overall tended to increase over continued sessions. Of the differences between Trial N *rf* and *nrf* autocorrelations where the correlations matched the reinforcement contingencies (#A-B for green and tA-B for red), eight of eighteen were greater than zero; of those that did not match (tA-B for green and #A-B for red), nine of eighteen were greater than zero. Responding was maintained and differentiated throughout these procedures, but we were unable to capture an effect on response N+1 of whether response N had produced a reinforcer.

<<Insert Table 5 about here.>>

Variability, Stereotypy and the Genesis of Operant Classes

Of course we know that single reinforcers have effects. We see them in Skinner's earliest experiments (Skinner, 1930, 1938) as well as in more recent research (Hunter & Rosales-Ruiz, 2019; Iversen & Mogensen, 1988). We must take care, however, to distinguish between the stimulus effects of delivering reinforcers and the operant effects of the contingent relations between responses and reinforcers. For example, the delivery of food reinforcers may raise the likelihood of responses related to eating (in the pigeon, one of those is pecking), regardless of whether the relation between the response and the reinforcer is contingent (e.g., Staddon & Simmelhag, 1971). "In such a bird as the pigeon, pecking has a certain genetic unity" (Ferster & Skinner, 1957, p. 7). In his treatment of superstition Skinner (1948) emphasized the effects of a single contiguity, but we know now that behavior is ordinarily not well maintained by noncontingent reinforcers (e.g., Catania & Keller, 1981; Katz & Catania, 2005). Long-lasting effects can follow from just a few instances of reinforcement (e.g., Neuringer, 1970), but in cases

we're tempted to label as superstitious it may be that the relation between the reinforced response and the reinforcer simply isn't arbitrary enough. Furthermore, when these concerns are extended to human behavior we must also be cautious about the possible role of verbal behavior (e.g., Shimoff & Catania, 1998).

We see immediate effects of reinforcers whenever we engage in the successful shaping of responses such as a pigeon's key pecks or a rat's lever presses. These are examples of selection at work, and selection at the level of operants shares features with Darwinian selection at the level of phylogeny (Catania, 1978, 1987). It remains tempting to look for the effects of single reinforcers, but implicit in that enterprise is the assumption that the effect of a reinforcer is to generate a new response just like the one that produced it. That assumption drove the studies I've reported here. But if that's all that reinforcement does, it's as inadequate for enabling operant selection as Mendelian genetics was for providing the variations necessary for Darwinian natural selection to work. The assumption seems innocuous enough but may not have served us well either within or outside our discipline. As Sidman (1989) illustrated in his treatment of coercion and its fallout, our characterizations of our basic phenomena can have consequences far beyond our discipline.

Sidman's examples involved aversive control, but problems have also arisen from misguided applications to human behavior of what we know about reinforcement and shaping (e.g., Catania, 2014). A case in point is research by Schwartz (1981), which, according to his own account (Schwartz, 1986), was designed to undermine Skinner's position by showing that reinforcement engenders stereotyped behavior and opposes variability. He concluded that variable sequences of a pigeon's pecks on an array of keys cannot be shaped (Schwartz, 1982), a conclusion refuted by a demonstration of the opposite effect after artifactual constraints on the

required sequences were removed (Page & Neuringer, 1985). Furthermore, his demonstration of an allegedly similar effect in human behavior (Schwartz, 1982) failed to take into account the likely verbal governance (Skinner, 1969) of the stereotyped behavior.

More important, Neuringer's correction to Schwartz's findings (Neuringer, 2002, 2004, 2014; Neuringer & Jensen, 2010) does far more than refute questionable conclusions. It bears on a fundamental aspect of our discipline: the definition of our basic units. The generic nature of operant classes (Skinner, 1935) implies variations. The details of contingencies and their stability over time may push distributions toward broader variability or toward narrower stereotypies, as illustrated by the extreme differences between Neuringer's and Schwartz's outcomes (see also Neuringer, 1992). Particular instances are always embedded in populations or classes of responses. We should no more be troubled by ambiguities about their boundaries than Darwin (1859) was about the classes called species: ". . . we shall have to treat species in the same manner as those naturalists treat genera, who admit that genera are merely artificial combinations made for convenience" (p. 485); Darwin went on to refer to "the vain search for the undiscovered and undiscoverable essence of the term species" (p. 485). Like many other significant classes, both species and operants have fuzzy boundaries (Catania, 1978).

As Cumming (1999) put it in discussing the evolution of Thorndike's Law of Effect, "The response must be emitted before it can be reinforced" (p 430). When the first instances of an operant emerge from a pre-existing population, we begin our measurements with those first instances and often ignore their precursors. When reinforcement contingencies are repeated, they also pull those variations together into the classes we call operants (Catania, 1973; Skinner, 1935). Operant classes are created by common contingencies. It would be nice to decompose the properties of responding during early acquisition so that we could examine the details

quantitatively (but see Sidman, 2012) . As these studies have testified, however, it isn't easy and maybe it isn't even feasible.

For example, none of my three studies took into account the possibility that the effect of a reinforcer isn't limited to its relation to the most recent response (usually the one that produced it). We've long known that each reinforcer can contribute to continued responding based on its relation to all of the responses leading up to it (e.g., Catania, 1971; Catania et al., 2015; Dews, 1960, 1962; Killeen, 2011). If a pigeon's reinforced responding is determined not just by the most recent response but also by the population of responses reinforced over some extended period, then the effects of the most recent reinforcer on current responding may remain visible only over the very few first instances of reinforcement in a given setting, when that population is still exceedingly small.

