The copyist model of response emission

Takayuki Tanno · Alan Silberberg

Published online: 7 June 2012
© Psychonomic Society, Inc. 2012

Abstract The variety of different performances maintained by schedules of reinforcement complicates comprehensive model creation. The present account assumes the simpler goal of modeling the performances of only variable reinforcement schedules because they tend to maintain steady response rates over time. The model presented assumes that rate is determined by the mean of interresponse times (time between two responses) between successive reinforcers, averaged so that their contribution to that mean diminishes exponentially with distance from reinforcement. To respond, the model randomly selects an interresponse time from the last 300 of these mean interresponse times, the selection likelihood arranged so that the proportion of session time spent emitting each of these 300 interresponse times is the same. This interresponse time defines the mean of an exponential distribution from which one is randomly chosen for emission. The response rates obtained approximated those found on several variable schedules. Furthermore, the model reproduced three effects: (1) the variable ratio maintaining higher response rates than does the variable interval; (2) the finding for variable schedules that when the reinforcement rate varies from low to high, the response rate function has an ascending and then descending limb; and (3) matching on concurrent schedules. Because these results are due to an algorithm that reproduces reinforced interresponse times,

responding to single and concurrent schedules is viewed as merely copying what was reinforced before.

Keywords Variable ratio · Variable interval · Choice · Matching · Simulation · Interresponse time reinforcement · Copyist model

In a variable-ratio (VR) schedule, the number of responses required to deliver a reinforcer varies between successive reinforcers. The mean of these interreinforcement ratios (the number of responses between successive reinforcers) defines the schedule's value. For example, if the first reinforcer was delivered after, say, an interreinforcement ratio of 50, the second after 10, and the third and final reinforcer after 30, the schedule would be defined as a VR 30 (90 responses/3 reinforcers). In a variable-interval (VI) schedule, a reinforcer is delivered for the first response following completion of an interreinforcement interval (IRI). In the case of the first reinforcer, this interval is timed from the beginning of the experiment. For subsequent reinforcers, it is timed from the completion of the delivery of the prior reinforcer. For example, a VI 30-s schedule that delivered three reinforcers could be created by requiring that (1) 50 s elapse since the session's beginning before a response could deliver a reinforcer and (2) the next two reinforcers be delivered for the first response following passage of, say, 10 and 30 s, respectively (90 s of cumulative IRI time/3 reinforcers).

When these two schedules provide the same rate of reinforcement, the response rate is typically higher to the VR than to the VI (e.g., Baum, 1993; Catania, Matthews, Silverman, & Yohalem, 1977; Ferster & Skinner, 1957, pp. 399–407; Peele, Casey, & Silberberg, 1984). For example, Peele et al. recorded the IRIs generated by pigeons on a VR

T. Tanno (✉)
Department of Psychological Sciences,
Kwansei Gakuin University,
1-1-155 Uegahara,
Nishinomiya-shi, Hyogo 662-8501, Japan
e-mail: tantantan01@gmail.com

A. Silberberg
American University,
Washington, DC, USA

schedule and then used those IRIs to create VI intervals. When response rates in their first experiment were averaged across subjects, they found that rates were 27 % higher on the VR than on the VI even though each schedule provided approximately the same rate of reinforcement.

The finding of a rate difference between VR and VI schedules that provide the same rate of reinforcement has generated considerable interest among operant-oriented researchers (Baum, 1973, 1981, 1993; Bowers, Hill, & Palya, 2008; Catania et al., 1977; Cole, 1994, 1999; Dawson & Dickinson, 1990; Ferster & Skinner, 1957; Morse, 1966; Peele et al., 1984; Reed, Hildebrandt, DeJongh, & Soh, 2003; Reed, Soh, Hildebrandt, DeJongh, & Shek, 2000). Empirical and theoretical reasons drive this interest. In terms of the former, variable schedules may be simpler to explain than periodic schedules like fixed interval (FI) and fixed ratio (FR) because variable schedules require less consideration of time-dependent changes in performance. For example, an FI schedule, which reinforces a response following the passage of a specified interval, may produce two, separable states over its IRI: a period of inaction following reinforcement, followed by high rates of responding to the delivery of the next reinforcer (B. A. Schneider, 1969). Since variable schedules tend not to do this, they are likely to be simpler to explain.

Because of Skinner, the VR–VI response rate difference is also of considerable theoretical interest to operant-oriented researchers. Skinner (1932a, b) claimed that when a rat's propensity to press a lever for a food reinforcer was viewed over time, it corresponded to the temporal patterning of the rat's eating reflex. Moreover, when hunger motivation was altered, both rate of response and food consumed changed in like manner. In his view, these outcomes rationalized the use of response rate as an index of the strength of the eating reflex (Skinner, 1938, 1950). Subsequently, Skinner made little mention of reflex strength and its measurement through response rate. Although our explanation of his comparative silence is admittedly speculative, it may well be because of data he helped generate. Ferster and Skinner (1957) found that VR schedules support higher response rates than do VI schedules when both schedules provide the same rate of reinforcement. If one assumes that the eating reflex does not differ as a function of the schedule that provides it, this between-schedule difference in response rate is incompatible with the idea that response rate measures strength. The obvious explanation for their result is that features of how VR and VI schedules arrange reinforcement shape the different response rates obtained.

While it is apparent that VR and VI schedules that provide the same reinforcement rate shape different response rates, why do they do so? Two kinds of accounts have been proposed. One is a molecular account that attributes this rate difference to between-schedule differences in the relation between the time between two successive responses

(interresponse time, or IRT) and the probability of reinforcement (e.g., Morse, 1966). In particular, on VIs the probability of reinforcement is an increasing and bounded function of IRT duration, while on VRs the probability does not change with the duration of the IRT. If response emission is controlled by this relation, the differential reinforcement of long IRTs on VI, but not VR, should result in longer IRTs on VI than on VR and, in consequence, lower rates to the former schedule, because IRT duration and response rate are inversely related.

The second account is often considered to be molar because it, unlike the IRT reinforcement account offered above, is based on aggregations of IRTs (response rate) and reinforcements over time (reinforcement rate). This model attributes the VR–VI rate difference to differences in the response-rate, reinforcement-rate feedback functions for these two schedules (Baum, 1973, 1981). Specifically, on VR, but not VI, marginal increases in response rates produce higher rates of reinforcement on VR than they do on VI. If animals are sensitive to this feedback function difference, the higher rates seen on VR schedules can be rationalized.

Beginning with Peele et al. (1984), two-link tandem schedules have been used to evaluate whether response emission on variable schedules is primarily controlled at a molecular or a molar level. In a tandem schedule, two or more schedules are presented in sequence in the presence of a single discriminative stimulus. Meeting the response contingency in one schedule is required to access a subsequent schedule. Only the terminal schedule in a schedule sequence provides reinforcement.

To illustrate use of tandem schedules, consider Experiment 1 from Peele et al. (1984). In phase 1 of this experiment, VR and VI schedules that provided the same rate of reinforcement were serially presented, each schedule identified by a different discriminative stimulus. As was expected, they found on this multiple schedule that response rates were consistently higher to VR than to VI. In the next phase, the VI was replaced by a tandem VI differential reinforcement of high rates (DRH) schedule. With this new arrangement, once the IRI defined by the VI had elapsed, reinforcement required emission of an IRT that approximated the ones reinforced on the VR. Peele et al. found that the rate differences apparent in the prior phase now were reduced but not eliminated—a result they interpreted as demonstrating the contribution of molecular factors to the control of rate emission on VI schedules. In their next experiment, they provided a complementary demonstration of this effect by showing that a VR–VI rate difference in a multiple VR VI schedule virtually disappears when reinforcement on a VR requires emission of IRTs similar to those reinforced on VI. They imposed this contingency in the VR multiple-schedule component by replacing that schedule with a tandem VR differential reinforcement of low rates (DRL) schedule that required, after the VR schedule requirement was met, emission of an IRT in the range of those normally found on VIs to produce reinforcement. Peele et al. emphasized that

the IRT-based contingency in the second link of the tandem schedule seemed largely determinative of the response rate obtained even though the rates shaped by its IRT contingency were opposed to those predicted by the molar feedback functions of the VR and VI schedules defining the first link (hence, the name tandem “counterpoise” test).

This conclusion is complemented by several specific schedules that have been created to evaluate whether, and in what ways, response rate is sensitive to molar control. In some of these schedules, the rate of reinforcement was positively correlated with response rate (Cole, 1999; McDowell & Wixted, 1986; Reed et al., 2003; Reed et al., 2000); in others, there was an inverse relation between these variables (Ettinger, Reid, & Staddon, 1987; Jacobs & Hackenberg, 2000; Reed & Schachtman, 1991; Tanno & Sakagami, 2008; Vaughan, 1987; Vaughan & Miller, 1984); and in still others, reinforcer rate was largely independent of response rate (Dawson & Dickinson, 1990; Kuch & Platt, 1976; Tanno & Sakagami, 2008). With few exceptions (Dawson & Dickinson, 1990), these studies fail to show control of rates by molar variables.

While these results are consistent with the idea that, on VR and VI schedules, molecular IRT contingencies primarily control rate, rather than the molar relation between response rate and reinforcement rate, analysis of Peele et al.’s (1984) data and that of subsequent work suggests that this molecular control may be more complex than supposed or, possibly, selectively incomplete. In Peele et al. and several other reports that have used tandem-schedule counterpoise tests, tandem VI DRH consistently falls short of matching the rates supported by a standard VR schedule (Cole, 1994; Reed et al., 2000, Experiment 4; Tanno & Sakagami, 2008).

Why does tandem VR DRL, but not tandem VI DRH, offer a clear-cut endorsement of a molecular account of response rate? Suggestive evidence can be found in Table 1, which is based on Tanno and Sakagami (2008). It compares across three conditions the median duration of IRTs prior to the one reinforced with the reinforced IRT on VR, VI, and tandem VI DRH in their study. Note that, for all subjects, the median duration of

the IRT preceding the reinforced IRT on the tandem VI DRH schedule was longer than those on target VR and VI schedules. If we assume that the backward action of reinforcement extends beyond the reinforced IRT to include its predecessors in a fashion akin to a delay-of-reinforcement gradient, a simple explanation for the intermediacy of tandem VI DRH rates emerges. That is, while the reinforced IRT on this tandem schedule is VR-like in duration, the predecessor IRT is VI-like, and thus response emission might be a blend of VR- and VI-like rates—the result apparently obtained. By this reasoning, rate control is still viewed largely as molecular, but the molecule is enlarged to consist of perhaps the two IRTs most proximal to reinforcement. This idea is supported by prior work that shows that control by reinforcement can extend over successive IRTs and response sequences in choice (Angle, 1970; Catania, 1971; Shimp, 1973).

