SELECTION DYNAMICS IN JOINT MATCHING TO RATE AND MAGNITUDE OF REINFORCEMENT

J. J. McDOWELL, ANDREI POPA, AND NICHOLAS T. CALVIN

EMORY UNIVERSITY

Virtual organisms animated by a selectionist theory of behavior dynamics worked on concurrent random interval schedules where both the rate and magnitude of reinforcement were varied. The selectionist theory consists of a set of simple rules of selection, recombination, and mutation that act on a population of potential behaviors by means of a genetic algorithm. An extension of the power function matching equation, which expresses behavior allocation as a joint function of exponentiated reinforcement rate and reinforcer magnitude ratios, was fitted to the virtual organisms' data, and over a range of moderate mutation rates was found to provide an excellent description of their behavior without residual trends. The mean exponents in this range of mutation rates were 0.83 for the reinforcement rate ratio and 0.68 for the reinforcer magnitude ratio, which are values that are comparable to those obtained in experiments with live organisms. These findings add to the evidence supporting the selectionist theory, which asserts that the world of behavior we observe and measure is created by evolutionary dynamics.

Key words: selection by consequences, behavior dynamics, computational modeling, generalized matching law, magnitude of reinforcement, concatenated matching law

A selectionist, or evolutionary, theory of behavior dynamics can be stated as a set of simple rules that cause selection, reproduction, and mutation to occur in a population of potential behaviors by means of a genetic algorithm (McDowell, 2004). McDowell, Caron, Kulubekova, & Berg (2008) studied the behavior of virtual organisms animated by this theory on concurrent schedules of reinforcement and found that it was well-described by the power function matching equation,

$$\frac{B_1}{B_2} = b \left(\frac{r_1}{r_2} \right)^a, \quad (1)$$

where the B s and r s represent response and reinforcement rates, the numerical subscripts refer to the components of a concurrent schedule, and a and b are parameters that represent sensitivity to the reinforcement rate ratio, and bias for one or the other component of the schedule. McDowell et al. also found that across a range of moderate mutation rates, the exponent in Equation 1 for their virtual organisms varied around a value of about 0.8, which is the value that is typically found in experiments

with live organisms (Baum, 1974b, 1979; McDowell, 1989; Myers & Myers, 1977; Wearden & Burgess, 1982). McDowell et al. concluded that both the form of Equation 1 and the exponent's value of about 0.8 were emergent properties of the evolutionary dynamics that animated their virtual organisms.

The effects of properties of reinforcement other than its rate of delivery have also been studied in concurrent-schedule experiments using live organisms. For example, when different reinforcer magnitudes are arranged in the two components, Equation 1, where the ratio of the magnitudes is taken to be the independent variable, has been found to provide a reasonable account of responding, but with an exponent that is usually much smaller than that typically obtained when the reinforcement rate ratio is the independent variable (Davison & Baum, 2003; Landon, Davison, & Elliffe, 2003; Maguire, Hughes, & Pitts, 2007; Schneider, 1973; Todorov, 1973; Todorov, Hanna, & Bittencourt de Sa, 1984).

The joint variation of reinforcement rate and reinforcer magnitude has also been studied in experiments with live organisms. An extension of Equation 1,

$$\frac{B_1}{B_2} = b \left(\frac{r_1}{r_2} \right)^{a_r} \left(\frac{m_1}{m_2} \right)^{a_m}, \quad (2)$$

(Baum & Rachlin, 1969; Davison & McCarthy, 1988; Killeen, 1972; Rachlin, 1971) has been

We thank William Baum, Michael Davison, and Nicolas Raillard for their helpful comments on an earlier version of this paper.

Address correspondence to J. J. McDowell, Department of Psychology, Emory University, Atlanta, GA 30322 (email: psyjjmd@emory.edu).

doi: 10.1901/jeab.2012.98-199

used to describe data from these experiments. In Equation 2 the *ms* represent reinforcer magnitudes, the subscripts on the exponents refer to the ratios they exponentiate, and all other quantities are the same as in Equation 1. Equation 2 essentially pulls the exponentiated magnitude ratio out of the bias parameter in Equation 1.

In a recent review, Cording, McLean, & Grace (2011) examined 59 sets of data from 11 experiments in which reinforcer magnitude or both reinforcer magnitude and reinforcement rate were varied in concurrent schedules. Their fits of Equation 2 to the latter data sets

accounted for 71% to 98% of the data variance, with a median of 92% of the variance accounted for by the equation. The quality of these fits, together with the authors' residual analyses, led them to conclude that Equation 2 described the data from these experiments well. Estimates of a_m (listed in their Tables 2 and 3) ranged from 0.26 to 1.82, with an average of 0.65, confirming earlier reports that the exponent on the reinforcer magnitude ratio was smaller than the exponent on the reinforcement rate ratio.

The purpose of the experiment reported in this article was to study the behavior of virtual

Fig. 1. The evolutionary theory of behavior dynamics begins with an initial population of potential behaviors from which a behavior is emitted at random. If the behavior is reinforced, then parents are chosen from the population on the basis of their similarity to the behavior that was just reinforced. If the behavior is not reinforced, then parents are chosen at random. In either case, parent genotypes are recombined to produce child behaviors that populate the next generation of potential behaviors. The new population is subjected to a small amount of mutation, and then the cycle repeats. Gray hatched boxes indicate the Darwinian processes of selection, reproduction, and mutation. S^{R+} = reinforcement.

organisms animated by the evolutionary theory on concurrent schedules where both the rate and magnitude of reinforcement in the schedule components were varied. In other words, the purpose was to test the theory under this more complicated set of experimental conditions. Before describing the details of the experiment, the theory will be reviewed briefly.