An emitted response has many properties. If a lever press that produces a reinforcer has a certain force and duration and topography, all may be shaped by the contingency, but we must not assume that all will be equally affected. It isn't sufficient to look just at one dimension, such as IRT or location or number of pecks. Only repeated instances in which contingencies operate along one dimension but not the others can determine which are eventually selected. When we shape behavior we count not only on the replication of some variation of the response we just reinforced but also on selection specific to the response dimension along which shaping is arranged. This progression from responses as units to studies that concurrently examined different dimensions of responding was the one followed by my three studies. For that reason, I'm no longer quite sure that they produced solely negative results.

This fittingly brings us back to Sidman and another of his major contributions. He extended our behavioral taxonomy by showing us how to arrange higher-order contingencies that

created the stimulus classes he called equivalences (Sidman, 1994). His own methodological tools served him well as he extended our operant taxonomy (e.g., Sidman, 1971; Sidman et al., 1974; Sidman & Stoddard, 1967). One key element was his taking advantage of the extended relations allowed by expanding three-term contingencies to four or more terms, in what have been called conditional discriminations. These higher-order contingencies are not common in natural environments, and we might be reminded of Skinner's (1957) argument that the contingencies creating the verbal classes we call abstractions must be arranged by verbal communities. Like operants and species, these classes too have fuzzy boundaries (I could offer up my increasingly illegible handwriting over time as an example).

Just as lever presses may include those with left paw, right paw, both paws, chin or rump, all held together within an operant class by common contingencies, so also equivalence classes can include many arbitrary variants (e.g., complex visual properties of stimuli), again all held together by common contingencies. The written members of the class we call the letter A (or *a*) do not share common topographies and their spoken equivalents have no visual properties at all. Yet they remain equivalent members of a class. They can also share membership in other overlapping classes (upper versus lower case, spoken versus written, cursive versus block script, etc.), so that we can treat complex behavior in terms of networks of interlocking contingencies. And we can extend the account to Sidman's own corpus of work, as we find commonalities running across his treatments of methodology and aversive control and equivalence classes. The common consequences that hold it together include the work we do that builds upon his fundamental contributions.

References

- Anger, D. (1956). The dependence of interresponse times upon the relative reinforcement of different interresponse times. *Journal of Experimental Psychology*, 52(3), 145-161. <https://doi.org/10.1037/h0041255>
- Catania, A. C. (1971). Reinforcement schedules: The role of responses preceding the one that produces the reinforcer. *Journal of the Experimental Analysis of Behavior*, 15(3), 271-287. <https://doi.org/10.1901/jeab.1971.15-271>
- Catania, A. C. (1972). Switching circuitry (Journal Supplement Abstract Service Manuscript No. 53). *JSAS Catalog of Selected Documents in Psychology*, 2(4), 67 pp.
- Catania, A. C. (1973). The concept of the operant in the analysis of behavior. *Behaviorism*, 1(2), 103-116. <https://www.jstor.org/stable/27758804>
- Catania, A. C. (1978). The psychology of learning: some lessons from the Darwinian revolution. *Annals of the New York Academy of Sciences*, 309, 18-28. <https://doi.org/10.1111/j.1749-6632.1978.tb29439.x>
- Catania, A. C. (1987). Some Darwinian lessons for behavior analysis. A review of Peter J. Bowler's "The eclipse of Darwinism.". *Journal of the Experimental Analysis of Behavior*, 47(2), 249-257. <https://doi.org/10.1901/jeab.1987.47-249>
- Catania, A. C. (2002). The watershed years of 1958-1962 in the Harvard Pigeon Laboratory. *Journal of the Experimental Analysis of Behavior*, 77(3), 327-345. <https://doi.org/10.1901/jeab.2002.77-327>
- Catania, A. C. (2011). Model versus model or model versus data? A commentary on Berg and McDowell's "Quantitative, steady-state properties of Catania's computational model of

the operant reserve.”. *Behavioural Processes*, 87(1), 84-87.

<https://doi.org/10.1016/j.beproc.2010.12.006>

Catania, A. C. (2014). Behavior analysis as a biological science: An alternate history. *European Journal of Behavior Analysis*, 15(1), 25-30.

Catania, A. C., & Keller, K. J. (1981). Contingency, contiguity, correlation, and the concept of causation. In P. Harzem & M. D. Zeiler (Eds.), *Predictability, correlation, and contiguity. Advances in analysis of behaviour, Volume 2* (pp. 125-167). Chichester: Wiley.

Catania, A. C., Reilly, M. P., Hand, D., Kehle, L. K., Valentine, L., & Shimoff, E. (2015). A quantitative analysis of the behavior maintained by delayed reinforcers. *Journal of the Experimental Analysis of Behavior*, 103(2), 288-331. <https://doi.org/10.1002/jeab.138>

Catania, A. C., & Reynolds, G. S. (1968). A quantitative analysis of the behavior maintained by interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 11(3), 327-383. <https://doi.org/10.1901/jeab.1968.11-s327>

Cumming, W. (1999). A Review of Geraldine Jonçich’s The Sane Positivist: A biography of Edward L. Thorndike. *Journal of the Experimental Analysis of Behavior*, 72(3), 429. <https://doi.org/10.1901/jeab.1999.72-429>

Darwin, C. (1859). *On the origin of species by means of natural selection*. London: John Murray.

Dews, P. B. (1960). Free-operant behavior under conditions of delayed reinforcement: I. CRF-type schedules. *Journal of the Experimental Analysis of Behavior*, 3(3), 221-234. <https://doi.org/10.1901/jeab.1960.3-221>

Dews, P. B. (1962). The effect of multiple S^{Δ} periods on responding on a fixed-interval schedule. *Journal of the Experimental Analysis of Behavior*, 5(3), 369-374.