To summarize, while several reports can be interpreted as showing that response rate is controlled in part by multiple IRTs proximal to reinforcement (e.g., Angle, 1970; Peele et al., 1984; Reed et al., 2000; Shimp, 1973; Wearden & Clark, 1988), rather than the molar control of response-rate, reinforcer-rate correlations, none models how these IRTs might contribute to the VR–VI rate difference. The present article remedies this by providing a new model of behavioral output based on an adaptation to a response-generating algorithm first outlined by Peele et al. To anticipate the results of our first simulation, this new model stands alone in accommodating the intermediate rate effect Tanno and Sakagami (2008) and others have noted on tandem VI DRH. In addition, it predicts Peele et al.’s finding of a match between rates on VI and tandem VR DRL schedules.

The copyist model

In Peele et al.’s (1984) model, a simulated bird responded to a schedule by randomly sampling from an exponential distribution of IRTs with a mean duration of 1.2 s. Once 300

Table 1 Data from Table 1 in Tanno and Sakagami (2008)

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
		Subject Number								M	(M)	
Condition	Schedule	T1	(T1)	T2	(T2)	T3	(T3)	T4	(T4)	IRT	IRT	resp / min
4	Tand VI DRH	0.31	0.88	0.50	1.20	0.48	1.11	0.61	1.09	0.48	1.07	61.71
5	VR	0.41	0.39	0.58	0.60	0.82	0.77	0.78	0.82	0.65	0.64	65.91
6	VI	1.41	0.58	1.42	0.66	1.83	0.98	1.39	0.92	1.51	0.78	52.46

reinforcers occurred, control of response emission shifted to an IRT selection algorithm that sampled from a continuously updated list of the 300 most recently reinforced IRTs. The likelihood that a particular IRT would be selected from this list equaled the reciprocal of its duration divided by the sum of the reciprocals of all 300 reinforced IRTs in the list. Once an IRT was selected, its duration defined the mean value of an exponential distribution of IRTs. The IRT actually emitted was randomly selected from this distribution, with the constraints that (1) IRTs less than 0.25 s were replaced by a 0.25-s IRT and (2) following simulated reinforcement, IRT emission was delayed for 1 s (simulated postreinforcement pause).

In Peele et al.'s (1984) model, each IRT emitted is a copy of a previously reinforced IRT. The harmonic weighting scheme is simply a device to ensure that all copies of previously reinforced IRTs occupy the same proportion of session time (Shimp, 1969, 1974; Shimp, Fremouw, Ingebritsen, & Long, 1994). To illustrate the role this weighting plays, imagine that, in a session, a simulated performance consists of just 2-s and 4-s IRTs. On the basis of the response selection algorithm, the probability of emission of 2-s IRTs is twice that of 4-s IRTs (i.e., the frequency of one reinforced 2-s IRT times the reciprocal of its length [$1 \times \frac{1}{2}$] is twice that of the 4-s IRT [$1 \times \frac{1}{4}$]). Note that by weighting reinforced IRTs by the reciprocals of their lengths, each of the two reinforced IRTs in this example occupies the same proportion of session time. That is, in a 1-h session, 30 min of responding would be allocated to each of the two IRT classes, an outcome that illustrates the role played by harmonically weighting reinforced IRTs. The result is a model that repeats prior reinforced IRTs, with each IRT in the 300-member reinforcement memory occupying 1/300th of session time.

The revised copyist model presented in this report makes only two changes in Peele et al.'s (1984) rendition. First, selected IRTs that were less than 0.25 s were no longer emitted as 0.25-s IRTs. Instead, all IRTs less than the criterion value (IRT-min) were discarded, and a new selection was made from the exponential distribution of IRTs until the one selected exceeded IRT-min. The rationale for this change was that Peele et al.'s selection method naturally led to an overabundance of 0.25-s IRTs. By discarding IRTs less than 0.25 s, our method does not. Second, when IRT memory is updated by reinforcement, it is not the reinforced IRT that is selected for inclusion in reinforcement memory. Instead, it is the exponentially weighted mean of all of the IRTs composing the IRI. It is this change that critically distinguishes this model from Peele et al.'s.

The exponential weighting scheme is defined by the following equation:

$$\text{weight}(i) = \exp(-\lambda d_i), \quad (1)$$

where i signifies each IRT in the current IRI, d_i is the time between onset of the i th IRT and reinforcement, and λ is the

exponential parameter. To illustrate this equation's use, suppose that the sequence of IRTs defining an IRI was 0.8, 0.4, 2.0, and 0.4 s. In Peele et al.'s (1984) model, only the reinforced IRT, 0.4 s, would be updated to memory. However, in this adaptation of their work with $\lambda = 1$, the time between the onset of each IRT and reinforcement is calculated (3.6, 2.8, 2.4, and 0.4 s, respectively). On the basis of Equation 1, the weighting of each IRT would be 0.03, 0.06, 0.09, and 0.67, respectively. The mean of these weighted IRTs, 0.58 s, would be updated to reinforcement memory.

Tests of the copyist model

Simulation 1: Evidence of the predictive superiority of the copyist model

The goal of this first simulation was to test the adequacy of the copyist model relative to the one offered by Peele et al. (1984). To this end, we simulated response rates for both models on VR, VI, tandem VI DRH, and tandem VR DRL schedules. Each simulation ended after 5,000 reinforcers and was conducted five times. All data presented are based on the last 1,000 reinforcers of each simulation averaged across the five simulations.

On the VR, each IRT was reinforced with $p = 1/30$, thereby creating a VR 30. The same VR 30 schedule was used in the VR component in the tandem VR DRL schedule. The schedule values of the single VI and the VI component in the tandem VI DRH schedule were equated to the mean IRI of the last 1,000 IRIs on the single VR schedule. The actual VI values for each IRI were defined by random sampling from the exponential distribution. Because of this arrangement, the variable schedules, both here and in subsequent simulations, may be described as *random ratio* and *random interval*, respectively.

In order to define the IRT criterion for DRH reinforcement, the distribution of the last 1,000 reinforced IRTs on VR were split into four classes that were equally sized in terms of membership. The IRT criterion for DRH reinforcement was randomly selected from one of these four classes. Reinforcement occurred if the simulated IRT was in the range of the selected IRT value (also see Tanno & Sakagami, 2008). The same technique was used to determine DRL values in the tandem VR DRL schedule, except that the distribution was based on the last 1,000 reinforced IRTs on the VI schedule.

Results and discussion of simulation 1 Figure 1 maps simulated response rate as a function of the exponential parameter of the copyist model and the predictions of Peele et al.'s (1984) reinforced-IRT account. Each panel presents performances of a *comparison* schedule (tandem VI DRH

Fig. 1 Simulated response rates for the copyist model as a function of the value of λ and the Peele, Casey, and Silberger (1984) reinforced-IRT model. Each panel compares a different comparison schedule with VR and VI target schedules

[top panel] or tandem VR DRL [bottom panel]) and is compared with the rates supported by *target* VR and VI schedules. In those cases where the target schedule is described in the panel as “VR-like,” Peele et al.’s account and ours predict that comparison rates should approximate that of the target VR. Similarly, when the target schedule is described as “VI-like,” these models predict rate approximation between that comparison schedule and the target VI seen in each panel.

Viewed together, the results from tandem VI DRH and tandem VR DRL seem anomalous because, while these schedules are generally complementary in construction, the rates they support are not complementary in outcome. In one case, that of tandem VI DRH, the rates this schedule supports covary with λ . But for tandem VR DRL, rates mimic those seen in their target schedule, VI, regardless of the value of λ . Despite this asymmetry, the results of both tandem schedules can be rationalized by the copyist model. The top panel of Fig. 1 successfully models the intermediate rate maintained by tandem VI DRH when $\lambda = 0.5$. Moreover, if one accepts λ as a constant applicable to both simulations, the copyist model at $\lambda = 0.5$ also correctly matches rates to those generated on a tandem VR DRL schedule (see the bottom panel). Finally, as would be

expected, when exponential decay is rapid (e.g., $\lambda = 5$), our model and Peele et al.’s (1984) make essentially identical predictions. This result suggests that our model’s second difference from Peele et al.’s—that selected IRTs below 0.25 s were discarded rather than emitted as 0.25-s IRTs—played a negligible role in accounting for rate differences between these models.

Table 2 presents for the simulation the average duration of the reinforced IRT and the predecessor IRT for all schedules and data points presented in Fig. 1. The following results obtained: (1) For VR schedules, the reinforced and predecessor IRTs were approximately of equal duration; (2) for VI, all IRTs were of longer duration than for VR, and the IRT producing reinforcement was longer than its predecessor IRT; (3) reinforced and predecessor IRTs on tandem VR DRL approximated the values for these IRT classes on its target schedule, VI; and (4) while the reinforced IRT on tandem VI DRH matched those of the VR target schedule, the preceding IRT was more akin to those seen on the VI schedule than on the VR. These results are consistent with our claims that (1) response rate to all four schedules can be synthesized by aggregating reinforced and predecessor IRTs and (2) the intermediate rates supported by tandem VI DRH at moderate rates of exponential decay (see the top panel of Fig. 1) can be rationalized by assuming that this schedule’s rates are a consequence of blending long predecessor IRTs with short reinforced IRTs.

Peele et al.’s (1984) rate simulations illustrated how an IRT-reinforcement model could accommodate the ordinal relation between VR and VI response rates. However, this success was not matched by the absolute rate predictions for the individual schedules from which this difference was calculated. Generally speaking, the Peele et al. model

Table 2 Simulated durations of the reinforced IRT and its predecessor for the copyist model and Peele, Casey, and Silberger’s (1984) reinforced-IRT account

Measure	λ	VR	VI	Tandem VI DRH (VR-like)	Tandem VR DRL (VI-like)
Mean reinforced IRT (s)	0.1	0.8	3.1	0.6	3.2
	0.3	0.7	2.6	0.4	2.8
	0.5	0.5	2.1	0.4	2.3
	1	0.5	1.8	0.3	2.6
	5	0.4	2.0	0.2	2.6
	Peele et al.	0.4	1.8	0.2	2.6
Mean predecessor IRT (s)	0.1	0.8	1.5	2.4	1.0
	0.3	0.6	1.0	1.5	0.9
	0.5	0.6	0.9	1.1	0.7
	1	0.5	0.7	0.9	0.7
	5	0.4	0.6	0.5	0.5
	Peele et al.	0.3	0.6	0.5	0.4

predicted response rates that were substantially above those typically obtained on VR and VI schedules (see Wearden & Clark, 1989). As is shown below, the present adaptation of their model fares better, not only maintaining the proper qualitative predictions of VR rates greater than VI rates, but also predicting their absolute levels.