In the evolutionary theory, a population of potential behaviors evolves under the selection pressure of reinforcement from the environment. Each behavior in the population is represented by a decimal integer, which is the behavior's phenotype, and by the binary representation of that integer, which is the behavior's genotype. A behavior is emitted at random from this population at each tick of time and the resulting stream of behavior can be studied just as if it were the behavior of a live organism. If the emitted behavior is reinforced, then so-called parent behaviors are chosen from the population on the basis of their fitness, and their binary genotypes are recombined to produce child behaviors that constitute the population at the next time tick. If the emitted behavior is not reinforced, then parent behaviors are chosen at random. The fitness of a behavior is defined as the difference between its phenotype and the phenotype of the behavior that was just reinforced. A one-parameter probability density function, called the parental selection function, is used to choose parents for mating on the basis of their fitness. This single parameter is the mean of the density function and reflects the strength of the selection event, which is analogous to reinforcer magnitude and can be taken to represent it (McDowell, 2004; McDowell et al., 2008). Once a new population of potential behaviors is formed, it undergoes a small amount of mutation. This cycle, which is illustrated in Figure 1, then repeats, beginning with the emission of a behavior at random from the new population.

Specific details of the theory's implementation will be described in the Procedure section below. It is important to emphasize that these details are identical to those used in McDowell et al.'s (2008) concurrent schedule experiment, with one exception. McDowell et al. defined the fitness of a behavior as the difference between the behavior's phenotype and the phenotype at the midpoint of the

target class of behaviors from which the just-reinforced behavior was emitted. In the present experiment, the fitness of a behavior was defined as the difference between the behavior's phenotype and the phenotype of the just-reinforced behavior itself. These definitions of fitness differ only minimally, and the latter is preferable because it has greater face validity. Experiments in our laboratory indicate that the two definitions do not generate materially different behavioral outputs. In summary, the experiment described in this article did not entail special settings for or tunings of the genetic algorithm that implemented the theory. The ranges of reinforcement rates and reinforcer magnitudes studied were also comparable to those studied by McDowell et al.

METHOD

Subjects

Virtual organisms with repertoires of 100 potential behaviors at each tick of time served as subjects. Decimal integer phenotypes of the potential behaviors ranged from 0 through 1023. The corresponding binary genotypes were padded with leading zeros as necessary to obtain 10-character bit strings for all genotypes. The phenotype range was taken to be circular, that is, it wrapped around upon itself such that 0 was as close to 1023 as to 1. A virtual organism emitted a behavior at random from its repertoire of 100 potential behaviors at each tick of time. Specifically, every behavior in the repertoire was equally likely to be emitted at each time tick. Subjects were naive at the start of each concurrent schedule in the sense that their populations of potential behaviors consisted of 100 phenotypes chosen at random from the 0 through 1023 range.

Apparatus and Materials

Experiments were run on commercially available off-the-shelf desktop and notebook computer hardware. The program that created and animated the virtual organism, called the organism code, was written by the first author in VB 2010 implemented in the Microsoft Visual Studio 2010 Integrated Development Environment (IDE). A separate piece of code, called the laboratory code, was

written by the third author, also in VB 2010 implemented in the Microsoft Visual Studio 2010 IDE. The laboratory code arranged the concurrent schedules and recorded the virtual organisms' behavior. At each tick of time it detected the emitted behavior. If this behavior came from one of the target classes (defined below), the laboratory code consulted the component schedules to determine whether reinforcement was available, and if so, delivered it with the appropriate magnitude to the virtual organism. The organism code was responsible for selecting parent behaviors from its population of potential behaviors, recombining parent pairs to produce child behaviors, adding mutation, and emitting a new behavior. Data were stored in standard databases and were analyzed with standard software.

Procedure

Two target classes of behavior were defined by the integer phenotypes 471 through 511 and 512 through 552. Behaviors in these classes may be taken to represent successful left-alternative and right-alternative responses. The classes are separated by the maximum Hamming distance (Russell & Norvig, 2003) for 10-character bit strings of 10. The Hamming distance between two bit strings of equal length is the number of bits that must be changed to convert one string into the other. The Hamming distance separating two classes of behavior has the same function as a changeover delay (Popa & McDowell, 2010).

The fitness of every behavior in the population of potential behaviors was calculated each time a reinforcer was delivered. Fitness was defined as the absolute value of the difference between a behavior's phenotype and the phenotype of the just-reinforced behavior, or 1024 minus this absolute value, whichever was smaller. Hence, behaviors lying closer to the just-reinforced behavior along the circular phenotype range were fitter. Notice that this method assigns smaller fitness values to fitter behaviors because the distance between these behaviors and the just-reinforced behavior is smaller. For example, the fittest behavior, which would be the behavior that was just reinforced, has a fitness value of zero.

After reinforcement, 100 pairs of parents were chosen on the basis of their fitness using

a linear function that expresses the probability density, p , associated with a parent being chosen for mating as a function of that parent's fitness, x ,

$$p(x) = -\frac{2}{9\mu^2}x + \frac{2}{3\mu}, \quad (3)$$

where $0 \leq x \leq 3\mu$, and μ is the mean of the density function. This equation, which is referred to as the parental selection function, is the simplest linear density function that can be written with a single parameter (McDowell, 2004). The smaller the mean, μ , of this function, the stronger the selection event. This is because smaller means result in fitter parents being selected for mating. Hence larger reinforcer magnitudes can be arranged by using parental selection functions with smaller means. Parents were selected by drawing a fitness value at random from the density function using the procedure described by McDowell (2004), and then searching the population of potential behaviors for a behavior with that fitness. If none was found, then another fitness value was drawn at random, and so on, until a parent behavior with the sampled fitness value was found. Parents were drawn from the population with replacement, but with the proviso that each parent pair consisted of two distinct behaviors.

If the emitted behavior was not reinforced, then fitness values were not calculated and parents were chosen randomly from the population with replacement, but with the proviso that each parent pair consisted of two distinct behaviors.

Once 100 pairs of parent behaviors were chosen for mating, a child genotype was constructed from each parent pair's genotypes by taking a bit at each of the 10 bit-string locations either from the father or the mother behavior with equal probability and then inserting it into the child's genotype at that location. In other words, a coin toss determined whether the first bit of the child's genotype was the same as the first bit of the father's genotype or the first bit of the mother's genotype, and so on for the remaining nine locations in the bit string. This is a form of multi-point crossover recombination. The 100 child behaviors so obtained then constituted the population of potential behaviors at the next time tick.