<https://doi.org/10.1901/jeab.1962.5-369>

Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement*. New York: Appleton-Century-Crofts.

Hunter, M., & Rosales-Ruiz, J. (2019). The power of one reinforcer: The effect of a single reinforcer in the context of shaping. *Journal of the Experimental Analysis of Behavior*, 111(3), 449-464. <https://doi.org/10.1002/jeab.517>

Iversen, I. H., & Mogensen, J. (1988). A multipurpose vertical holeboard with automated recording of spatial and temporal response patterns for rodents. *J. Neurosci. Methods*, 25(3), 251-263. [https://doi.org/10.1016/0165-0270\(88\)90140-9](https://doi.org/10.1016/0165-0270(88)90140-9)

Katz, I., & Catania, A. C. (2005). Concurrent performances: Extinction, noncontingent reinforcement and variably delayed reinforcement. *European Journal of Behavior Analysis*, 6(1), 95-108. <https://doi.org/10.1080/15021149.2005.11434254>

Keister, W., Ritchie, A. E., & Washburn, S. H. (1951). *The design of switching circuits*. Princeton, NJ: Van Nostrand.

Keller, F. S., & Schoenfeld, W. N. (1949). The psychology curriculum at Columbia College. *American Psychologist*, 4(6), 165-172. <https://doi.org/10.1037/h0057770>

Keller, F. S., & Schoenfeld, W. N. (1950). *Principles of psychology*. New York: Appleton-Century-Crofts.

Killeen, P. R. (2011). Models of trace decay, eligibility for reinforcement, and delay of reinforcement gradients, from exponential to hyperboloid. *Behavioural Processes*, 87(1), 57-63. <https://doi.org/10.1016/j.beproc.2010.12.016>

Killeen, P. R., & Sitomer, M. T. (2003). MPR. *Behavioural Processes*, 62(1-3), 49-64.

[https://doi.org/10.1016/S0376-6357\(03\)00017-2](https://doi.org/10.1016/S0376-6357(03)00017-2)

Neuringer, A. (1970). Superstitious key pecking after three peck-produced reinforcements.

Journal of the Experimental Analysis of Behavior, 13(2), 127-134.

<https://doi.org/10.1901/jeab.1970.13-127>

Neuringer, A. (1992). Choosing to vary and repeat. *Psychological Science*, 3(4), 246-250.

<https://doi.org/10.1111/j.1467-9280.1992.tb00037.x>

Neuringer, A. (2002). Operant variability: Evidence, functions, and theory. *Psychonomic Bulletin*

and Review, 9(4), 672-705. <https://doi.org/10.3758/BF03196324>

Neuringer, A. (2004). Reinforced variability in animals and people. *American Psychologist*,

59(9), 891-906. <https://doi.org/10.1037/0003-066X.59.9.891>

Neuringer, A. (2014). Operant variability and the evolution of volition. *International Journal of*

Comparative Psychology, 27(2), 204-223. <https://escholarship.org/uc/item/0s78k28c>

Neuringer, A., & Jensen, G. (2010). Operant variability and voluntary action. *Psychological*

Review, 117(3), 972-993. <https://doi.org/10.1037/a0019499>

Page, S., & Neuringer, A. (1985). Variability is an operant. *Journal of Experimental Psychology:*

Animal Behavior Processes, 11(3), 429-452. <https://doi.org/10.1037/0097-7403.11.3.429>

Schwartz, B. (1981). In pursuit of B. F. Skinner. *Swarthmore Alumni Bulletin(March)*, 12-16.

Schwartz, B. (1982). Failure to produce response variability with reinforcement. *Journal of the*

Experimental Analysis of Behavior, 37(2), 171-181. [https://doi.org/10.1901/jeab.1982.37-](https://doi.org/10.1901/jeab.1982.37-171)

[171](https://doi.org/10.1901/jeab.1982.37-171)

Schwartz, B. (1986). *The battle for human nature*. New York: Norton.

Shimoff, E., & Catania, A. C. (1998). The verbal governance of behavior. In K. A. Lattal & M. Perone (Eds.), *Handbook of research methods in human operant behavior* (pp. 371-404). New York: Plenum.

Sidman, M. (1953). Two temporal parameters in the maintenance of avoidance behavior by the white rat. *Journal of Comparative and Physiological Psychology*, 46(4), 253-261.
<https://doi.org/10.1037/h0060730>

Sidman, M. (1958). By-products of aversive control. *Journal of the Experimental Analysis of Behavior*, 1(3), 265-280. <https://10.1901/jeab.1958.1-265>

Sidman, M. (1960). *Tactics of scientific research*. New York: Basic Books.

Sidman, M. (1971). The behavioral analysis of aphasia. *Journal of Psychiatric Research*, 8(3-4), 413-422. [https://doi.org/10.1016/0022-3956\(71\)90034-3](https://doi.org/10.1016/0022-3956(71)90034-3)

Sidman, M. (1987). Two choices are not enough. *Behavior Analysis*, 22(1), 11-18.

Sidman, M. (1989). *Coercion and its fallout*. Boston: Authors Cooperative.

Sidman, M. (1994). *Equivalence relations and behavior: A research story*. Boston: Authors Cooperative.