Simulation 2: Adequacy of the copyist model in predicting response rates on variable schedules

The first simulation made a specific test: Can the copyist model predict the finding that, in counterpoise tests with tandem schedules, tandem VR DRL more closely approximates the rate of a VI schedule than tandem VI DRH does of a VR? As is shown in simulated data in Fig. 1, this outcome can be rationalized by our copyist model, but not by the model from Peele et al. (1984). Moreover, in terms of the IRT analysis presented in Table 2, simulated performances on tandem VI DRH replicate the results reported by Tanno and Sakagami (2008; see Table 1). That is, for both the simulation and their rats, predecessors IRTs were longer than reinforced IRTs only on the latter schedule.

This accomplishment notwithstanding, simulation 1 models performances by a computer, not by real subjects. Moreover, its tests are limited to only four of the schedules used by Peele et al. (1984) to create and test the VR–VI rate difference: VR, VI, tandem VR DRL, and tandem VI DRH. In the present simulation, we expand this list to include every schedule of which we are aware that has been used in VR–VI rate tests. Our rationale is that these schedules share properties with VR and VI beyond variability in reinforcement and often counterpoise a test for molecular and molar control of response emission in fundamentally different ways. In our view, such a collection of schedules constitutes a robust test of the copyist model. To this end, we add five schedules to the four already tested. In all cases, the model's predictions are compared not with the model's simulations of VR and VI rates, but with those generated by the rat, pigeon, and human subjects exposed to this schedule set.

The nine schedules to be evaluated are the four seen in simulation 1 (VR, VI, tandem VR DRL, and tandem VI DRH) plus tandem VR VI, tandem VI VR, regulated probability interval (RPI), VI with linear feedback loop (VI+), and negative feedback (NF). Tandem schedules were defined earlier. For a full account of how these other, new schedules are created, we refer the reader to the referenced articles. However, to describe their properties briefly, RPI schedules have the molar feedback properties of a VI but the IRT-reinforcement characteristics of a VR; VI + operates at a molar level like a VR but reinforces at the molecular level of the IRT like a VI; and the NF schedule used reinforces IRTs in a fashion identical to a VR (a rate-enhancing feature

if rates are controlled molecularly) but lowers the overall reinforcement rate (a rate-decreasing contingency if rates are controlled at a molar level).

Table 3 presents the data set used to evaluate the predictive adequacy of the copyist model. The composition of this table is restricted to studies in which both VR and VI schedules and reinforcement rates of the VI were equated with those produced on a VR schedule, except for Peele et al.'s (1984) Experiment 3. The table lists the data source within each publication and the species and value of the target VR schedule. All rates presented were averaged across subjects at a particular VR value, except those marked by asterisks, which were averaged not only across subjects, but also across VR ratio sizes. All data for comparison schedules are presented in responses/minute.

The simulation procedure was the same as that in simulation 1, except that, to determine the best-fitting value of λ and IRT-min, most simulations of a comparison schedule's response rate were stepped through a factorial combination of these parameters. For IRT-min, the range covered was 0.05 – 1.00 s in 0.05-s steps. For λ , the range and step size were 0.1 – 3.0 and 0.1, respectively. Definition of the best-fitting value was resolved statistically as the combination of IRT-min and λ values that showed minimal deviance from actual response rates. Three exceptions to this arrangement were simulations from Peele et al. (1984), Baum (1993), and Cole (1999), where it was possible to discern IRT-min from their data sets. Additionally, we tested the copyist model without exponential weighting to evaluate the necessity of this scheme.

Results and discussion of simulation 2 Just how capable is the copyist model in accommodating absolute response rates of real subjects on this diverse set of variable schedules? Table 4 summarizes the predicted response rate from the copyist model, and the upper panel of Fig. 2 shows obtained versus predicted response rate data, each of which is based on Tables 3 and 4. Whether based on one or two free parameters (depending on whether actual or best-fit values defined IRT-min), the copyist model accommodates approximately 90 % of the variance seen in response rates.

The bottom panel of Fig. 2 presents outcomes from the copyist model without exponential weighting. R^2 decreased considerably, an outcome that supports the necessity of an exponential-weighting scheme in explaining VR–VI rate differences.

Simulation 3: Modeling the relation between response rate and reinforcement rate on VR and VI schedules

In our view, the copyist model does a credible job in predicting absolute response rates on several variable schedules. In this simulation, we extend its application to a new

Table 3 Responses/minute obtained in previous studies

Research	Subject	VR value	Schedule								
			VR	VI	VI-DRH	RPI	VR-DRL	VI+	NF	VR-VI	VI-VR
Catania et al.(1977)	Pigeon	25	91	49							
Matthews et al. (1977)	Human	30	210	119							
Peele et al. (1984, E.1)	Pigeon	200	106	86	97						
Peele et al. (1984, E.2)	Pigeon	200	67	48			49				
Peele et al. (1984, E.3)	Pigeon	100								48	76
Dawson & Dickinson (1990, E.2)	Rat	20	70	22		33					
Baum (1993)*	Pigeon	57	151	125							
Cole (1994, E. 1)	Rat	10	80	59	68						
Cole (1999)*	Rat	35	50	26				29			
Reed et al. (2000, E.1)	Rat	30	46	22				18			
Reed et al. (2000, E.2)	Rat	30	64	44				32			
Reed (2001, E.1 & 4)	Human	8	51	35	44						
Reed (2007a, E. 1)	Human	30	137	76				117			
Reed (2007b, E.2, low force)	Rat	30	60	37				63			
Reed (2007b, E.2, high force)	Rat	30	41	6				11			
Tanno & Sakagami (2008)	Rat	30	61	52	62	62	61				
Mean		54	86	54	68	48	49	45	61	48	76

E=experiment. Data sources are Catania, Matthews, Silverman, and Yohalem (1977), Table 1; Matthews, Shimoff, Catania, and Sagvolden (1977), Table 1; Peele, Casey, and Silberberg (1984, E. 1), Table 2; Peele et al. (1984, E. 2), Table 4; Peele et al. (1984, E. 3), Table 5; Dawson and Dickinson (1990, E. 2), Fig. 3; Baum (1993), Fig. 4; Cole (1994, E. 1), Table 1; Cole (1999), Fig. 1, all-rats data; Reed, Soh, Hildebrandt, DeJongh, and Shek, (2000, E. 1), Fig. 1, last two-session block; Reed et al. (2000, E. 2), Fig. 2, last two-session block; Reed (2001, E. 1 and 4), averaged across Figs. 1 and 4, last trial; Reed (2007a, E. 1), Fig. 1, using VI[yVR] condition for VI rate; Reed (2007b, E. 2, low force; high force), Fig. 2, last three sessions, averaged VI+ rates; Tanno & Sakagami (2008), Table 1. Baum (1993) and Cole (1999) use several VR values. Table contains data averaged across those VR values and are indicated by asterisks.

problem: modeling VR and VI response rate changes as a function of changes in the density of reinforcement.

The sole comprehensive data set of these relations is offered by Baum (1993). In his study, 4 pigeons responded on multiple VR VI schedules in which the rate of VI reinforcement was equated with the rate of reinforcement obtained on the VR. He varied the size of the VR schedule from 1 to 512 across conditions. He found two major effects: (1) As reinforcement rates increased from VR 512 through VR 1, response rates to both schedules tended to increase with schedule size up to approximately VR 32 and then to decrease with larger ratios, drawing thereby bitonic VR and VI functions (see Baum, 1993, Fig. 4), and (2) the VR–VI rate differences maintained at moderate reinforcement rates diminished with increasing reinforcement rates until they actually reversed at the two richest VR schedule sizes (VR 1 and VR 2).

To model Baum's (1993, Fig. 4) data, we set as constants the IRT-min (0.05 s) and the mean, across-subject postreinforcement pause data he obtained at each VR size and corresponding yoked VIs (increased from 0.67 to 1.39 s as VR value increases) and then varied λ in the fashion described in simulation 2. Therefore, in this simulation, the

copyist model functioned with a single free parameter. The value of λ that minimized predictive error was 0.5.

Results and discussion of simulation 3 The top panel of Fig. 3 presents the results of this simulation excluding Baum's (1993) data for VR 1 and VR 512. While a rationale could be offered for excluding VR 1 from our analysis on the basis of the fact that such a schedule is usually not considered to be a VR, our rationale differs: Because the copyist model assigns a role to predecessor IRTs in defining rate and because, on a VR 1 schedule, there cannot be predecessor IRTs, since all IRTs are reinforced, a VR 1 schedule falls outside the scope of performances this model was intended to simulate. Performances on VR 512 were also excluded because only one of 4 subjects responded on this schedule. We were in a quandary as to whether this schedule should be viewed as a VR or extinction.

The top panel of Fig. 3 illustrates that the model has some success in accommodating variance, with the variance-accommodated statistic equaling 0.77. More impressive, in our view, is the correspondence between the model's qualitative features and the data from Baum's (1993) Fig. 4. In particular, the simulation and data share the following

Table 4 Predicted responses/min and estimated values of the copyist model

Research	Subject	VR value	Schedule									Estimated parameter	
			VR	VI	VIDRH	RPI	VRDRL	VI+	NF	VR-VI	VI-VR	IRT min	λ
Catania et al. (1977)	Pigeon	25	86	49								0.15	0.7
Matthews et al. (1977)	Human	30	192	123								0.05	1.0
Peele et al. (1984, E.1)	Pigeon	200	127	57	94							0.05*	0.3
Peele et al. (1984, E.2)	Pigeon	200	85	38			51					0.05*	0.1
Peele et al. (1984, E.3)	Pigeon	100								58	82	0.05*	0.2
Dawson & Dickinson (1990, E.2)	Rat	20	62	29		62						0.25	1.0
Baum (1993)	Pigeon	57	157	127								0.05*	0.5
Cole (1994, E. 1)	Rat	10	77	55	68							0.15	0.4
Cole (1999)	Rat	35	50	24				24				0.3*	0.7
Reed et al. (2000, E.1)	Rat	30	44	19				19				0.35	0.7
Reed et al. (2000, E.2)	Rat	30	68	42				39				0.15	0.3
Reed (2001, E.1 & 4)	Human	8	50	37	46							0.25	0.3
Reed (2007a, E. 1)	Human	30	139	80				71				0.10	2.0
Reed (2007b, E.2, low force)	Rat	30	66	42				39				0.15	0.3
Reed (2007b, E.2, high force)	Rat	30	42	21				15				0.35	0.5
Tanno & Sakagami (2008)	Rat	30	66	44	57	68			67			0.15	0.3
Mean		54	87	52	66	65	51	34	67	58	82	0.19	0.6

E=experiment. Each IRT-min identified by an asterisk is drawn from actual data. Sources are Peele, Casey, and Silberberg (1984, E.1), Figs. 2 and 3; Peele et al. (1984, E.2), Figs. 4 and 5; Peele et al. (1984, E.3), Fig. 7; Baum (1993), Fig. 8; Cole (1999), Fig. 1

characteristics: (1) a standard VR–VI rate difference at moderate reinforcement rates, (2) a tendency for these differences to be reduced and then to reverse as reinforcement density increases, and (3) bitonicity of VR- and VI-rate functions. In our view, these results, in conjunction with those in the earlier simulations, support the application of the copyist model not just to variable schedules that differ in type, but also to those that differ in rate of reinforcement.