After the new population was constructed, it was subjected to a small amount of mutation by flipping a random bit in a percentage of randomly chosen behaviors.

The foregoing describes the genetic algorithm that was used in this experiment in its entirety. It entailed no other features, properties, equations, distributions, or parameters.

Phase 1. In the canonical experiment, virtual organisms worked on independent concurrent random-interval (RI) random-interval schedules that arranged all combinations of five component RI schedule pairs and five reinforcer magnitude (parental selection function mean) pairs. The RI schedule pairs were RI 20 RI 120, RI 45 RI 95, RI 70 RI 70, RI 95 RI 45, and RI 120 RI 20, where the schedule values are expressed in ticks of time, and the reinforcer magnitude pairs were 75-25, 62-38, 50-50, 38-62, and 25-75, recalling that smaller values of the parental selection function mean represent larger reinforcer magnitudes. Intervals for the RI schedules were drawn at random from exponential distributions with means equal to the RI value. The mutation rate in all conditions was 10%, which means that 10% of the behaviors in each newly formed population became mutants. Responding on each of the 25 combinations of RI schedule pair and reinforcer magnitude pair continued for 20,500 generations, or ticks of time, and each 20,500-generation run was repeated 10 times. This is analogous to running 10 virtual organisms in the 25 conditions of the experiment, which produced a total of about 5 million responses.

Phase 2. The purpose of Phase 2 was to study the effect of a range of mutation rates on the behavior of the virtual organisms. As in Phase 1, virtual organisms animated by the evolutionary theory worked on independent concurrent RI RI schedules that arranged all combinations of five component RI schedule pairs and five reinforcer magnitude pairs at each of 12 mutation rates. The RI schedule pairs and reinforcer magnitude pairs were the same as in Phase 1. The 12 mutation rates were 3%, 5%, 7.5%, 8%, 12%, 14%, 16%, 18%, 20%, 30%, 40%, and 50%. Recall that the mutation rate specifies the percentage of behaviors in each newly constituted population that become mutants. For the 3% and 5% mutation rates, responding on each of the 25 combinations of RI schedule pair and rein-

forcer magnitude pair continued for 100,500 generations, and each 100,500-generation run was repeated 10 times. For all other mutation rates, responding in the 25 conditions continued for 20,500 generations, and each 20,500-generation run was repeated 10 times. About 100 million responses were generated in this phase of the experiment.

Phase 3. The ranges of RI schedules and reinforcer magnitudes studied in Phases 1 and 2 were comparable to those often studied in experiments with live organisms. The purpose of Phase 3 was to explore extended ranges of these independent variables. In the previous phases, the maximum ratios of scheduled reinforcement rates and reinforcer magnitudes were 6/1 and 3/1 respectively. In Phase 3, these maximum ratios were doubled. As in Phase 1, all combinations of five component RI schedule pairs and five reinforcer magnitude pairs were arranged. The RI schedule pairs were RI 15 RI 180, RI 56 RI 139, RI 98 RI 98, RI 139 RI 56, and RI 180 RI 15. The reinforcer magnitude pairs were 15-90, 34-71, 52-52, 71-34, and 90-15. The mutation rate in all conditions was 10%. Responding in each of the 25 conditions continued for 20,500 generations and each condition was repeated 10 times, generating a total of about 5 million responses.

RESULTS

Responses and reinforcements in the target response classes were collected in successive 500-generation (or tick) blocks. Responding on the component schedules stabilized within the first 500-generation block. The first block from each repetition of a condition was discarded, and the response and reinforcement frequencies in the remaining blocks were averaged. This yielded average response and reinforcement rates (per 500-tick block) for the two target responses in each repetition. All analyses were conducted on these average rates.

Phase 1

Equation 2, in logarithmic transformation,

$$\log\left(\frac{B_1}{B_2}\right) = a_r \log\left(\frac{r_1}{r_2}\right) + a_m \log\left(\frac{m_1}{m_2}\right) + \log b, \quad (4)$$

was fitted by the method of least squares to the data from each repetition of the 25 combinations of component RI schedules and component reinforcer magnitudes. This equation describes a plane in three-dimensional space where the two horizontal axes are the logs of the reinforcement rate and reinforcer magnitude ratios, and the vertical axis is the log of the response rate ratio. Because smaller parental selection function means correspond to larger reinforcer magnitudes, the reciprocals of the parental selection function means were used to represent reinforcer magnitudes in Equation 4.

Parameter estimates from the fits, along with the proportion of variance they accounted for (*pVAF*) are listed in Table 1 for each repetition of the experiment. As shown in the last column of Table 1, the equation accounted for essentially all the variance in the virtual organisms' behavior. When response and reinforcement rates were averaged over fewer than 20,000 500-generation blocks, the *pVAF*s were smaller and more variable, and hence more like values that would be obtained from an experiment with live organisms.

As shown in Table 1, the exponent on the reinforcement rate ratio varied around a mean of 0.83 (range: 0.79 - 0.88), and the exponent on the reinforcer magnitude ratio varied around a mean of 0.68 (range: 0.62 - 0.74). The ranges of these exponents increased when response and reinforcement rates were averaged over fewer than 20,000 500-generation blocks. Notice that for every repetition listed in Table 1, the rate exponent was larger than the magnitude exponent. On average the rate exponent was 23% larger than the magnitude exponent. Estimates of the bias parameter, *b*, in Equation 4 are also listed in Table 1, and varied around a mean of 1.01, as would be expected in an environment that entailed no asymmetries other than the experimentally arranged reinforcement property asymmetries.