Sidman, M. (2012). Why mathematical treatments of operant behavior? *European Journal of Behavior Analysis*, 13(2), 255-257. <https://doi.org/10.1080/15021149.2012.11434436>

Sidman, M., Cresson, O., Jr., & Willson-Morris, M. (1974). Acquisition of matching to sample via mediated transfer. *Journal of the Experimental Analysis of Behavior*, 22(2), 261-273.
<https://doi.org/10.1901/jeab.1974.22-261>

Sidman, M., & Stoddard, L. T. (1967). The effectiveness of fading in programming a simultaneous form discrimination for retarded children. *Journal of the Experimental Analysis of Behavior*, 10(1), 3-15. <https://doi.org/10.1901/jeab.1967.10-3>

Sidman, M., Wynne, C. K., Maguire, R. W., & Barnes, T. (1989). Functional classes and equivalence relations. *Journal of the Experimental Analysis of Behavior*, 52(3), 261-274.

<https://doi.org/10.1901/jeab.1989.52-261>

Skinner, B. F. (1930). On the conditions for elicitation of certain eating reflexes. *Proceedings of the National Academy of Sciences*, 16(6), 433-438. <https://doi.org/10.1073/pnas.16.6.433>

Skinner, B. F. (1935). The generic nature of the concepts of stimulus and response. *Journal of General Psychology*, 12, 40-65. <https://doi.org/10.1080/00221309.1935.9920087>

Skinner, B. F. (1938). *The behavior of organisms: An experimental analysis*. New York: Appleton-Century-Crofts.

Skinner, B. F. (1948). "Superstition" in the pigeon. *Journal of Experimental Psychology: Animal Behavior Processes*, 38(2), 168-172. <https://doi.org/10.1037/h0055873>

Skinner, B. F. (1956). A case history in scientific method. *American Psychologist*, 11(5), 221-233. <https://doi.org/10.1037/h0047662>

Skinner, B. F. (1957). *Verbal behavior*. New York: Appleton-Century-Crofts.

Skinner, B. F. (1969). An operant analysis of problem solving. In *Contingencies of reinforcement* (pp. 133-157). New York: Appleton-Century-Crofts.

Staddon, J. E. R., & Simmelhag, V. L. (1971). The "superstition" experiment: A reexamination of its implications for the principles of adaptive behavior. *Psychological Review*, 78(1), 3-43. <https://doi.org/10.1037/h0030305>

Urcuioli, P. J. (2008). Associative symmetry, antisymmetry, and a theory of pigeons' equivalence-class formation. *Journal of the Experimental Analysis of Behavior*, 90(3), 257-282. <https://doi.org/10.1901/jeab.2008.90-257>

Running head: Tactical Contingencies

Wright, A. A. (2012). Memory processing. In T. R. Zentall & E. A. Wasserman (Eds.), *The Oxford handbook of comparative cognition* (pp. 239-260). Oxford: Oxford University Press.

Associate Editor: Iver Iversen

Editor-in-Chief: Mark Galizio

Table 1. Interresponse time (IRT) distributions before and after reinforced long and reinforced short IRTs for Bird 201, an experimental naive pigeon in its second VI session, which was the first one that satisfied the setup-to-reinforcer time criterion. The reinforcement cells of the distribution are identified by underlined italics; no adjustments were made for the unequal cell durations. *Setup to rfr* shows time in *s* from a reinforcer setup to production of that reinforcer by a response. Details in text.

BIRD 201: Relative Frequencies -- VI 180-s, 30-s window									
seconds		<.5	.5-1	1-2	2-3	3-4	>4	N	Setup to rfr (s)
Before		.34	.08	.28	.12	.09	.09	188	1.6
After		<u>.41</u>	.10	.22	.14	.06	.08	274	
Before		.37	.13	.23	.12	.06	.10	254	6.0
After		.39	.09	.23	<u>.15</u>	.05	.09	259	

Table 2. Interresponse time (IRT) distributions and IRTs per opportunity (IRTs/op) before and after reinforced long and reinforced short IRTs for Birds 278 and 279, both of which had extensive VI histories (Catania & Reynolds, 1968). The reinforcement cells of the distribution are identified by underlined italics; no adjustments were made for the unequal cell durations. *Setup to rfr* shows time in s from a reinforcer setup to production of that reinforcer by a response. Details in text.

BIRD 278: Relative Frequencies -- VI 150-s, 30-s window									
seconds		<.5	.5-1	1-2	2-3	3-4	>4	N	Setup to rfr (s)
Before		.48	.25	.23	.02	.00	.00	444	2.3
After		<u>.33</u>	.24	.28	.06	.04	.05	269	
Before		.51	.26	.21	.01	.01	.00	468	3.4
After		.40	.26	<u>.22</u>	.05	.03	.05	286	

BIRD 278: IRTs/op -- VI 150-s, 30-s window									
seconds		<.5	.5-1	1-2	2-3	3-4	>4	N	Setup to rfr (s)
Before		.47	.48	.87	.73	--	--	444	2.3
After		<u>.32</u>	.36	.66	.41	.43	--	269	
Before		.50	.53	.89	.58	.80	--	468	3.4
After		.39	.43	<u>.64</u>	.40	.38	--	286	

BIRD 279: Relative Frequencies -- VI 180-s, 60-s window									
seconds		<.5	.5-1	1-2	2-3	3-4	>4	N	Setup to rfr (s)
Before		.44	.17	.24	.13	.02	.00	705	3.7
After		<u>.39</u>	.26	.16	.13	.05	.01	677	
Before		.41	.14	.30	.12	.02	.00	659	8.4
After		.38	.19	.20	.17	.05	.02	597	

BIRD 279: IRTs/op --VI 180-s, 60-s window									
seconds		<.5	.5-1	1-2	2-3	3-4	>4	N	Setup to rfr (s)
Before		.43	.30	.61	.86	.87	--	705	3.7
After		<u>.39</u>	.43	.47	.68	.82	--	677	
Before		.40	.24	.67	.82	.89	--	659	8.4
After		.36	.31	.47	.72	.76	--	597	

Table 3. Assignment of stimuli to the *a*, *b* and *c* sets for each pigeon and highest accuracies (%) obtained with each in the last of five of about thirty arbitrary-matching sessions with a correction procedure. With four comparison stimuli, chance accuracy was 25%.