The conclusion we reach may surprise readers familiar with Reynolds and McLeod (1970). They claimed to show analytically that rate changes on variable schedules cannot be due to differential IRT reinforcement, as required by our model, because the distributions of IRTs these schedules reinforce are in a fixed mathematical relationship to emitted IRTs that is independent of schedule value (see also Wearden & Clark, 1988, 1989). How do we reconcile their work with ours? In part, we turn to Casey (1984). He denied the validity of Reynolds and McLeod's criticisms as they relate to VI by demonstrating mathematically and by simulation that (1) the distribution of IRTs emitted on VI schedules can shape IRT emission to conform to the reinforced IRT distribution and (2) increases in VI reinforcement density change the distribution of reinforced IRTs in a way that should shape response rate increases on VI schedules.

Casey's (1984) arguments aside, the results of our simulations mean that Reynolds and McLeod's (1970) work is

not applicable to our model. A likely reason for the failure of their work to accommodate our results (aside from Casey's critique of the VI case) is that IRT emission on our simulated VR and VI schedules violate a tenet of Reynolds and McLeod's analysis—that is, that IRT emission be exponentially distributed. Although the rule for response emission in the copyist model is indeed based on an exponential, many of the brief IRTs that would be required to create an exponential distribution are gated out by the model's IRT-min requirement. In consequence, the tests of Reynolds and McLeod, even if immune from criticism as offered by Casey, may not make contact with the model as presently structured. Given the fact that there truly is a minimum IRT that an organism can emit, and given the fact that its existence violates the generation of exponentially distributed responding in our simulations, we need not accommodate predictions that follow from Reynolds and McLeod's assumption of exponentially distributed response emission.

A portion of the credence given to simulations relates to their compatibility with intuition. For example, the descending limb of the VR response rate function created when VR reinforcer rates move from moderate to high in the top panel of Fig. 3 is easy to intuit because richer VRs naturally produce more post-reinforcement pauses than do leaner VRs, an outcome that lowers the mean response rate. What we find more difficult to intuit, however, is why the VR response rate has an ascending limb as VR reinforcer rates move from low to moderate levels.

Fig. 2 Obtained versus predicted response rates from the copyist model with (top panel) and without (bottom panel) exponential-weighting scheme. Diagonal line indicates perfect matching between these two rates. Closed and open circles identify data points based on actual IRT-min data and estimated IRT-min data drawn from a best-fit determination, respectively. The R^2 statistics are presented for one- (“actual IRT-min”) and two-parameter (“estimated IRT-min”) versions of the model, as well as both versions combined

In this case, the frequency of response-rate-reducing postreinforcement pauses increases across the range under consideration and should, therefore, lower response rates from what they would otherwise be. Why, then, do response rates go up as lean VR schedules get richer?

Imagine that a probabilistic emitter responds at the mean rate of once per second to a VR 10 and a VR 150. Each of the IRT strings composing each schedule’s IRI is exponentially weighted and stored in a 300-member reinforcement memory. The variance of the distribution of the VR 10 in reinforcement memory will be greater than that of the VR 150 because there will be more occasions for the VR 10 where sampling leads to a predominance of short or long IRTs that result in reinforcement than for the VR 150. If the variance in the distribution equals, say, 0.5 for the VR 10 and 0.1 for the VR 150, the respective harmonic mean IRTs of the distributions equal 0.54 and 0.99 s. When expressed as response rates, these means are reciprocated, resulting in higher response rates to the VR 10

than to the VR 150. This interpretation rationalizes the ascending limb apparent in Fig. 3 when VR reinforcement moves from lean to moderate levels.

As was shown earlier in the top panel of Fig. 3, the copyist model does fairly well in accommodating Baum’s (1993) VR–VI rate data. How does its predictive adequacy compare with that of other models that predict the VR–VI rate difference? While models that make such a prediction come in several forms (e.g., Allison, 1983; Baum, 1981; Catania, 2005; Killeen, 1994; Peele et al., 1984; Rachlin & Burkhard, 1978; Staddon, 1979), only Killeen’s mathematical principles of reinforcement (MPR) account has, to our knowledge, been evaluated quantitatively in terms of its predictive adequacy. Since we claim that a hallmark of the copyist model is its capacity to approximate actual data sets, we have selected MPR, which has also been applied to real data, as an appropriate comparison model.

In order to make a between-model comparison, we first define the MPR equations for VR and VI schedules, respectively:

$$B = \frac{\rho\lambda N}{1 + \lambda N} \cdot \frac{1}{\delta} - \frac{N}{A} \quad (2)$$

and

$$B = \frac{(\rho/\delta) \cdot R}{R + 1/A} - \frac{R}{\lambda}, \quad (3)$$

where B , R , and N represent response rate, reinforcer rate, and VR schedule value, respectively; A represents the time for which a reinforcer activates behavior; λ measures rate of decay of short-term memory; δ equals the reciprocal of the maximum response rate; and ρ defines the proportion of short-term memory occupied by target responses. Next, we fit Equation 2 to the VR data from Fig. 3 of Baum (1993), with ρ set at 1 and with A , δ , and λ as free parameters. Finally, we fit Equation 3 to the VI data from Fig. 3 of Baum (1993), with the values of A , δ , and λ for Equation 3 equaling the best-fit estimates used in modeling the VR results with Equation 2. In this fitting, ρ is the sole free parameter.

The bottom panel of Fig. 3 shows the response rates predicted by MPR for Baum’s (1993) data based on the following parameter estimates: $A = 180$, $\delta = 0.32$, $\lambda = 1.35$, and $\rho = 0.77$ for the VI schedule. Viewed qualitatively, MPR seems better than the copyist model (top panel of Fig. 3) in predicting VR–VI rate differences when reinforcer rates are low, but the copyist model does better when reinforcer rates are high. In terms of the R^2 statistic, MPR and the copyist model have equivalent explanatory power, although the copyist model does this with fewer free parameters.

MPR and the copyist model provide different views of the VR–VI rate difference. In MPR, ascending and

Fig. 3 Top panel: Response rates as a function of reinforcement rates. Baum's (1993) VR and VI data are presented by closed and open circles, respectively. The predictions of the copyist model are delineated by solid and dashed lines, respectively. Bottom panel: The predictions of MPR are delineated by solid and dashed lines, respectively

descending limbs of the response rate in Fig. 3 are explained by low arousal and low response–reinforcer association, respectively. In contrast, the copyist model explains these two limbs by the high variability of reinforced IRTs and the larger number of postreinforcement pauses, respectively. This means not only that the copyist model explains Baum's (1993) data with fewer free parameters than does MPR, but also that the hyperbolic response-rate, reinforcement-rate function observed in single VI schedules can be explained by IRT reinforcement (Catania & Reynolds, 1968). Further studies are needed to test which variable, reinforced-IRT variability or the density of postreinforcement pausing is more responsible for the hyperbolic form this function depicts.

Simulation 4: Modeling performances on concurrent VI VI schedules

In the copyist model, animals are posited to accept as an input to memory the exponentially weighted mean IRT created by responses between successive reinforcers and then to produce as outputs IRTs drawn from a distribution of IRTs whose mean equals the previously calculated mean IRT, their likelihood of selection being weighted by the reciprocals of their lengths. While we view the account as described above as simple, the language in which it is couched may mask that fact.

Fig. 4 Simulated relative response and time allocations (top left panel) and relative local response rates (bottom left panel) as a function of the relative reinforcement rate for VI-*b* in concurrent VI *a* VI *b* schedules for the copyist model. The diagonal line in the top panel defines perfect matching. The horizontal line in the bottom left panel defines equal local response rates to each schedule. The right panels present the same data as the left panels, except that they are from the target publications (Herrnstein, 1961, Fig. 1; Stubbs & Pliskoff, 1969, Fig. 1) the left column of data is intended to simulate

Fig. 5 Simulated relative response and time allocations (top left panel) and relative local response rates (bottom left panel) as a function of the relative reinforcement rate for VI in concurrent VR VI schedules for the copyist model. The diagonal line in the top panel defines perfect matching. The horizontal line in the bottom left panel defines equal local response rates to each schedule. The right panels present the same data as the left panels, except that they are from the target publication (Herrnstein & Heyman, 1979, Table 2) the left column of data is intended to simulate

To underscore its psychological simplicity, we introduce the notion of *dwell time*. A dwell time is the amount of session time an animal spends emitting a particular IRT. As was noted earlier and will be shown in more detail below, the purpose of the IRT-weighting scheme in the copyist model is to ensure that in simulated performances, each reinforced IRT in reinforcement memory has the same dwell time. Given that there are 300 reinforced-IRT values in reinforcement memory, the copyist model arranges response emission so that 1/300th of session time, on average, is spent emitting each remembered, previously reinforced IRT.

To show the connection between dwell time and the copyist model's IRT-weighting scheme, imagine that a bird has a high value of λ so that the exponentially weighted mean value of the reinforced IRT equals that of the reinforced IRT. Also imagine that, in a session, a simulated performance consists of just two IRIs and that the IRT that results in reinforcement is 2 s for one IRI and 4 s for the other. On the basis of the response algorithm of the copyist model, the probability of emission of 2-s IRTs is twice that of 4-s IRTs (i.e., the frequency of one reinforced 2-s IRT times the reciprocal of its length [$1 \times \frac{1}{2}$] is twice that of the 4-s IRT [$1 \times \frac{1}{4}$]).