Data from four of the repetitions listed in Table 1 are plotted in Figure 2, along with a plot of the best fitting plane described by Equation 4. The logs of the reinforcement rate ratios are located on the right horizontal axes of the plots, the logs of the reinforcer magnitude ratios are on the left horizontal axes, and the logs of the response rate ratios are plotted along the vertical axes. An imag-

inary diagonal extending from the front corner of each fitted (gray) plane to its rear corner illustrates the increasing preference for a component as the magnitude and rate of delivery of reinforcement in that component both become more favorable. Also shown in Figure 2 are planes of perfect matching (white), where the three parameters in Equation 4 equal 1. Comparison of the lines of intersection of the gray and white planes with the left rear wall of each plot shows the undermatching that occurred with respect to the reinforcement rate ratio. Comparison of the analogous intersections with the right rear wall shows the undermatching that occurred with respect to the reinforcer magnitude ratio. As is clear from Figure 2, and from the results listed in Table 1, the Equation 4 planes provided an excellent description of the virtual organisms' behavior.

Residuals from the fits of Equation 4 were examined by plotting the standardized residuals (*z*-scores) against the predicted logarithms of the response rate ratios, pooled across the 10 repetitions of the experiment. Visual inspection of this plot, which is shown in Figure 3, revealed no clear trends in the residuals, and statistical analysis (described later) confirmed the absence of significant trends.

Changeovers were examined by plotting the average number of changeovers per 500-generation block as a joint function of the logarithm of the reinforcement rate ratio and the logarithm of the reinforcer magnitude ratio. Changeovers were first averaged over the 40 500-generation blocks in each repetition of

Table 1.

Parameters from and proportion of variance accounted for (*pVAF*) by fits of Equation 4 to ten repetitions of the canonical experiment.

Repetition	<i>a_r</i>	<i>a_m</i>	<i>b</i>	<i>pVAF</i>
R1	0.80	0.66	0.98	0.99
R2	0.79	0.62	1.02	0.99
R3	0.88	0.68	1.01	1.00
R4	0.84	0.71	1.04	0.99
R5	0.87	0.62	1.02	0.99
R6	0.84	0.74	1.01	0.99
R7	0.81	0.69	0.98	1.00
R8	0.84	0.68	0.98	0.99
R9	0.85	0.72	1.02	0.99
R10	0.83	0.64	1.02	0.99
Mean	0.83	0.68	1.01	0.99

Fig. 2. Equation 4 (gray planes) fitted to four repetitions of the canonical experiment (mag = magnitude in the left horizontal axis title). Black data points fall on or above the plane; gray data points fall below the plane. The exponents, bias parameter, and proportion of variance accounted for by the equation ($pVAF$) are given for each fit. White planes show perfect bivariate matching where all parameters in Equation 4 equal 1.

the experiment, and then were averaged over the 10 repetitions. The average changeover rates are plotted in Figure 4, along with the best least-squares fit of the general quadric surface,

$$f(x, y) = Ax^2 + 2Bxy + Cy^2 + 2Dx + 2Ey + G, \quad (5)$$

where x and y are the logs of the reinforcement rate and reinforcer magnitude ratios.

This surface can assume a variety of specific forms depending on the values of its six parameters, which in turn are determined by the data being fitted. For the average changeovers in this experiment, the surface became an elliptic paraboloid, which is a quadric surface with vertical cross sections that are parabolas, and horizontal cross sections that are ellipses. The elliptic paraboloid shown in the figure accounted for 98% of the variance in the average changeover rates. Gray circles in

Fig. 3. Standardized residuals from fits of Equation 4 to data generated by the evolutionary theory. Residuals are plotted against the logarithms of the response rate ratios predicted by the best fitting Equation 4, and are pooled across the 10 repetitions of the canonical experiment.

the plots are projections of the data points onto the rear walls of the 3-D space.

The data and plot in Figure 4 indicate that, in general, as reinforcement rates became more similar in the two components of the concurrent schedules, and/or reinforcer magnitudes became more similar (i.e., as the rate and magnitude coordinates moved toward the center of the plot's floor), the changeover rate increased. Conversely, as these properties of

reinforcement became dissimilar, that is, as they came to favor one or the other component (with rate and magnitude coordinates moving away from the center of the plot's floor), the changeover rate decreased. In addition, the major (longer) axis of the elliptical cross section of the fitted surface was approximately parallel to the log magnitude ratio axis, which means that vertical parabolic cross sections parallel to this axis were wider than those parallel to the log rate ratio axis. This property of the fitted elliptic paraboloid indicates that the changeover rate decreased more rapidly as the reinforcement rate ratios became dissimilar, than as the reinforcer magnitude ratios became dissimilar. Hence, just as behavior allocation was less sensitive to the reinforcer magnitude ratio than to the reinforcement rate ratio, as indicated by the lower exponent on the former, so too the changeover rate was less sensitive to differences in reinforcer magnitude than to differences in reinforcement rate in the schedule components. This was not an artifact of the range of reinforcer magnitude ratios used in the experiment. The range of the log magnitude ratios was smaller than the range of the log rate ratios, as is evident from the horizontal extent of the projected data points shown in Figure 4. The finding that the fitted surface was nevertheless wider along the

Fig. 4. The left panel shows a general quadric surface fitted to the rate of changeovers (COs) averaged over repetitions as a joint function of the logarithm of the reinforcement rate ratio (right horizontal axis) and the logarithm of the reinforcer magnitude ratio (left horizontal axis; mag = magnitude). The surface is an elliptic paraboloid with its major (longer) horizontal axis approximately parallel to the log magnitude-ratio axis. The gray circles are projections of the data points onto the rear walls of the 3-D space. The right panel shows the same plot rotated 90° counterclockwise about the vertical axis of the coordinate system.

log magnitude ratio axis indicates that the changeover rate changed more slowly along that axis than along the log rate ratio axis, even given the smaller range of log magnitude ratios that was studied.

Phase 2

Equation 4 was fitted to the data from individual repetitions of the 25 combinations of component RI schedules and component reinforcer magnitudes at each mutation rate. The estimates of the parameters in the equation and the proportion of variance accounted for were then averaged over the 10 repetitions. These averages are shown in Table 2; averages for the 10% mutation rate were taken from Table 1. In all cases, the standard errors of the means of the average parameters were 0.01 or less.

The $pVAF$ s in Table 2 show that Equation 4 described the response allocations of the virtual organisms well at all mutation rates, although the $pVAF$ s were slightly lower at the two highest mutation rates. Estimates of b across mutation rates varied around 1.0, as would be expected in concurrent schedules that were symmetric other than in the experimental variables.