Bird	Stimuli (<i>a</i> , <i>b</i> , <i>c</i>)	Maximum Accuracy (%)	
		<i>ab</i>	<i>bc</i>
76	form, color, texture	68%	84%
78	color, texture, form	95%	83%
80	texture, form, color	51%	67%
77	color, form, texture	92%	76%
79	form, texture, color	54%	64%
81	texture, color, form	86%	92%

Note that each possible *ab-bc* combination is represented.

Table 4. Trial N+1 pecking showing p(peck same key) and p(peck same stimulus) given trial N nonreinforced (nonrf) or reinforced (rf), on the first and fifth day of *ab* training and the first day of *bc* training for six pigeons.

Bird 76		<i>ab</i> Day 1: 11% corr		<i>ab</i> Day 5: 31 % corr		<i>bc</i> Day 1: 14% corr	
		nonrf	rf	nonrf	rf	nonrf	rf
p(peck stim)	%same	50%	20%	26%	23%	37%	12%
p(peck key)			28%		32%		25%
p(peck stim)	#trials	238	12	34	14	352	7
p(peck key)			17		19		15

Bird 78		<i>ab</i> Day 1: 10% corr		<i>ab</i> Day 5: 43% corr		<i>bc</i> Day 1: 11% corr	
		nonrf	rf	nonrf	rf	nonrf	rf
p(peck stim)	%same	51%	23%	38%	23%	46%	18%
p(peck key)			28%		35%		35%
p(peck stim)	#trials	266	14	29	14	220	11
p(peck key)			17		21		21

Bird 80		<i>ab</i> Day 1: 6% corr		<i>ab</i> Day 5: 22% corr		<i>bc</i> Day 1: 14% corr	
		nonrf	rf	nonrf	rf	nonrf	rf
p(peck stim)	%same	55%	39%	30%	25%	47%	28%
p(peck key)			21%		37%		33%
p(peck stim)	#trials	310	15	62	15	164	17
p(peck key)			8		22		20

Bird 77		<i>ab</i> Day 1: 13% corr		<i>ab</i> Day 5: 53% corr		<i>bc</i> Day 1: 15% corr	
		nonrf	rf	nonrf	rf	nonrf	rf
p(peck stim)	%same	43%	27%	40%	28%	40%	30%
p(peck key)			30%		28%		33%
p(peck stim)	#trials	163	16	21	17	128	18
p(peck key)			18		17		20

Bird 79		<i>ab</i> Day 1: 2% corr		<i>ab</i> Day 5: 25% corr		<i>bc</i> Day 1: 16% corr	
		nonrf	rf	nonrf	rf	nonrf	rf
p(peck stim)	%same	54%	25%	34%	27%	41%	33%
p(peck key)			33%		37%		30%
p(peck stim)	#trials	287	3	61	16	121	20
p(peck key)			4		22		18

Bird 81		<i>ab</i> Day 1: 8% corr		<i>ab</i> Day 5: 43% corr		<i>bc</i> Day 1: 12% corr	
		nonrf	rf	nonrf	rf	nonrf	rf
p(peck stim)	%same	50%	26%	25%	25%	45%	20%
p(peck key)			23%		25%		35%
p(peck stim)	#trials	400	18	20	15	182	12
p(peck key)			16		15		21

Table 5. Autocorrelations of Trial N to Trial N+1 values with shaping of number of A pecks before a reinforced B peck (Green) or time from last A peck to reinforced B peck (Red), for three pigeons. Adapted from Shimoff & Catania (1998, Table 12.2, p. 400). Details in text.

	Session	Green (#A-B)				Red (tA-B)			
		Number		Time		Number		Time	
		nrf	rf	nrf	rf	nrf	rf	nrf	rf
Pigeon 37	a	.00	-.22	.44	.36	-.15	.00	-.16	.00
	b	.26	.37	-.15	.56	.25	.07	.13	-.14
	c	.41	.92	-.02	.26	.18	.43	.28	.22
Pigeon 72	a	-.04	.18	-.04	-.03	-.06	.59	.21	-.16
	b	-.14	-.45	.00	.28	.00	-.01	-.06	.21
	c	.21	.36	.47	.55	-.13	-.32	.03	-.33
Pigeon 98	a	-.02	-.43	.70	.28	.04	-.05	-.09	-.41
	b	.42	.07	.08	-.03	-.19	.33	.11	.23
	c	.70	.83	.16	-.23	.09	.00	.01	-.08

nrf = no reinforcer on trial N; rf = reinforcer on trial N.

a = 1st session with enough responses to calculate correlations; b = session including trial 1000;

c = session including trial 2000.

IRT study

Bird # 1298

Date 3/29/62

VI 3'

IRT distribution

	<1"	1"-1"	1"-2"	2"-3"	3"-4"	>4"	TOTAL	RESPS DURING RESET	RESPS	RFT	TIME FROM SETUP TO RFT
1' BEFORE RFT SETUP	629	170 431	192 261	56 69	11 13	2	1060 1072	348	61	16	10.3
1' AFTER RFT	415	128 391	147 263	72 116	25 44	19	806 811	251	7817	34	
RFT for IRT of <1/2"											
1' BEFORE RFT SETUP	671	173 416	177 243	51 66	9 15	6	1037 1052	333	1390	16	17.46
1' AFTER RFT	413	111 369	156 258	61 102	20 41	21	782 795	235	3026	66	
RFT for IRT of 1/2"-3"											

Figure 1. Sample data sheet showing customized layout for entries from electromechanical counters. This example also includes entries added later as checksums and for calculating other measures such as IRTs/op. Details in text.