When stated in terms of IRT frequencies, the algorithm-defined emission likelihood for the 2- and 4-s IRTs is, respectively, two to one, even though their representation in reinforcement memory is on a one-to-one basis. At first blush, this algorithmic rule may seem an arbitrary device to

weight simulated IRT selection toward shorter IRTs. But when stated in terms of dwell times, it becomes clear that the algorithm's IRT-weighting rule is used to ensure that in simulated performances, the same portion of session time is devoted to each reinforcer received. For example, if the session lasted 2 min and half of this session time was devoted to emitting 2-s IRTs and the other half to 4-s IRTs, thirty 2-s IRTs and fifteen 4-s IRTs would be obtained in the session. In other words, weighting reinforced IRTs by the reciprocals of their length ensures that each reinforced IRT in reinforcement memory occupies the same proportion of session time. This dwell time transformation of reinforced-IRT distributions makes response emission copyist: Animals are modeled as repeating the IRTs that were reinforced, with each reinforced IRT in reinforcement memory occupying the same amount of session time as any other. In our view, this is a psychologically simple notion indeed.

The copyist model posits that responding to single variable schedules is little more than a replay of the behaviors that were previously reinforced. So far we have shown that a model that posits that animals replay previously reinforced behaviors can accommodate a range of single-schedule effects. Now we test whether it can also predict choice behavior on concurrent VI VI schedules. To make such an evaluation, the copyist model has been modified so that all IRT emission is tagged not just in terms of its duration and proximity to reinforcement, but also in terms of the VI schedule to which it is assigned.

As in the earlier simulations, the last 300 reinforced IRTs are stored in reinforcement memory. To apply an

Fig. 6 On the left, simulated relative VR-*b* response rates and relative VR-*b* time allocations as a function of (1) the relative VR-*b* reinforcement probability (top panel) and the relative rate of VR-*b* reinforcement (bottom panel) in concurrent VR *a* VR *b* schedules for the copyist model. The vertical line in the top panel defines indifference in relative reinforcement probability. The diagonal in the bottom panel defines equation between the relative *x*- and *y*-axis measures. The right panels present the same data as the left panels, except that they are from the target publication (Herrnstein & Loveland, 1975, top panel of Fig. 1; top four panels in Fig. 3) the left column of data is intended to simulate

exponential weighting to an IRT, each schedule's IRTs are viewed in isolation. For example, if, in simulated choice performances, IRTs are emitted in the sequence “schedule *a*, schedule *a*, schedule *b*, schedule *a*” and the last response to schedule *a* is reinforced, only the calculation of the exponentially weighted mean IRT for schedule *a* is entered into reinforcement memory, while the durations of predecessor schedule *b* IRTs are included in calculations of temporal distance from reinforcement. Next, a schedule-tagged, exponentially weighted mean IRT is randomly drawn from reinforcement memory. It defines the mean of an exponential distribution of IRTs from which one is randomly selected to be made to the schedule for which the distribution is tagged.

A changeover delay (COD), a contingency sometimes used to ensure that the first response to a different schedule does not result in reinforcement (see, e.g., Herrnstein, 1961), was not used in this or subsequent simulations. There were three reasons for this step. First, it was unclear how this should be added to the simulation algorithm since, as the algorithm is currently defined, it is ill-equipped to handle a sequence of a period of extinction prior to access to a VI schedule—the very contingency that a COD such as Herrnstein's (1961) imposes. Second, even were this problem resolved, we consider it improbable that the revised algorithm would accommodate the robust and perplexing finding that local response rates within the COD period are higher for the alternative with the lower reinforcement rate (see Silberberg & Fantino, 1970). And finally, there is a

rationale for bypassing these complexities, because there is evidence that choice distributions are not affected much by assessing choice without a COD (see, e.g., the 0-s COD condition in Stubbs & Pliskoff, 1969). In any case, while a COD was not used in our choice simulations, changeover time between spatially distinct manipulanda was modeled by adding 0.2 s to an IRT-min parameter when the model dictated a change in the locus of choice.

The simulation engine was exposed to all possible paired combinations of VI 20-s, 40-s, 60-s, and 80-s schedules, with both schedules independently arranged (e.g., Herrnstein, 1961). The values of λ and IRT-min equaled 0.6 and 0.19 s, respectively, the means across all studies shown in Table 4 (so there was no free parameter in this simulation). Each simulation ended after 50,000 responses, the last 10,000 of which were used for data analysis. In other regards, this simulation is the same as those that precede it.

Often, choice between schedules is analyzed in terms of its concordance with the predictions of the generalized matching law:

$$\log(B_a/B_b) = \alpha \log(R_a/R_b) + \log \beta, \quad (4)$$

or

$$\log(T_a/T_b) = \alpha \log(R_a/R_b) + \log \beta, \quad (5)$$

where *B*, *R*, and *T* denote responses, reinforcers, and time allocations to each alternative (*a* and *b*), respectively. The

parameter α measures the *sensitivity* of the response or time ratio to the reinforcer ratio, while β measures *bias* between alternatives (see Baum, 1974, for further discussion of these measures). As reviewed by Baum (1974, 1979), sensitivity in concurrent VI VI schedules is usually less than one when measured in terms of responses but approximately one when measured in terms of time.

Results and discussion of simulation 4 The top left panel of Fig. 4 presents the simulated preference data in terms of relative-response (responses to VI *b*/all responses) and relative-time (time in the presence of VI *b*/session time) allocation for the 16 different conditions of relative reinforcement used in the present simulations. The diagonal defines matching—the equation between relative measures of time and response and relative rate of reinforcement. As is apparent, points of both dependent variables are positioned around the diagonal line.

The bottom left panel of Fig. 4 presents the relative local response rates from the copyist model. To make this calculation, the local response rates to each alternative, defined as the responses to an alternative divided by the time in that alternative's presence, were determined. Then the local response rate to VI *b* was divided by the sum of the local rates to VI *a* and VI *b*. The figure plots these data as a function of relative reinforcement rate. As is apparent, the local response rates to both schedules were approximately the same regardless of relative reinforcement rate.

To evaluate the adequacy of model predictions, we plot in the top right panel the relative response- and time-allocation data from two concurrent VI VI studies, those from Herrnstein (1961, Fig. 1) and Stubbs and Pliskoff (1969, Fig. 1). The bottom right panel presents the relative local response rates seen in Stubbs and Pliskoff.

In terms of the top panels, there is a correspondence between the actual data shown in the right panel and the simulated results in the left panel. Moreover, when fit to Equations 4 and 5, the simulated performances mapped out in the left panel are consistent with Baum's (1974, 1979) description of choice on these schedules, with the estimates of α and β being .90 and $-.01$, respectively, for Equation 4 and .99 and .00 for Equation 5.

Relative local response rates are not available for Herrnstein (1961). However, they are available for Stubbs and Pliskoff (1969), who reported across a range of VI schedule pairs that provided different relative rates of reinforcement that local rates did not differ between schedules as a function of the schedule pair (see bottom right panel). The simulated performances in this bottom panel are echoed in the data from our simulation. In sum, this figure shows that the copyist model does an effective job in reproducing standard outcomes seen on concurrent VI VI schedules. To our knowledge, no other model can make such a claim.

Simulation 5: Modeling performances on concurrent VR VI schedules

As was suggested above, application of the copyist model to concurrent VI VI schedules generates representative choice data and, in this regard, raises the possibility that this model integrates single- and concurrent-schedule performances within a unitary scheme in which choice seems derivative. The goal of the present simulation is to model choice behavior on concurrent VR VI schedules and, by so doing, offer a test the generality of this conclusion. To this end, the program used in simulation 4 was modified by replacing one of its VI schedules with a VR. The schedule values of the VR were 20, 40, 60, and 80. In other regards, the simulation engine and procedure were unchanged from those used in the prior simulation.

In Herrnstein and Heyman's (1979) report on concurrent VR VI schedules, relative response and time allocation approximately track the slope of the matching loci; however, a bias is evident in the unequal distribution of response and time data around the matching line. In particular, when measured in terms of time, they found choice allocation biased toward the VI; however, when measured in terms of responses, the bias was toward the VR—an outcome they attributed to the higher local response rate on their VR schedules.

Results and discussion of simulation 5 The top left panel of Fig. 5 presents the results of our simulation. When fit by Equations 4 and 5, α and β were .91 and $-.21$, respectively, in terms of relative response rate and .89 and $-.03$ in terms of relative time allocation. As reflected in these statistics and in a comparison between these simulated performances and actual data from Herrnstein and Heyman (1979, Table 2), the simulation reproduced two empirical effects seen in the target study: (1) The slopes of the relative-response and relative-time data approximated the slope of matching, and (2) response, but not time, data were biased toward the VR.

The bottom panels of Fig. 5 present local response rate as a function of relative reinforcement rate for simulated performances (left side) and for Herrnstein and Heyman (1979, right side). Except when preference was exclusive, local response rates were approximately 1.5 times higher on VR than on VI, a result similar to Herrnstein and Heyman's. In sum, the copyist model seems to do a reasonable job in reproducing the effects found on concurrent VR VI schedules.

Although not central to the purposes of this report, it is an interesting aside to look at the interpretation Herrnstein and Heyman (1979) lent their report and the implications of our simulated outcomes for it. An active topic at the time of that publication was whether matching as outcome

could be attributed to matching as process or, alternatively, maximizing as process. Since they found matching with bias that took a form that was incompatible with maximizing reinforcer rates, they argued for the primacy of a matching process. To the extent that the present simulation is viewed as successful, it makes an interesting comment on this debate: By virtue of its copyist nature, neither matching nor maximizing governs choice on concurrent VR VI schedules. Instead, choice is due to the reproduction of dwell time allocations that had been reinforced in the past to these concurrently available schedules (see also Tanno, Silberberg, & Sakagami, 2010).

Simulation 6: Modeling performances on concurrent VR VR schedules

In Herrnstein and Loveland (1975), pigeons chose across conditions between ten different pairs of VR schedules that varied both in the density of reinforcement they provided and in their relative probability of reinforcement. They reported that pigeons showed nearly exclusive preference for the alternative with higher reinforcement probability (lower VR value). The present simulation evaluates whether the copyist model, adapted to choice, can reproduce the major effects they reported. To this end, the algorithm used in the prior simulation was modified so that simulated choice was between concurrent VR VR schedules. Responding in this simulation was to all possible paired combinations of VR 20, VR 40, VR 60, and VR 80 schedules. In other regards, this simulation was procedurally identical to its predecessor.

Results and discussion of simulation 6 The top panels of Fig. 6 present the relative response and time allocations to one of the two alternatives as a function of the relative scheduled likelihood of reinforcement. To illustrate definition of the independent variable, imagine that alternatives *a* and *b* provide reinforcement according to a VR 40 and a VR 60 schedule, respectively. On the basis of these VR pairs, the relative probability of reinforcement for each alternative-*b* response is $1/60 / (1/40 + 1/60)$, or .4. In these panels, a vertical line is placed at .5 on the *x*-axis. Points that fall on this vertical line maximize reinforcement rate. Points to the left maximize when they have an *x*-axis value of 0, while points to the right maximize at a value of 1.0. Generally speaking, in both the simulation (left side) and Herrnstein and Loveland (1975, right side), whose work provided the target data set, the results look similar and often consistent with maximizing.

