

Temporal generalization gradients following an interdimensional discrimination protocol

Ana Catarina Vieira de Castro^{1,2,3}, Marco Vasconcelos^{3,4}, and Armando Machado³

¹Instituto de Investigação e Inovação em Saúde, Universidade do Porto, Porto, Portugal

²IBMC– Instituto de Biologia Molecular e Celular, Universidade do Porto, Porto, Portugal

³School of Psychology, University of Minho, Braga, Portugal

⁴Department of Zoology, University of Oxford, Oxford, UK

(Received 31 December 2014; accepted 2 August 2015; first published online 25 November 2015)

We investigated the effects of interdimensional discrimination training in the temporal generalization gradient. In a matching-to-sample task, pigeons learned to choose key S after a T -s houselight sample and key NS in the absence of the houselight sample. For one group of pigeons, $T = 20$ s; for another, $T = 10$ s. Subsequently, houselight duration was varied to obtain temporal generalization gradients. Results showed that (a) proportion S increased as houselight duration ranged from 0 s to T s and then remained high for houselight durations longer than T ; (b) the gradients were well described by negative-exponential functions; (c) these non-flat gradients were present from the beginning of testing, and; (d) the average gradients obtained with $T = 20$ s and $T = 10$ s overlapped when plotted in relative time. We conclude that temporal control does not require explicit discrimination training along the temporal dimension, and that temporal generalization gradients obtained with an interdimensional protocol show the scalar property of timing. We discuss how these findings challenge current models of timing.

Keywords: Interdimensional training; Pigeon; Scalar property; Temporal generalization; Timing models.

Guttman and Kalish's (1956) classic study inspired a large body of research on stimulus generalization. Researchers examined the generalization gradients obtained by different species including rats, pigeons, and humans, along stimulus dimensions such as hue, brightness, line orientation, auditory

frequency, and intensity, and following distinct training protocols (for reviews see Ghirlanda & Enquist, 2003; Honig & Urcuioli, 1981; Riley, 1968; Rilling, 1977; Terrace, 1966). One of the issues they examined was how the training protocol affects the shape of the generalization gradient. In

Correspondence should be addressed to Ana Catarina Vieira de Castro, IBMC– Institute for Molecular and Cell Biology (Laboratory Animal Science Group), Rua do Campo Alegre 823, 4150–171, Porto, Portugal. E-mail: castro.anacatarina@gmail.com

We are grateful to the members of the Animal Learning and Behavior Laboratory of University of Minho for their helpful comments on prior versions of this paper.

This work was supported by Portuguese Foundation for Science and Technology (FCT) [Grant number PTDC/MHC-PCN/3540/2012] to Armando Machado. Ana Catarina Vieira de Castro was supported by an FCT doctoral grant [Grant number SFRH/BD/43398/2008], by ERDF funds through the Operational Competitiveness Programme–COMPETE, and by National Funds through FCT under project FCOMP-01-0124-FEDER-029527 [Grant number PTDC/MHC ETI/4890/2012], and Marco Vasconcelos was supported by an FCT Investigator Grant [Grant number IF/01624/2013].

the present study, we extend this line of research to the dimension of stimulus duration.

Although training protocols may be classified in two broad groups—those involving differential reinforcement and those involving non-differential reinforcement—here we focus on the former because it has been the most widely used in studies of temporal control (e.g., Church & Deluty, 1977; Elsmore, 1971; Mellgren, Mays, & Haddad, 1983, Experiment 2; Russell & Kirkpatrick, 2007, Experiment 1; Siegel, 1986; Spetch & Cheng, 1998, Experiment 1; Stubbs, 1968; Vieira de Castro & Machado, 2012; Vieira de Castro, Machado, & Tomanari, 2013). In fact, to our knowledge no study using non-differential protocols has obtained non-flat temporal generalization gradients (Elsmore, 1971; Spetch & Cheng, 1998, Experiment 2).

In differential protocols, reinforcement is correlated with the stimulus values along the relevant dimension. To illustrate, food may be provided in the presence of a green light, the S^+ , but not in the presence of a red light, the S^- ; or food may be provided in the presence of a light, S^+ , but not in its absence, S^- . The correlation typically occasions some form of stimulus discrimination, with the subject responding significantly more in the presence of the S^+ than in the presence of the S^- .

In addition, on the basis of the relationship between the discriminative stimuli used during training and the stimulus dimension varied during subsequent generalization tests, differential protocols may be further classified as intradimensional, interdimensional, or extradimensional (Switalski, Lyons, & Thomas, 1966; Thomas, 1969). In intradimensional protocols the S^+ and the S^- differ in a single dimension (e.g., wavelength), which is also the dimension varied during the generalization tests. In the interdimensional protocol, the S^- is either the absence of the S^+ (e.g., the absence of a light) or a stimulus defined by dimensions orthogonal to the dimensions of S^+ (e.g., S^- is a tone and S^+ is a light). In testing, one of the S^+ dimensions (e.g., wavelength) is varied. Finally, in extradimensional protocols the S^+ and S^- may differ also in a single dimension (wavelength), but testing proceeds along a novel, previously untrained dimension (e.g., line-tilt).