The estimates of a_r and a_m in Table 2 (± 1 standard error) are plotted in Figure 5 as a function of mutation rate. As is evident from the figure, both parameters varied with mutation rate. The rate exponent, a_r , first increased, reaching a maximum of about 0.83 in the mutation rate range of 7.5% to 12%,

and then decreased as mutation rate increased further. Similar results were reported in earlier tests of the evolutionary theory on concurrent schedules (McDowell et al., 2008; McDowell & Popa, 2010). In contrast, a_m decreased roughly continuously as mutation rate increased. Except at the lowest and highest mutation rates, a_r was always larger than a_m .

Residuals from the fits of Equation 4 were examined as in Phase 1, pooled across repetitions at each mutation rate. At least small cubic polynomial trends were present in all residual plots. Analyses of variance using a Bonferroni-corrected cumulative alpha of 0.05 detected significant cubic polynomial trends at the two lowest mutation rates and at mutation rates greater than 14%. Significant cubic polynomial trends were not found at the five mutation rates in the moderate range of 7.5% through 14%.

Rates of changeover were examined as in Phase 1 by fitting a general quadric surface (Equation 5) to the changeover rates averaged across repetitions. The results of the fits are listed in Table 3; results for the 10% mutation rate are from Phase 1. In all cases, the quadric surface became an elliptic paraboloid which, as shown by the $pVAF$ s in Table 3, described the changeover rates well (median $pVAF$ = 0.97). The vertex of the fitted surface, and hence the overall rates of switching, increased with mutation rate. The eccentricity of the elliptic horizontal cross sections, which is the ratio of the distance between the foci of the ellipses and the length of their major (longer) axes, is a measure of the ellipses' elongation. Zero eccentricity corresponds to a circle; as the eccentricity increases toward the limiting value of 1, the ellipses become more elongated. The eccentricities listed in Table 3 show that the elliptic cross sections of the surfaces tended to become more elongated as mutation rate increased, although there was considerable variability in this trend, especially at low mutation rates. This increasing trend suggests that the smaller sensitivity of changeover rate to magnitude discrepancies, than to rate discrepancies found in Phase 1, became more pronounced as mutation rate increased. However, the increasing eccentricity was complicated by the fact that the surface also tended to rotate counterclockwise about its vertical axis as mutation rate increased, which is shown by the degrees of rotation (counter-

Table 2.

Parameters from and proportion of variance accounted for ($pVAF$) by fits of Equation 4, averaged over 10 repetitions at each mutation rate.

Mutation rate (%)	a_r	a_m	b	$pVAF$
3	0.75	0.77	1.02	1.00
5	0.81	0.73	1.01	1.00
7.5	0.83	0.71	0.99	0.99
8	0.85	0.71	0.98	0.99
10	0.83	0.68	1.01	0.99
12	0.83	0.64	1.00	0.99
14	0.78	0.64	0.99	0.99
16	0.76	0.60	0.99	0.99
18	0.74	0.62	1.00	0.99
20	0.69	0.57	1.00	0.99
30	0.51	0.46	1.00	0.98
40	0.38	0.36	1.00	0.97
50	0.29	0.29	1.00	0.95

Fig. 5. Values of a_r (filled circles) and a_m (unfilled circles) from least-squares fits of Equation 4 plotted as a function of mutation rate. Error bars represent ± 1 standard error of the mean.

clockwise from the x -axis) listed in Table 3. As a result of this rotation, the parabolic vertical cross sections aligned with the log rate ratio and log magnitude ratio axes were no longer approximately coincident with the major and minor axes of the surfaces' elliptical cross sections. Because the axis-aligned cross sections reflect the sensitivity of changeover rate to reinforcement rate and reinforcer magnitude discrepancies in the schedule components, it is necessary to compare the lengths of the latus recta of these specific cross sections directly. The latus rectum of a parabola, which is a measure of its width, is the chord that passes through its focus and is perpendicular

to its axis of symmetry. The ratio of the lengths of the latus recta, listed in the last column of Table 3, tended to approach 1 as mutation rate increased, which indicates that the axis-aligned parabolic cross sections tended to become more, not less, similar at higher mutation rates. Hence, in spite of the increasing elongation of the surfaces as mutation rate increased, their counterclockwise rotation caused the axis-aligned parabolic cross sections to become more similar. This means that as mutation rate increased, the sensitivity of changeover rate to discrepancies in reinforcer magnitude became more, not less, similar to the sensitivity to discrepancies in reinforcement rate.

Phase 3

Equation 4 was fitted to the data from each repetition of the 25 combinations of component RI schedules and reinforcer magnitudes in this phase of the experiment. Parameter estimates from the fits, along with the proportion of variance they accounted for, are listed in Table 4. Again, the fits of Equation 4 were excellent, and the ranges and means of the parameter estimates were comparable to those obtained in Phase 1 (Table 1). In addition, the exponent on the reinforcement rate ratio was larger than the exponent on the magnitude ratio in every repetition of the experiment. Evidently, the behavior of the virtual organisms over the extended ranges of scheduled reinforcement rates and reinforcer magni-

Table 3.

Proportion of variance accounted for ($pVAF$) and properties of quadric surfaces fitted to the average changeover rates at each mutation rate. Eccentricity refers to the elliptic horizontal cross sections of the surface. The ratio of latus recta lengths was calculated for the parabolic vertical cross sections aligned with the log rate ratio and log magnitude ratio axes.

Mutation rate (%)	$pVAF$	Vertex (COs per 500 generations)	Eccentricity	Rotation ($^{\circ}$ counter clockwise)	Ratio of latus recta
3	0.96	1.1	0.57	-149.3	1.1
5	0.97	2.7	0.83	-8.5	3.0
7.5	0.96	4.8	0.75	3.1	2.3
8	0.98	5.1	0.93	4.8	6.6
10	0.98	6.8	0.63	13.0	1.6
12	0.98	8.2	0.65	22.2	1.5
14	0.99	9.3	0.83	9.9	2.8
16	0.98	10.2	0.85	11.5	2.8
18	0.97	11.2	0.78	17.1	1.8
20	0.99	12.1	0.76	34.3	1.3
30	0.97	14.8	0.80	41.4	1.3
40	0.92	16.4	0.86	26.3	1.7
50	0.93	17.4	0.87	75.0	1.2

tudes was similar to their behavior over the ranges studied in Phase 1.