The bottom panels of Fig. 6 present the same choice data from our simulation (left side) and Herrnstein and Loveland (1975) as a function of the relative obtained frequency of reinforcement. The diagonal through the graph defines

perfect matching. As is apparent, on the basis of this measure, matching obtains on concurrent VR VR schedules in both the simulation and target data.

As had been the case with concurrent VR VI, our reading of Herrnstein and Loveland (1975) suggests that their interest in concurrent VR VR schedules was motivated by their perceived potential in resolving the maximizing–matching debate. While their data were often consistent with both maximizing and matching, they note that in several instances in their data set, pigeons' preferences for the richer VR schedule were not exclusive. In this regard, pigeons occasionally failed to maximize. However, by their analysis and as shown in the bottom right panel, all choice allocation in their study was compatible with matching predictions. Such a result could be viewed as evidence of the primacy of matching processes in choice, even when maximizing is in evidence. Whatever the merits of this argument might be, the conclusion supported by the copyist model's reproduction of their results is clear-cut: Maximizing (top panels) and matching (bottom panels) are predicted by the copyist model of response emission.

This reproduction of Herrnstein and Loveland's (1975) results suggests that the copyist model's response algorithm is descriptive of the process governing choice on concurrent VR VR schedules; and since this success echoes those in our earlier choice simulations, it now appears that no matter how VR and VI schedules are combined, a copyist process may govern choice.

Simulation 7: Melioration

Vaughan (1981) attributed the finding of matching on concurrent schedules to a process he called *melioration*: the equation of local rates of reinforcement between schedules. To illustrate, imagine an animal chooses in a 60-min session so that 30 min are spent responding to one schedule (a VI 1-min schedule that provides 60 reinforcers/h) and 30 min to another (a VI 2-min schedule that provides 30 reinforcers/h). Given that the time allocations are 30 min to each schedule in the concurrent pair, the local rate of reinforcement to the VI 1-min schedule should be twice that to the VI 2-min schedule (60 reinforcers/30 min, or 2 reinforcers/min, vs. 30 reinforcers/30 min, or 1 reinforcer/min). To equate the rates of reinforcement as dictated by the melioration process, the animal should change preference to allocate twice as much time to the VI 1-min schedule as to the VI 2-min schedule (i.e., 60 reinforcers/40 min equals 30 reinforcers/20 min). Such a two-to-one preference for the VI 1-min schedule in time allocation is identical to the predictions of matching.

Vaughan (1981) wished to determine whether melioration could be responsible for a matching outcome in a

choice context in which matching outcomes could not be attributed to maximizing rates of reinforcement. To that end, he arranged a choice procedure for pigeons in which each VI operated only while its associated key was chosen, and the rate of reinforcement each key provided was based on pigeons' time allocations during a prior period (see his report for more procedural details). By virtue of this unusual arrangement, he created a procedure that separated the predictions of reinforcer-rate maximizing from those of melioration. The results he obtained were compatible with a melioration account, but not those of maximizing.

In the present simulation, the choice allocation rules used in the prior simulation were used on the procedures of Vaughan (1981) and Silberberg and Ziriax (1985), who also used Vaughan's (1981) design in a melioration assay. All conditions within each report were simulated save conditions 4 and 5 from Silberberg and Ziriax. They were excluded because they did not list the schedule values used in these conditions. In other regards, the methods used in prior choice simulations operated here as well.

Results and discussion of simulation 7 The top panel of Fig. 7 presents simulated relative VI-*b* response rates as a function of their relative VI-*b* reinforcement rates obtained using the Vaughan (1981) and Silberberg and Ziriax (1985) procedures. The bottom panel presents the simulated relative VI-*b* response rate as a function of the actual relative VI-*b* response rates obtained in these two studies. For both panels, data points seem acceptably close to the diagonal defining the loci of matching and melioration. In sum, this figure suggests that procedures designed to demonstrate melioration outcomes are compatible with a copyist choice algorithm. Given that Silberberg and Ziriax interpreted their findings as demonstrating a version of maximizing, the results of this choice simulation join our earlier simulations applied to choice in showing that matching, melioration, and maximizing outcomes may reflect the control of behavior by a rule in which animals simply replay previously reinforced behavior.

General discussion

Summary of tests and predictions of the copyist model

This article began by noting that counterpoise tests have been used by operant analysts to evaluate whether the difference in response rates between VR and VI schedules that provide the same reinforcement rate should be attributed primarily to between-schedule differences in IRT reinforcement (e.g., Peele et al., 1984) or to each schedule's molar feedback function (e.g., Baum, 1973). For example, all other features being equal, including reinforcement rates and the duration of tandem schedule

Fig. 7 Simulated relative response rate to VI *b* as a function of (1) the simulated relative VI-*b* reinforcement rate (top panel) and (2) the obtained relative VI-*b* response rate from the target publications (Silberberg & Ziriax, 1985, Table 2; Vaughan, 1981, Table 1). The diagonals in each figure define equation between these variables

components, a tandem VR VI should support a lower response rate than should a tandem VI VR if response emission is controlled at the level of the IRT but should generate equivalent rates if emission is controlled at the molar level of the relation between reinforcement rate and response rate (see Peele et al., 1984, Experiment 3). As was noted earlier, the outcomes of this test and several others suggest that an IRT-based model is predictively advantaged over one based on molar feedback functions.

Despite this fact, our review points to a clear-cut exception to the generalization that molecular factors control response rate: Rates on tandem VI DRH schedules consistently fall short of matching the rates maintained by a VR schedule that provides the same rate of reinforcement (Cole, 1994; Reed et al., 2000, Experiment 4; Tanno & Sakagami, 2008). It was this outcome that motivated our attempt to revise Peele et al.'s (1984) rate emission model. Toward this end, we modified the Peele et al. account to enable it to make its IRT selection not from a memory composed solely of reinforced IRTs, but from an average of sequential IRTs in which the weight assigned to each IRT composing the mean diminishes exponentially as a function of its backward distance from reinforcement.

The copyist model this change creates, along with other minor modifications in the original Peele et al. (1984) account, had felicitous effects. As is shown in Fig. 1, it now accommodates both the VR–VI rate difference and the finding that on tandem VI DRH, rates are typically lower than on a target VR schedule. In addition, it lowers simulated response rates into a range that is representative of what animals produce, a capacity the Peele et al. model lacked.

To test the generality of the copyist model in accommodating counterpoise tests, we compared its simulated performances against those of every such schedule of which we are aware. As is shown in Fig. 2, the copyist model does a credible job in reproducing performances on the several different types of schedules that have been used to test molecular and molar accounts of the VR–VI rate difference. This result gives credence to the claim that the rate differences VR and VI schedules produce can be viewed as a consequence of between-schedule differences in how reinforcement contacts the stream of IRTs proximal to it.

The copyist model was designed to address the failure of molecular accounts to accommodate the rate shortfall in tandem VI DRH response rates when they were compared with a VR. Given its success in this counterpoise test, its application to other counterpoise tests, as shown in Fig. 2, was a necessary outcome. Indeed, for the model to have utility, it also must succeed, as it did, in explaining outcomes on these other types of schedules. But where the success of the copyist model was not necessary for claimed utility was in its application to the different problem posed by Baum's (1993) data—the mapping of response-rate, reinforcement-rate relations on VR and VI schedules.

When reinforcer rates are systematically varied on VR and VI schedules that provide the same rate of reinforcement, Baum (1993) found that (1) a standard VR–VI rate difference at moderate reinforcement rates, (2) a tendency for these differences to be reduced and then to reverse as reinforcement density increases, and (3) bitonicity of VR-and VI-rate functions. As is shown in the bottom panel of Fig. 3, these features of his results were reproduced qualitatively by the copyist model. Moreover, the model did a fair job of predicting the absolute response rates obtained (see the top panel of Fig. 3).

We thought that there is no reason this copyist process should be limited to single schedules and, therefore, conducted simulations when choice was between concurrent VI VI, VR VI, and VR VR schedules, as well as melioration test schedules. To do this, we altered the copyist model used with single schedules to record an additional feature—the locus of choice.

When this was done, the performances approximated major features of the Herrnstein (1961), Stubbs and Pliskoff (1969), Baum (1974), Herrnstein and Heyman (1979), Herrnstein and Loveland (1975), Vaughan (1981), and Silberberg and Zirrax (1985) studies that were targeted for simulation. This model, and no other, provides a comprehensive account of both matching and local response rate in choice procedures. Moreover, this outcome raises the prospect that the matching phenomenon itself is derived from the response emission algorithm on which the copyist model is based.

In sum, the copyist model (1) is parsimonious (it can often be expressed with a single free parameter), (2) provides breadth of coverage (it can model adequately a range of single and concurrent variable-schedule effects), and (3) is conceptually simple (it says that animals equate the portions of session time emitting each reinforced IRT that the animal recalls). To our knowledge, no extant account of operant output shares these successes.

Limitations of the copyist model

The text heretofore stresses the strengths of the copyist account. Yet it has weaknesses as well. One limitation is its inadequacy in explaining a subset of single-schedule effects in simulations 1 and 2, where some evidence of molar control of responding can be discerned (Dawson & Dickinson, 1990; Reed, 2007a, 2007b, Experiment 2, low-force data; see Tables 2 and 3). Because of these and other data sets where molar factors seem operative (Heyman & Tanz, 1995; Reed, 2006; Sakagami, Hursh, Christensen, & Silberberg, 1989; Shurtleff & Silberberg, 1990; Silberberg, Thomas, & Berendzen, 1991; Soto, McDowell, & Dallery, 2006), it appears that operant performances are not always controlled only at a molecular level. That having been said, the operation of molar factors in responding, at least to a degree, need not violate the copyist model, because it blends molar and molecular accounts by considering predecessor IRTs along with reinforced IRT in its definition. While the exponential-weighting scheme used in its IRT selection algorithm biases it toward being a molecular account, the model's incompatibility with the concurrent operation of molar processes is not absolute. And this fact invites consideration of the degree to which rate effects in studies evidencing molar schedule effects stand in opposition to our model.