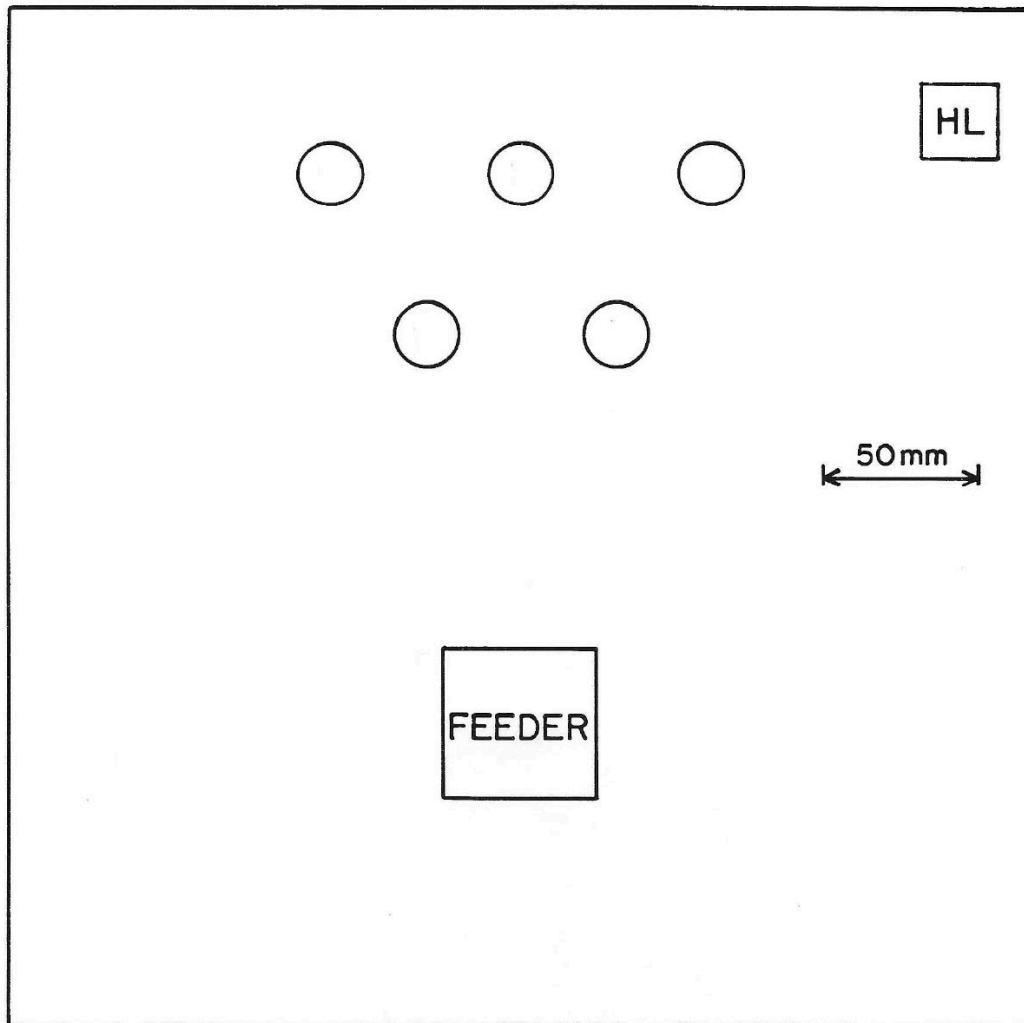


Figure 2. Schematic of pigeon panel used for arbitrary matching. A peck on the sample stimulus on the top middle key turned on four comparison stimuli on the remaining four keys (see Figure 3). A peck on the appropriate comparison key produced a reinforcer. A houselight (HL) remained lit during both trials and intertrial intervals. Details in text.