As in Phase 1, the residuals from the fits of Equation 4 were examined by plotting the standardized residuals against the predicted logarithms of the response rate ratios, pooled across the 10 repetitions of the experiment. The plot was similar to that shown in Figure 3 and showed no significant trends.

Changeovers were examined as in Phase 1 by fitting the general quadric surface described by Equation 5 to the average changeover rates. The fitted surface became an elliptic paraboloid, which accounted for 96 % of the variance in the average changeover rates. This was slightly less than the percentage of variance accounted for by the fit to the Phase 1 data. In addition, the vertex of the surface was slightly higher (7.3 vs. 6.8 changeovers per 500 generations), its eccentricity (elongation) was larger (0.83 vs. 0.63), and its counterclockwise rotation was smaller (3 vs. 13 degrees), than the surface fitted in Phase 1. The greater eccentricity and smaller degree of rotation caused the ratio of latus recta of the axis-aligned parabolic vertical cross sections to be much larger for this surface than for the Phase 1 surface (3.4 vs. 1.6). This means that the different (i.e., reduced) sensitivity of switching to changes in the reinforcer magnitude ratio as compared to changes in the reinforcement rate ratio, was more pronounced over the extended ranges of scheduled reinforcement rates and reinforcer magnitudes than over the ranges studied in Phase 1.

DISCUSSION

The results of this experiment show that in the moderate mutation rate range of 7.5% through 14%, bivariate matching (Equation 2), without residual trends, is an emergent property of the evolutionary theory of behavior dynamics. The average a_r in this range was 0.83, which is comparable to values typically found in concurrent schedule experiments with live organisms. The average a_m in this range of mutation rates was 0.68, which compares favorably to the average value of 0.65 reported by Cording et al. (2011) in their review of data from experiments with live organisms. Evidently, these specific exponent values are also emergent properties of the

Table 4.

Parameters from and proportion of variance accounted for ($pVAF$) by fits of Equation 4 to the extended range data from Phase 3.

Repetition	a_r	a_m	b	$pVAF$
R1	0.84	0.62	0.99	0.99
R2	0.78	0.61	1.01	0.99
R3	0.82	0.64	1.02	0.99
R4	0.83	0.63	1.01	0.99
R5	0.83	0.65	1.02	0.99
R6	0.81	0.62	1.01	0.99
R7	0.83	0.63	0.98	0.99
R8	0.83	0.63	0.93	0.99
R9	0.85	0.65	1.02	0.99
R10	0.83	0.68	0.97	0.99
Mean	0.83	0.64	1.00	0.99

theory. It may be worth emphasizing that the results reported here were obtained without materially changing the evolutionary algorithm that produced Equation 1 as an emergent property in an experiment where only the reinforcement rate ratios were varied (McDowell et al., 2008).

As discussed in earlier articles (McDowell, 2004; McDowell & Caron, 2007; McDowell et al., 2008; McDowell & Popa, 2010), specific steady-state equation forms or parameter values were not, and probably could not be, built into the evolutionary theory. Hence, the quantitative form and properties of the steady-state behavior the theory generates can only be the result of the low level Darwinian processes of variation (recombination and mutation) and selection. Just as these processes produce in organic evolution a variety of species with an array of body forms, organ systems, color patterns, and other phenotypic details, they evidently also are able to produce an array of orderly and quantitatively detailed behavioral phenomena, which are indistinguishable from phenomena observed in laboratory experiments with live organisms, and from behavioral phenomena observed in natural environments (McDowell & Caron, 2010a, 2010b). The results of the present experiment, together with results from earlier tests of the evolutionary theory (McDowell, 2004; McDowell & Caron, 2007; Kulube-kova & McDowell, 2008; McDowell et al., 2008; McDowell & Popa, 2010; Popa & McDowell, 2010) are consistent with the assertion that the world of behavior that we observe and measure is created by evolutionary dynamics. These dynamics are formally represented in the

evolutionary theory, and may be implemented in the nervous systems of biological organisms by specific neural dynamics that entail the same processes of variation and selection (McDowell, 2010).

The switching behavior of the virtual organisms in this experiment was similar to that observed by McDowell *et al.* (2008) in their test of the evolutionary theory using concurrent schedules that varied only reinforcement rates. They found that switching was greatest when reinforcement rates were roughly equal in the two components, and decreased as the rates became discrepant. This qualitative result has also been reported in experiments with live organisms (Alsop & Elliffe, 1988; Baum, 1973, 1974a, 1975, 1976; Brownstein & Pliskoff, 1968; Herrnstein, 1961; Heyman, 1979). Changeover patterns have been studied quantitatively in a variety of ways (Baum, 1975, 1976; Hunter & Davison, 1978), but no consensus has emerged on how best to analyze them. McDowell *et al.* found that a quadratic polynomial described the pattern of changeovers exhibited by their virtual organisms well, which is why a general quadric surface was chosen to fit the changeovers in the present experiment. Alsop & Elliffe (1988) presented plots of changeovers for their pigeons (their Figure 4) that appear to be consistent with a parabolic form. Evidently, the evolutionary theory requires these changeover profiles to be parabolic for concurrent schedules that vary only reinforcement rates in the schedule components, and it requires a general quadric surface to become an elliptic paraboloid when fitted to changeover rates in concurrent schedules that vary both the rate and magnitude of reinforcement. The properties of these surfaces (listed in Table 3) varied as a function of mutation rate but, except for the vertices, they did so with considerable noise. Hence, further experiments will be required to understand exactly how these properties are related to mutation rate.