To illustrate, when operant performances from human and nonhumans on similar schedules are compared, differences are sometimes apparent. For example, negative-feedback schedules, where reinforcement rates decrease as response rates increase, affect rate emission in humans

(Jacobs & Hackenberg, 2000), but not in rats (Ettinger et al., 1987; Reed & Schachtman, 1991; Tanno & Sakagami, 2008). In this circumstance, the predictive adequacy of the copyist model is necessarily impaired (it cannot predict species differences in outcome). Nevertheless, this model remains predictively useful ($R^2 = .91$) even if the data under consideration are limited to those from humans (see Table 2). Such an outcome is consistent with the idea that molecular processes are still operative in human schedule effects, albeit less exclusively than in other animals. Moreover, this underscores our observation above that the incompatibility between molar and molecular accounts in the copyist model may be more a matter of degree than of kind.

Another possible weakness of the model is that it assumes that animals are continuous response emitters. There is some debate as to whether this assumption is true. For example, Shull (2004) has shown that behavior often seems discontinuous, alternating between periods of action and inaction, while Bowers et al. (2008) have argued that response emission on some schedules seems well characterized by assuming that emission is continuous. Evaluation of the integrity of the copyist model's assumption that response emission is continuous awaits further work. Should the conclusions that follow show that schedule performances are generally episodic, the emission engine of the copyist model would require modification.

Another possible threat to the integrity of the copyist account comes from data sets showing that in some circumstances, response emission is “unitized”—that is, sequences of responses or behavioral patterns may define a behavioral unit that is itself amenable to schedule control (e.g., Shimp, 1979; Zeiler, 1986). If the response unit should, in fact, be a behavior stream and not a single IRT, the relevance of the copyist model, dependent as it is on modeling individual IRT selection, would be seriously weakened. We sound a cautionary note regarding endorsing this criticism of the copyist model. Generally speaking, procedures structured to demonstrate the shaping of behavioral patterns do not provide evidence that behavioral patterning is a defining characteristic of behavioral output in experiments not so structured. In particular, a common characteristic of studies demonstrating the occurrence of behavioral patterning is that reinforcement of the pattern does not occur for successive responses and they tend to make use of a fixed number of behaviors to define the sequence (e.g., Pisacreta, 1982; Reid, Chadwick, Dunham, & Miller, 2001; S. M. Schneider & Morris, 1992; Schwartz, 1982). This is hardly surprising, since reinforcing successive behaviors prohibits defining a multibehavior unit. Yet the schedules we model often produce reinforcers for successive IRTs. Complicating unit definition still more, the sequences between successive reinforcers often differ in length in our work. The problem variability in sequence length creates for “unitized”

modeling is clear: How can an animal learn a behavior sequence when the number of behaviors composing the sequence is variable?

A clear limitation of our model is its microstructural mischaracterization of IRT emission, at least for pigeons. Ample data show that pigeons' VR and VI IRT distributions tend to be multimodal (see, e.g., Peele et al., 1984, Fig. 2), but IRT distributions generated by the copyist model are unimodal. Finally, the present form of the copyist model cannot accommodate temporally patterned performances, such as FI, or reproduce acquisition and extinction functions.

Copyist model and Guthrie's contiguity account

We view IRT emission on schedules as due to a copyist process. This notion may remind the reader of a contiguity account from Guthrie (1930, 1933). As he stated often, “stimuli which are acting at the time of a response tend, on their recurrence, to evoke that response” (Guthrie, 1933, p. 355). His account and ours are conceptually similar because both are copyist: Animals are thought to reproduce the movements that occurred proximal to reinforcement.

To our reading, Guthrie (1930, 1933) brooks no exceptions to the idea that changes in behavior that accompany changes in independent variables are due to reproduction of prior movements that occurred in the presence of the stimuli cuing reinforcement. He can do this because he is silent about just how and whether those cuing stimuli should operate. To illustrate, consider Clark's (1958) finding that increasing hunger levels in rats responding on VI schedules results in increases in response rates. This result presents no problems for Guthrie's contiguity rule because, for any two hunger levels being compared, the stimulus conditions differ, and different stimulus conditions can support different response rates. Nevertheless, Guthrie's account would not be invalidated if these different stimulus conditions did not result in different response rates or, for that matter, supported rates opposite to those obtained. Such results would mean only that these stimuli, viewed no doubt by the scientist as important cues attending reinforcement that should (or should not) condition particular movements, were not so viewed by the subject.

Given this latitude, Guthrie (1930, 1933) can afford a “strong” contiguity rule—that is, a view that there are no exceptions to his rule and no additional processes in operation that are relevant to explain changes in behavior. The copyist model, on the other hand, lacks these latitudes in application because all of its predictions flow from its explicitly stated algorithm. Changes in hunger level have no representation in the algorithm, so, of course, if rates change

with changing hunger levels—a fact evidenced by Clark (1958)—the copyist model errs in prediction.

For this reason, we advance a “weak” version of copyist principles. According to this view, using, say, a VI schedule to measure the effects of an independent variable on response rate will dampen evidence of that variable’s operation because response emission tends to copy what was reinforced before. This is to say not that response rate will not change but, rather, that the changes, whatever they might be, will be muted to a degree by the copyist-inducing feature of IRT reinforcement on VI. While IRTs will tend toward invariance despite changes in hunger level in the Clark (1958) study, this manipulation may nevertheless affect response rate through, say, changes in the latency to begin responding in a session or the proportion of the session that an animal spends in a state of responding. Such a view is corroborated by Shull’s (2004) finding that changing rats’ hunger levels has a larger effect on how long and how frequently they are in a state of responding than it does on the properties of IRT emission.

Our enunciation that the copyist model is weak points to a future direction model making based on the copyist response engine might take—namely, finding a way to make it strong. Data such as Shull’s (2004) are consistent with the copyist model as articulated here. That is, response rate changes induced by variations in hunger levels have little effect on response emission within a response bout; however, they do affect when an animal begins a bout and what portion of session time is devoted to bout-based responding. For want of a better term, consider the elements of latency to respond and bout duration to be reflective of “strength,” a covariation between the portion of session time devoted to responding and a manipulation thought to affect the value of reinforcement. Clearly, this feature of behavior, ignored by the copyist model as currently stated, would make a weak model stronger were it added to the model. This conclusion follows from the fact that response rate is shaped by its relation to reinforcement, but the value of reinforcement is predicated on a second set of variables such as the hunger levels an experimental manipulation imposes. Obviously, additional work needs to be done to make this goal realizable.

Implications

In some cases, the similarities between the copyist model and Guthrie’s contiguity account make his views our views. One example with noteworthy implications is Guthrie’s (1930) argument that learning theorists err when they espouse that the vigor of behavior measures the strength of a reflex or the value of a reinforcer (see pp. 419–420). This view is also an implication of the

copyist model. For example, in the copyist model, response rates are higher on a VR than on a VI when reinforcer rates are equated, not because the VR is of more value or has greater associative strength, but because of differences in the IRTs reinforced by each of these schedules. It seems that in any strong version of the copyist accounts offered by Guthrie or us, these rate differences have little to do with value or strength.

Of course, as was noted earlier, we advance a weak version of the copyist model, one that acknowledges the operation of other components of response rate outside of IRT emission that, in practice, may be correlates of strength. However, since IRT emission is a major component in rate determination on many schedules, one should expect that response rate often would not be assessing a reinforcer’s value but, instead, the local feedback function between IRTs emitted proximal to the reinforcer and the way the schedule programs reinforcement.

This idea implies that the strength-scaling successes of two popular operant measures, those of response rate (Skinner, 1932b, 1938, 1950) and relative response rate (Herrnstein, 1970), are more apparent than real. As regards response rate, simulation 3 shows that, to a considerable degree, response rates on schedules like a VI and VR mirror the IRT distribution created when responding produces reinforcement. Given this simulation—which by virtue of its algorithm attributes response rate to the properties of the particular schedule maintaining behavior—why should response rate be imagined also to index effectively the strength-inducing properties of the reinforcer?

To the extent that the copyist model is veridical, the criticism of response rate offered above extends to the successor technique advanced by Herrnstein (1970), which measures strength not in terms of response rate to a schedule, but in terms of the relative response rates between schedules. According to the copyist model in simulations 4, 5, and 6, choice allocation in the form of matching is due to replaying the reinforced IRT distributions for each of the schedules offered in choice. As was argued above, this result does not scale relative value between alternatives but, instead, describes dwell time assignments drawn from the history of IRTs proximal to the reinforcers generated by these concurrently available schedules.

For some years, investigators have been in pursuit of alternatives to response rate and relative response rate as a means of scaling strength (e.g., Hodos, 1961; Hursh & Silberberg, 2008). Even if not acknowledged by these authors, we entertain that their efforts were at least implicitly driven by predictive inadequacies they saw in traditional measures. Perhaps a virtue of the present work is that it makes explicit what might have caused these inadequacies to appear.

Author Notes This work was supported by the Program of Global COE (Center of Excellence) #D029 in Japan, and Grand-in-Aid for JSPS Fellows #00237309. The authors thank Takayuki Sakagami, Jay Casey, Jonathan Katz, Peter Killeen, Charles Shimp, and reviewers for helpful comments on this work.

Some preliminary results of this article were presented in an annual meeting of the Japan Society for the Quantitative Analyses of Behavior.