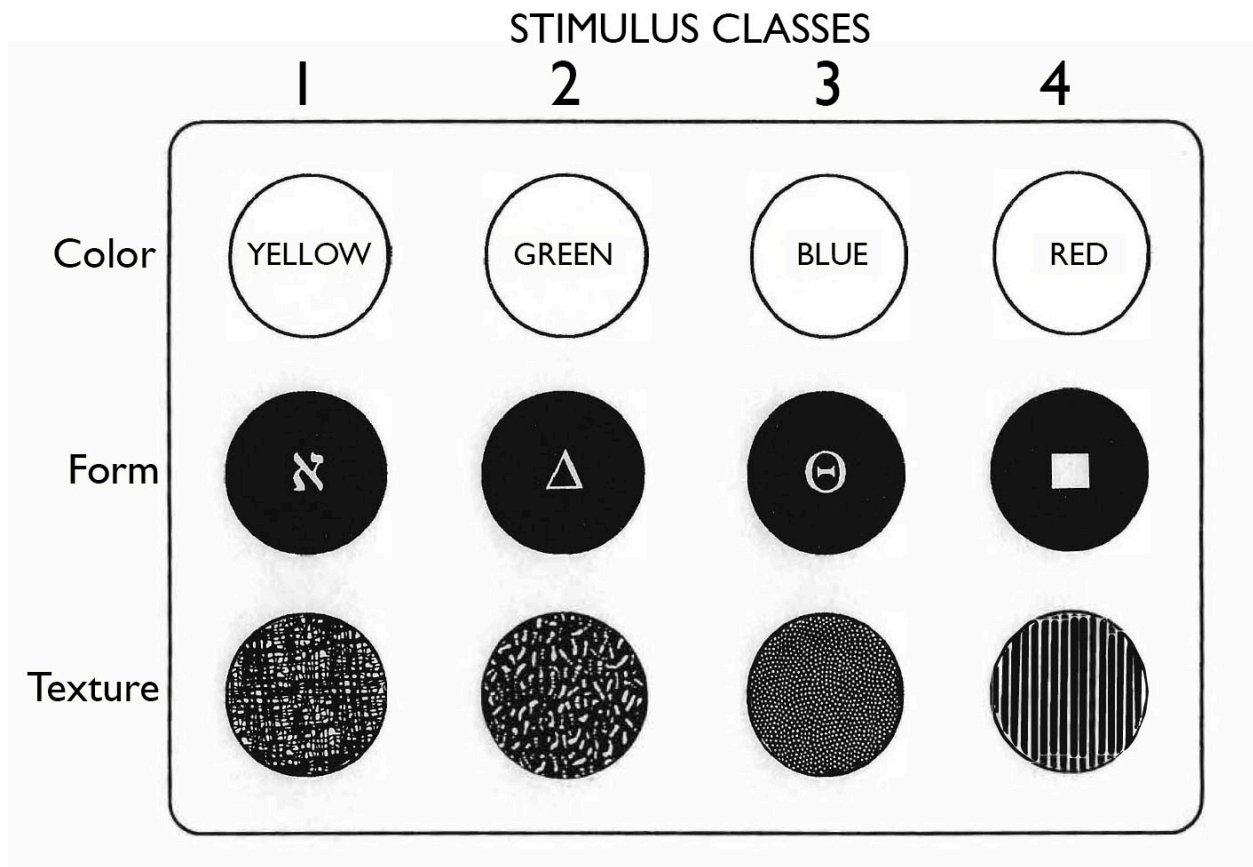


Figure 3. The four-stimulus sets (colors, forms, textures) used in training arbitrary matching.

The objective was to work with potential equivalence classes 1 through 4, with each class consisting of one color, one form and one texture (*abc*), as identified by the columns. The study was set up to examine *ab* and *bc* training, which was to be followed eventually by *ac* testing (transitivity). The stimuli were projected on the keys by standard inline projectors. As shown in Table 3, the assignment of stimuli to *a*, *b* and *c* was counterbalanced across six pigeons.