In Phase 2, fits of Equation 4 to data generated at low and high mutation rates showed small but statistically significant residual trends, although the residual mean squares were extremely small at all mutation rates (median = 0.0025, or 1/4 of 1 percent of a log unit). Only fits to data generated at a range of moderate mutation rates from 7.5% through 14% produced residuals that did not

show significant trends. In previous tests of the theory, a roughly comparable range of moderate mutation rates was identified as generating behavior that was similar to the behavior of typical live organisms. McDowell *et al.* (2008) argued that mutation rate may correspond to a fixed property of the nervous system that causes behavior to fall along a continuum from perseverative (lower mutation rates) to impulsive (higher mutation rates), with a middle range of mutation rates representing an adaptive balance between perseveration and impulsivity. This is reminiscent of Pavlov's (1960/1927, Lecture XVII, pp. 284ff.) identification of predispositions in his experimental animals to either inhibition or excitation. It follows that if the typical laboratory subject can be considered to show an adaptive balance between these tendencies, then the evolutionary theory requires its behavior on concurrent schedules, where both the rate and magnitude of reinforcement are varied, to be accurately described by Equation 4 without significant residual trends. However, fits of Equation 4 to data from subjects that exhibit generally perseverative or impulsive behavior would be expected to show small but significant residual trends, while still accounting for substantial proportions of variance.

The extended ranges of scheduled reinforcement rates and reinforcer magnitudes studied in Phase 3 magnified the difference between the effect of reinforcement rate ratio changes and magnitude ratio changes on the rate of switching between components. But otherwise, the extended ranges did not materially affect the virtual organisms' behavior. These ranges were comparable to ranges that have been used in much concurrent schedule research with live organisms, but they did not include extreme values, which have been arranged in some experiments. For example, Davison and Jones (1995) arranged scheduled reinforcement rate ratios as large as 160/1 and Baum, Schwendiman, and Bell (1999) arranged ratios as large as 532/1. Experiments with extreme values have typically been used to test theories such as Davison and Jenkins' (1985) contingency discriminability theory, and Baum's (Baum *et al.*, 1999; Aparicio & Baum, 2006) fix-and-sample account, both of which require large ranges of reinforcement rate ratios to test their predictions. Interestingly, Baum *et al.* (1999) and Davison & Jones

(1995) found that behavior did not conform exactly to power function matching (Equation 1) over the extreme ranges they studied, although the type of deviation from the equation differed in the two experiments. Clearly, tests of the evolutionary theory over extreme reinforcement rate ratio ranges, and also extreme magnitude ratio ranges, is an important project for future research.

Finally, and in a more general vein, it is interesting to note that the evolutionary theory of behavior dynamics is an instance of complexity theory, which is a modern alternative to traditional differential-equation based theory in science (McDowell & Popa, 2009; NOVA, July 2007). Differential-equation based theories typically simplify phenomena by extracting essential properties from them, measurements of which are then entered into equations that constitute the theory. In contrast, complexity theories take the perspective that complex natural phenomena may be caused by sets of relatively simple low level processes acting together over time. According to this perspective, to understand the phenomena it is necessary to discover and understand the simple processes that give rise to them. Holland (1992, 1995, 1998) and, more recently, Mitchell (2009), have discussed complexity theory in general terms. Bentley (2002) has described complexity theories in the biological sciences, and Miller and Page (2007) have discussed applications of complexity theory to social systems. Complexity theories have also appeared before in the behavior-analytic literature. The earliest was Shimp's (1966) theory of momentary maximizing. More recently, Catania (2005) proposed a complexity theory based on Skinner's idea of the operant reserve. Both of these theories have encountered difficulties (Berg & McDowell, 2011; Heyman, 1979; Nevin, 1979). The evolutionary theory of behavior dynamics so far has resisted disconfirmation, although there are many additional environments in which it has yet to be tested.

REFERENCES

- Alsop, B., & Elliffe, D. (1988). Concurrent-schedule performance: Effects of relative and overall reinforcer rate. *Journal of the Experimental Analysis of Behavior*, 49, 21–36.
- Aparicio, C. F., & Baum, W. M. (2006). Fix and sample with rats in the dynamics of choice. *Journal of the Experimental Analysis of Behavior*, 86, 43–63.
- Baum, W. M. (1973). Time allocation and negative reinforcement. *Journal of the Experimental Analysis of Behavior*, 20, 313–322.
- Baum, W. M. (1974a). Chained concurrent schedules: Reinforcement as situation transition. *Journal of the Experimental Analysis of Behavior*, 22, 91–101.
- Baum, W. M. (1974b). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior*, 22, 231–242.
- Baum, W. M. (1975). Time allocation in human vigilance. *Journal of the Experimental Analysis of Behavior*, 23, 45–53.
- Baum, W. M. (1976). Time-based and count-based measurement of preference. *Journal of the Experimental Analysis of Behavior*, 26, 27–35.
- Baum, W. M. (1979). Matching, undermatching, and overmatching in studies of choice. *Journal of the Experimental Analysis of Behavior*, 32, 269–281.
- Baum, W. M., & Rachlin, H. C. (1969). Choice as time allocation. *Journal of the Experimental Analysis of Behavior*, 12, 861–874.
- Baum, W. M., Schwendiman, J. W., & Bell, K. E. (1999). Choice, contingency discrimination, and foraging theory. *Journal of the Experimental Analysis of Behavior*, 71, 355–373.
- Bentley, P. J. (2002). *Digital biology*. New York: Simon & Schuster.
- Berg, J. P., & McDowell, J. J. (2011). Quantitative, steady-state properties of Catania's computational model of the operant reserve. *Behavioural Processes*, 87, 71–83.
- Brownstein, A. J., & Pliskoff, S. S. (1968). Some effects of relative reinforcement rate and changeover delay in response-independent concurrent schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 11, 683–688.
- Catania, A. C. (2005). The operant reserve: A computer simulation in (accelerated) real time. *Behavioural Processes*, 69, 257–278.
- Cording, J. R., McLean, A. P., & Grace, R. C. (2011). Testing the linearity and independence assumptions of the generalized matching law for reinforcer magnitude: A residual meta-analysis. *Behavioural Processes*, 87, 64–70.
- Davison, M., & Baum, W. M. (2003). Every reinforcer counts: Reinforcer magnitude and local preference. *Journal of the Experimental Analysis of Behavior*, 80, 95–129.
- Davison, M., & Jenkins, P. E. (1985). Stimulus discriminability, contingency discriminability, and schedule performance. *Animal Learning & Behavior*, 13, 77–84.
- Davison, M. D., & Jones, B. M. (1995). A quantitative analysis of extreme choice. *Journal of the Experimental Analysis of Behavior*, 64, 147–162.
- Davison, M., & McCarthy, D. (1988). *The matching law: A research review*. Hillsdale, N.J.: Erlbaum.
- Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, 4, 267–272.
- Heyman, G. M. (1979). A Markov model description of changeover probabilities on concurrent variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, 31, 41–51.