References

- Allison, J. (1983). *Behavioral economics*. New York: Praeger.
- Angle, H. V. (1970). The role of reinforcement in controlling sequential IRT dependencies. *Journal of the Experimental Analysis of Behavior*, 14, 145–151.
- Baum, W. M. (1973). The correlation-based law of effect. *Journal of the Experimental Analysis of Behavior*, 20, 137–153.
- Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior*, 22, 231–242.
- Baum, W. M. (1979). Matching, undermatching, and overmatching in studies of choice. *Journal of the Experimental Analysis of Behavior*, 32, 269–281.
- Baum, W. M. (1981). Optimization and the matching law as accounts of instrumental behavior. *Journal of the Experimental Analysis of Behavior*, 36, 387–403.
- Baum, W. M. (1993). Performances on ratio and interval schedules of reinforcement: Data and theory. *Journal of the Experimental Analysis of Behavior*, 59, 245–264.
- Bowers, M. T., Hill, J., & Palya, W. L. (2008). Interresponse time structures in variable-ratio and variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, 90, 345–362.
- Casey, J. (1984). A model explaining how animals produce variable-interval rates. *Dissertation Abstracts International*, 43(8), 2715B (UMI No. ATT 8425724).
- 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, 271–287.
- Catania, A. C. (2005). The operant reserve: A computer simulation in (accelerated) real time. *Behavioural Processes*, 69, 257–278.
- Catania, A. C., Matthews, T. J., Silverman, P. J., & Yohalem, R. (1977). Yoked variable-ratio and variable-interval responding in pigeons. *Journal of the Experimental Analysis of Behavior*, 28, 155–161.
- Catania, A. C., & Reynolds, G. S. (1968). A quantitative analysis of the responding maintained by interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 11, 327–383.
- Clark, F. C. (1958). The effect of deprivation and frequency of reinforcement on variable-interval responding. *Journal of the Experimental Analysis of Behavior*, 1, 221–228.
- Cole, M. R. (1994). Response-rate differences in variable-interval and variable-ratio schedules: An old problem revisited. *Journal of the Experimental Analysis of Behavior*, 61, 441–451.
- Cole, M. R. (1999). Molar and molecular control in variable-interval and variable-ratio schedules. *Journal of the Experimental Analysis of Behavior*, 71, 319–328.
- Dawson, G. R., & Dickinson, A. (1990). Performance on ratio and interval schedules with matched reinforcement rates. *Quarterly Journal of Experimental Psychology*, 42B, 225–239.
- Ettinger, R. H., Reid, A. K., & Staddon, J. E. R. (1987). Sensitivity to molar feedback functions: A test of molar optimality theory. *Journal of Experimental Psychology: Animal Behavior Processes*, 13, 366–375.
- Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement*. New York: Appleton-Century-Crofts.
- Guthrie, E. R. (1930). Conditioning as a principle of learning. *Psychological Review*, 37, 412–428.
- Guthrie, E. R. (1933). Association as a function of time interval. *Psychological Review*, 40, 355–367.
- 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.
- Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior*, 13, 243–266.
- Herrnstein, R. J., & Heyman, G. M. (1979). Is matching compatible with reinforcement maximization on concurrent variable interval, variable ratio? *Journal of the Experimental Analysis of Behavior*, 31, 209–223.
- Herrnstein, R. J., & Loveland, D. H. (1975). Maximizing and matching on concurrent ratio schedules. *Journal of the Experimental Analysis of Behavior*, 24, 107–116.
- Heyman, G. M., & Tanz, L. (1995). How to teach a pigeon to maximize overall reinforcement rate. *Journal of the Experimental Analysis of Behavior*, 64, 277–297.
- Hodos, W. (1961). Progressive ratio as a measure of reward strength. *Science*, 134, 943–944.
- Hursh, S. R., & Silberberg, A. (2008). Economic demand and essential value. *Psychological Review*, 115, 186–198.
- Jacobs, E. A., & Hackenberg, T. D. (2000). Human performance on negative slope schedules of points exchangeable for money: A failure of molar maximization. *Journal of the Experimental Analysis of Behavior*, 73, 241–260.
- Killeen, P. R. (1994). Mathematical principles of reinforcement. *The Behavioral and Brain Sciences*, 17, 105–172.
- Kuch, D. O., & Platt, J. R. (1976). Reinforcement rate and interresponse time differentiation. *Journal of the Experimental Analysis of Behavior*, 26, 471–486.
- Matthews, B. A., Shimoff, E., Catania, A. C., & Sagvolden, T. (1977). Uninstructed human responding: Sensitivity to ratio and interval contingencies. *Journal of the Experimental Analysis of Behavior*, 27, 453–467.
- McDowell, J. J., & Wixted, J. T. (1986). Variable-ratio schedules as variable-interval schedules with linear feedback loops. *Journal of the Experimental Analysis of Behavior*, 46, 315–329.
- Morse, W. H. (1966). Intermittent reinforcement. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 52–109). New York: Appleton-Century-Crofts.
- Peele, D. B., Casey, J., & Silberberg, A. (1984). Primacy of interresponse-time reinforcement in accounting for rate differences under variable-ratio and variable-interval schedules. *Journal of Experimental Psychology: Animal Behavior Processes*, 10, 149–167.
- Pisacreta, R. (1982). Some factors that influence the acquisition of complex, stereotyped, response sequences in pigeons. *Journal of the Experimental Analysis of Behavior*, 37, 359–369.
- Rachlin, H., & Burkhard, B. (1978). The temporal triangle: Response substitution in instrumental conditioning. *Psychological Review*, 85, 22–47.
- Reed, P. (2001). Schedules of reinforcement as determinants of human causality judgments and response rates. *Journal of Experimental Psychology: Animal Behavior Processes*, 27, 187–195.
- Reed, P. (2006). Effect of required response force on rats' performance on a VI + schedule of reinforcement. *Learning & Behavior*, 34, 379–386.
- Reed, P. (2007a). Human sensitivity to reinforcement feedback functions. *Psychonomic Bulletin & Review*, 14, 653–657.
- Reed, P. (2007b). Response rate and sensitivity to the molar feedback function relating response reinforcement rate on VI + schedules of reinforcement. *Journal of Experimental Psychology: Animal Behavior Processes*, 33, 428–439.
- Reed, P., Hildebrandt, T., DeJongh, J., & Soh, M. (2003). Rats' performance on variable-interval schedules with a linear feedback

- loop between response rate and reinforcement rate. *Journal of the Experimental Analysis of Behavior*, 79, 157–173.
- Reed, P., & Schachtman, T. R. (1991). Instrumental performance on negative schedules. *Quarterly Journal of Experimental Psychology*, 43B, 177–197.
- Reed, P., Soh, M., Hildebrandt, T., DeJongh, J., & Shek, W. Y. (2000). Free-operant performance on variable interval schedules with a linear feedback loop: No evidence for molar sensitivities in rats. *Journal of Experimental Psychology: Animal Behavior Processes*, 26, 416–427.
- Reid, A. K., Chadwick, C. Z., Dunham, M., & Miller, A. (2001). The development of functional response units: The role of demarcating stimuli. *Journal of the Experimental Analysis of Behavior*, 76, 303–320.
- Reynolds, G. S., & McLeod, A. (1970). On the theory of interresponse-time reinforcement. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in theory and research* (Vol. 4, pp. 85–107). New York: Academic Press.
- Sakagami, T., Hursh, S. R., Christensen, J., & Silberberg, A. (1989). Income maximizing in concurrent interval-ratio schedules. *Journal of the Experimental Analysis of Behavior*, 52, 41–46.
- Schneider, B. A. (1969). A two-state analysis of fixed-interval responding in the pigeon. *Journal of the Experimental Analysis of Behavior*, 12, 677–687.
- Schneider, S. M., & Morris, E. K. (1992). Sequences of spaced responses: Behavioral units and the role of contiguity. *Journal of the Experimental Analysis of Behavior*, 58, 537–555.
- Schwartz, B. (1982). Interval and ratio reinforcement of a complex, sequential operant in pigeons. *Journal of the Experimental Analysis of Behavior*, 37, 349–357.
- Shimp, C. P. (1969). The concurrent reinforcement of two interresponse times: The relative frequency of an interresponse time equals its relative harmonic length. *Journal of the Experimental Analysis of Behavior*, 12, 403–411.
- Shimp, C. P. (1973). Sequential dependencies in free-responding. *Journal of the Experimental Analysis of Behavior*, 19, 491–497.
- Shimp, C. P. (1974). Time allocation and response rate. *Journal of the Experimental Analysis of Behavior*, 21, 491–499.
- Shimp, C. P. (1979). The local organization of behavior: Method and theory. In M. D. Zeiler & P. Harzem (Eds.), *Reinforcement and the organization of behavior* (pp. 261–298). Chichester, U.K.: Wiley.
- Shimp, C. P., Fremouw, T., Ingebritsen, L. M., & Long, K. A. (1994). Molar function depends on molecular structure of behavior. *Journal of Experimental Psychology: Animal Behavior Processes*, 20, 96–107.
- Shull, R. L. (2004). Bouts of responding on variable-interval schedules: Effects of deprivation level. *Journal of the Experimental Analysis of Behavior*, 81, 155–167.
- Shurtleff, D., & Silberberg, A. (1990). Income maximizing on concurrent ratio-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 53, 273–284.
- Silberberg, A., & Fantino, E. (1970). Choice, rate of reinforcement, and the changeover delay. *Journal of the Experimental Analysis of Behavior*, 13, 187–197.
- Silberberg, A., Thomas, J. R., & Berendzen, N. (1991). Human choice on concurrent variable-interval variable-ratio schedules. *Journal of the Experimental Analysis of Behavior*, 56, 527–584.
- Silberberg, A., & Zirias, J. M. (1985). Molecular maximizing characterizes choice on Vaughan's (1981) procedure. *Journal of the Experimental Analysis of Behavior*, 43, 83–96.
- Skinner, B. F. (1932a). Drive and reflex strength. *The Journal of General Psychology*, 6, 22–37.
- Skinner, B. F. (1932b). Drive and reflex strength: II. *The Journal of General Psychology*, 6, 38–48.
- Skinner, B. F. (1938). *The behavior of organisms*. New York: Appleton-Century-Crofts.
- Skinner, B. F. (1950). Are theories of learning necessary? *Psychological Review*, 57, 193–216.
- Soto, P. L., McDowell, J. J., & Dallery, J. (2006). Feedback functions, optimization, and relation of response rate to reinforcer rate. *Journal of the Experimental Analysis of Behavior*, 85, 57–71.
- Staddon, J. E. R. (1979). Operant behavior as adaptation to constraint. *Journal of Experimental Psychology: General*, 108, 48–67.
- Stubbs, D. A., & Pliskoff, S. S. (1969). Concurrent responding with fixed relative rate of reinforcement. *Journal of the Experimental Analysis of Behavior*, 12, 887–895.
- Tanno, T., & Sakagami, T. (2008). On the primacy of molecular processes in determining response rates under variable-ratio and variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, 89, 5–14.
- Tanno, T., Silberberg, A., & Sakagami, T. (2010). Concurrent VR VI schedules: Primacy of molar control of preference and molecular control of response rates. *Learning & Behavior*, 38, 382–393.
- Vaughan, W., Jr. (1981). Melioration, matching, and maximization. *Journal of the Experimental Analysis of Behavior*, 36, 141–149.
- Vaughan, W., Jr. (1987). Dissociation of value and response strength. *Journal of the Experimental Analysis of Behavior*, 48, 367–381.
- Vaughan, W., Jr., & Miller, H. L., Jr. (1984). Optimization versus response-strength accounts of behavior. *Journal of the Experimental Analysis of Behavior*, 42, 337–348.
- Wearden, J. H., & Clark, R. B. (1988). Interresponse-time reinforcement and behavior under aperiodic reinforcement schedules: A case study using computer modeling. *Journal of Experimental Psychology: Animal Behavior Processes*, 14, 200–211.
- Wearden, J. H., & Clark, R. B. (1989). Constraints on the process of interresponse-time reinforcement as the explanation of variable-interval performance. *Behavioural Processes*, 20, 151–175.
- Zeiler, M. D. (1986). Behavioral units: A historical introduction. In T. Thompson & M. D. Zeiler (Eds.), *Analysis and integration of behavioral units* (pp. 1–12). Hillsdale, NJ: Erlbaum.