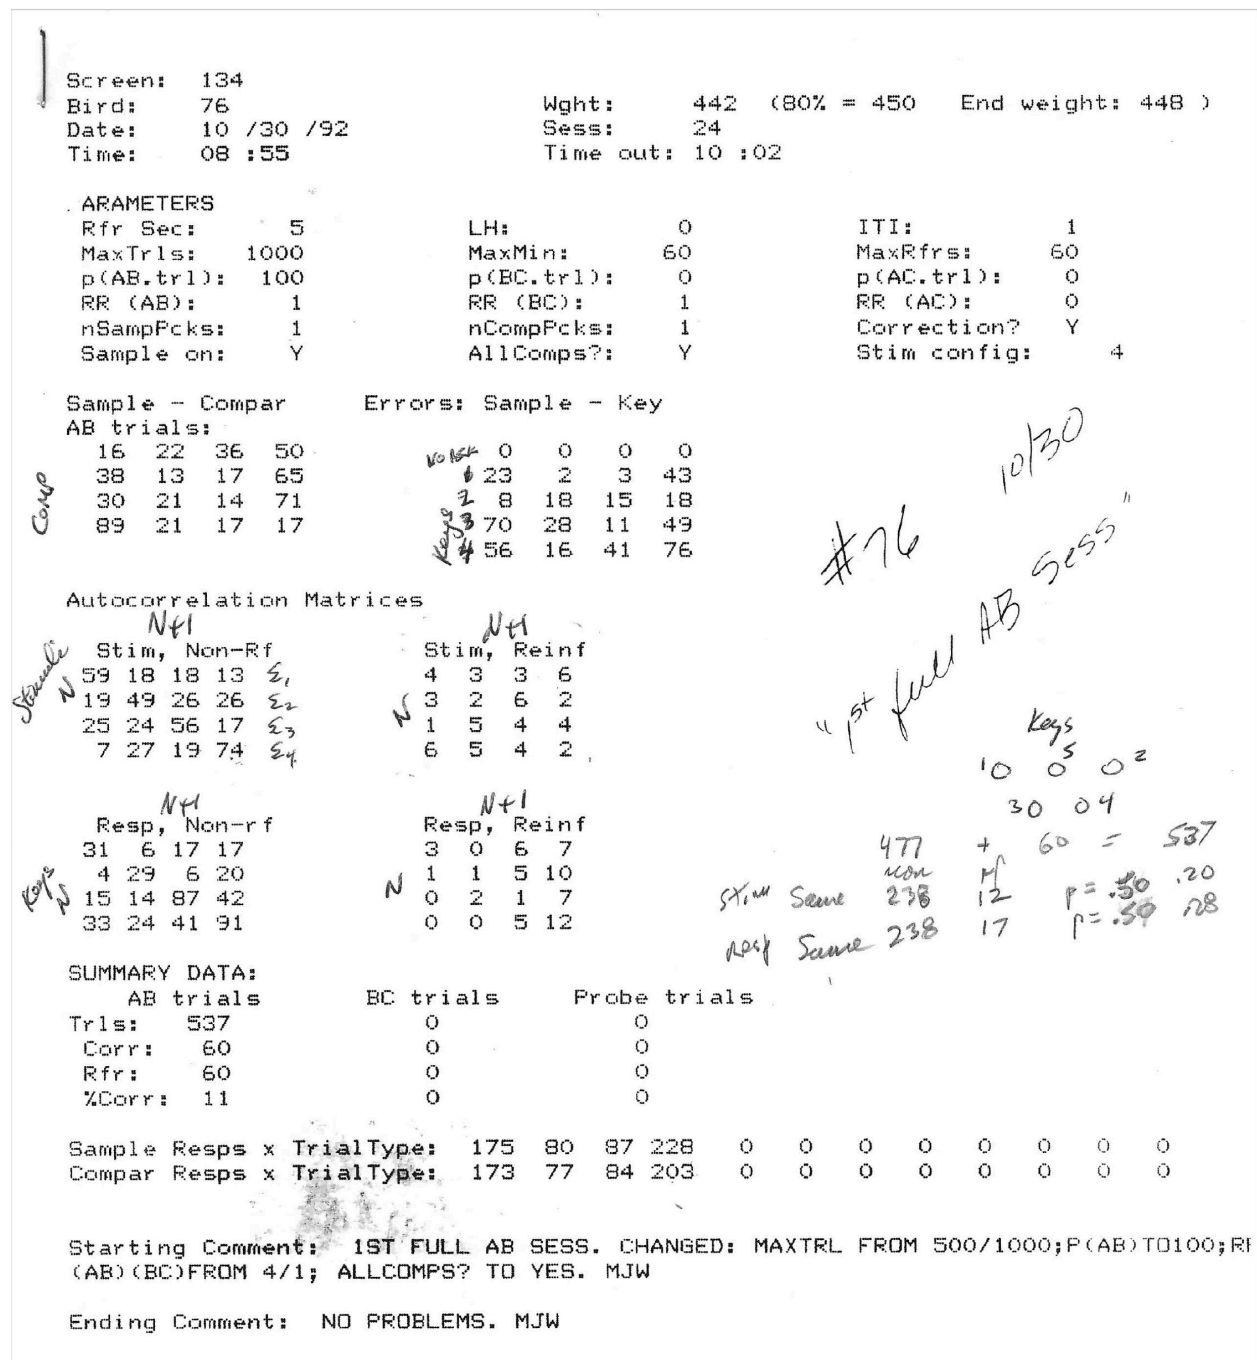


Figure 4. A sample data printout from the first *ab* session of arbitrary matching for Bird 76, showing the data matrixes from which probabilities were derived and some labels and calculations that were later added.

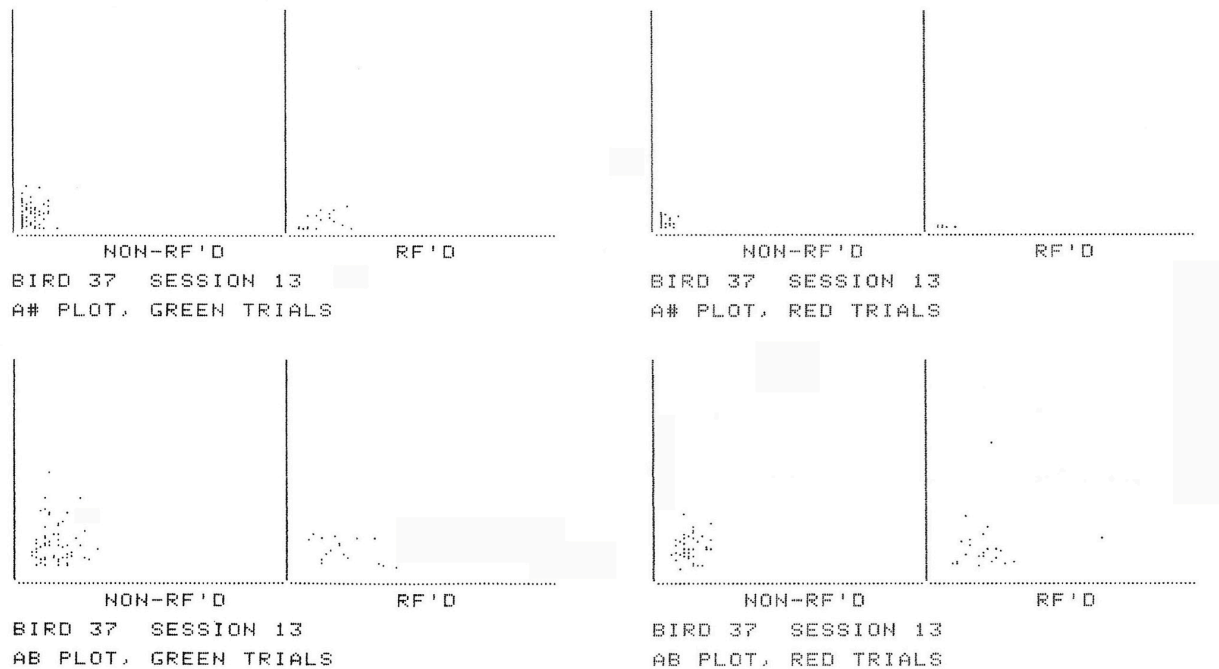


Figure 5. Autocorrelation scatterplots of trial N+1 responding (y-axis) versus trial N responding (x-axis) separated by whether trial N responding had been non-reinforced (left) or reinforced (right), for Pigeon 37, Session 13. Left graphs show responding during green, when the reinforcer depended on number of A pecks before a switch to B (#A-B); right graphs show responding during red, when the reinforcer depended on the time between the first A peck and the switch to B (tA-B). The top graphs plot #A-B within each trial; the bottom graphs plot tA-B. To provide a range of response values within each session, the data were obtained over the course of shaping longer #A-B runs during green and longer tA-B times during red. Details in text.