- Holland, J. H. (1992). *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence*, 2nd Ed. Cambridge, MA: MIT Press.
- Holland, J. H. (1995). *Hidden order: How adaptation builds complexity*. Cambridge, MA: Perseus.
- Holland, J. H. (1998). *Emergence: From chaos to order*. Cambridge, MA: Perseus.
- Hunter, I. W., & Davison, M. C. (1978). Response rate and changeover performance on concurrent variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, 29, 535–556.
- Killeen, P. R. (1972). The matching law. *Journal of the Experimental Analysis of Behavior*, 17, 489–495.
- Kulubekova, S., & McDowell, J. J. (2008). A computational model of selection by consequences: Log survivor plots. *Behavioural Processes*, 78, 291–296.
- Landon, J., Davison, M., & Elliffe, D. (2003). Concurrent schedules: Reinforcement magnitude effects. *Journal of the Experimental Analysis of Behavior*, 79, 351–365.
- Maguire, D. R., Hughes, C. E., & Pitts, R. C. (2007). Rapid acquisition of preference in concurrent schedules: Effects of reinforcement amount. *Behavioural Processes*, 75, 213–219.
- McDowell, J. J. (1989). Two modern developments in matching theory. *The Behavior Analyst*, 12, 153–166.
- McDowell, J. J. (2004). A computational model of selection by consequences. *Journal of the Experimental Analysis of Behavior*, 81, 297–317.
- McDowell, J. J. (2010). Behavioral and neural Darwinsim: Selectionist function and mechanism in adaptive behavior dynamics. *Behavioural Processes*, 84, 358–365.
- McDowell, J. J., & Caron, M. L. (2007). Undermatching is an emergent property of selection by consequences. *Behavioural Processes*, 75, 97–106.
- McDowell, J. J., & Caron, M. L. (2010a). Matching in an undisturbed natural human environment. *Journal of the Experimental Analysis of Behavior*, 93, 413–431.
- McDowell, J. J., & Caron, M. L. (2010b). Bias and undermatching in delinquent boys' verbal behavior as a function of their level of deviance. *Journal of the Experimental Analysis of Behavior*, 93, 469–481.
- McDowell, J. J., Caron, M. L., Kulubekova, S., & Berg, J. P. (2008). A computational theory of selection by consequences applied to concurrent schedules. *Journal of the Experimental Analysis of Behavior*, 90, 387–403.
- McDowell, J. J., & Popa, A. (2009). Beyond continuous mathematics and traditional scientific analysis: Understanding and mining Wolfram's *A New Kind of Science*. *Behavioural Processes*, 81, 343–352.
- McDowell, J. J., & Popa, A. (2010). Toward a mechanics of adaptive behavior: Evolutionary dynamics and matching theory statics. *Journal of the Experimental Analysis of Behavior*, 94, 241–260.
- Miller, J. H., & Page, S. E. (2007). *Complex adaptive systems: An introduction to computational models of social life*. Princeton, N.J.: Princeton University Press.
- Mitchell, M. (2009). *Complexity: A guided tour*. New York, N.Y.: Oxford University Press.
- Myers, D., & Myers, L. (1977). Undermatching: A reappraisal of performance on concurrent variable-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 27, 203–214.
- Nevin, J. A. (1979). Overall matching versus momentary maximizing: Nevin (1969) revisited. *Journal of Experimental Psychology: Animal Behavior Processes*, 5, 300–306.
- NOVA (July, 2007). *Emergence*. Public Broadcasting Service (PBS). Retrieved January 17, 2012 from <http://www.pbs.org/wgbh/nova/sciencenow/3410/03.html>.
- Pavlov, I. P. (1960). *Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex* (G. V. Anrep, Tr.). Oxford, England: Dover Publications. (Originally published in 1927).
- Popa, A., & McDowell, J. J. (2010). The effect of Hamming distances in a computational model of selection by consequences. *Behavioural Processes*, 84, 428–434.
- Rachlin, H. (1971). On the tautology of the matching law. *Journal of the Experimental Analysis of Behavior*, 15, 249–251.
- Russell, S., & Norvig, P. (2003). *Artificial intelligence: A modern approach*. Saddle River, NJ: Prentice-Hall.
- Schneider, J. W. (1973). Reinforcer effectiveness as a function of reinforcer rate and magnitude: A comparison of concurrent performances. *Journal of the Experimental Analysis of Behavior*, 20, 461–471.
- Shimp, C. (1966). Probabilistically reinforced choice behavior in pigeons. *Journal of the Experimental Analysis of Behavior*, 9, 443–455.
- Todorov, J. C. (1973). Interaction of frequency and magnitude of reinforcement on concurrent performances. *Journal of the Experimental Analysis of Behavior*, 19, 451–458.
- Todorov, J.C., Hanna, E.S., & Bittencourt de Sa, M.C. (1984). Frequency versus magnitude of reinforcement: New data with a different procedure. *Journal of the Experimental Analysis of Behavior*, 41, 157–167.
- Wearden, J. H., & Burgess, I. S. (1982). Matching since Baum (1979). *Journal of the Experimental Analysis of Behavior*, 38, 339–348.

Received: May 13, 2012

Final Acceptance: July 3, 2012