In the temporal dimension only a few studies have examined the effect of training protocols on the shape of the generalization gradient. Moreover, the majority of these studies have used intradimensional protocols; only one used an interdimensional protocol, and, to our knowledge, none has used extradimensional protocols. To illustrate, Elsmore (1971) used an intradimensional protocol wherein pigeons received food for pecking at a key following a 21-s timeout, the S^+ , but they did not receive food for pecking at the same key following a 9-s timeout, the S^- . The generalization gradients obtained in this and similar studies were highly consistent in the interval between the two training samples: The number of responses increases from a low value at the S^- to a high value at the S^+ . Outside the training interval, the shape of the gradients was less clear, in particular on the side of the S^+ opposite the S^- . In some cases the entire gradient resembled a ramp function, low close to the S^- , increasing from the S^- to the S^+ , and then remaining constant past the S^+ . In some other cases, the gradients were monotonic, with the highest response strength occurring at one or more durations past the S^+ (peak shift). In yet other cases, the gradients were bitonic, with response strength decreasing significantly for durations past the S^+ .

The fact that temporal generalization gradients are flat following non-differential training (Elsmore, 1971; Spetch & Cheng, 1998, Experiment 2) and not flat following intradimensional discrimination training suggests that the acquisition of temporal control requires some form of differential reinforcement with respect to stimulus duration. However, it does not specify exactly what this form must be—for example, whether temporal control requires differential reinforcement with respect to two distinct stimulus durations, or whether it suffices to have reinforcement in the presence of a stimulus with a given duration and no reinforcement in the absence of that stimulus. The issue can be studied with the interdimensional protocol pioneered by Jenkins and Harrison (1960), whereby an animal is trained in a present-absent form of differential reinforcement and then tested for generalization along one of the dimensions of the trained stimulus.

Only one study used the interdimensional protocol to investigate temporal control. Russell and Kirkpatrick (2007, Experiment 1) exposed a group of rats to the following matching-to-sample task. After a dark intertrial interval (ITI), the houselight was either turned on for 20 s (sample S1) or it was not turned on (sample S2). Two levers were then inserted into the chamber. Pressing lever R1 was reinforced following S1 and pressing lever R2 was reinforced following S2. During subsequent generalization tests, the houselight duration varied from 2 to 80 s, and the proportion of R1 choices was recorded. That proportion increased from approximately .4 to .9 as the houselight duration increased from 2 s to 20 s, and then it decreased from .9 to .75 as the houselight duration increased from 20 s to 80 s. The bitonic and strongly asymmetric gradient was wider than the gradients produced by two other groups exposed to intradimensional protocols (same S1 = 20 s, but S2 = 10 s or S2 = 40 s).

Although Russell and Kirkpatrick's (2007) study suggests that temporal control may be acquired through interdimensional training, one feature of the study makes its interpretation difficult. Because the generalization gradient was based on all 20 test sessions, it is not known whether a gradient was present from the first test sessions on or was acquired during testing. The latter hypothesis is plausible because the contingencies during testing are similar to the contingencies of a typical temporal generalization task—in both cases choices of the S1 lever are reinforced following one sample duration and extinguished following shorter and longer samples—and temporal generalization tasks yield bitonic, bell-shaped gradients (Church & Gibbon, 1982; Weisman et al., 1999). More generally, the contingencies during testing define an intradimensional protocol because reinforcement of S1 depends on sample duration. Hence, if the bitonic gradient is not present during the first test sessions and develops only during testing, we may conclude that the interdimensional protocol is insufficient to yield temporal control. But if the bitonic gradient is present from the first test sessions on, we may conclude that the interdimensional protocol is sufficient to yield temporal control.

Moreover, because Russell and Kirkpatrick (2007) used only one houselight duration during training, it is not known whether gradients obtained with the interdimensional protocol conform to the scalar property of timing, arguably the most robust property of temporal control (see Church, 2003; Gallistel, 1990; Gibbon, 1977, 1991; Lejeune & Wearden, 2006). The scalar property means that generalization gradients obtained with different training durations are superimposed when plotted in a relative scale. Although several studies have found the scalar property with intradimensional protocols (e.g., Church & Deluty, 1977; Church & Gibbon, 1982; Vieira de Castro & Machado, 2013; Weisman et al., 1999), it is not known whether it holds with interdimensional protocols.

In the present study we aimed to clarify and extend Russell and Kirkpatrick's (2007) findings and thereby improve our understanding of the conditions required to obtain temporal control. Specifically, we asked (a) whether pigeons, like rats, produced bitonic temporal generalization gradients following interdimensional training; (b) whether these gradients were present in the first test sessions; and (c) whether the gradients showed the scalar property of timing. We divided pigeons into two groups: Group 20 and Group 10. Group 20 followed closely Russell and Kirkpatrick's (2007) procedure: The pigeons received food for choosing one key following a 20-s houselight sample and another key following the absence of the houselight. Next, during test sessions, we varied the houselight duration and obtained a choice temporal generalization gradient. Group 10 replicated Group 20, with the standard duration changed from 20 s to 10 s and the test durations adjusted accordingly. We used the data from the two experiments to check, for the first time, the scalar property with an interdimensional protocol.

The data obtained with the interdimensional protocol may also have implications for timing theories. First, they bring to the forefront the issue of how and why temporal control is acquired in the absence of differential reinforcement along the time dimension. Second, most timing models,

including the Scalar Expectancy Theory (e.g., Gibbon, 1991; Gibbon, Church, & Meck, 1984), the Behavioural Theory of Timing (e.g., Killeen & Fetterman, 1988), the Learning-to-Time model (Machado, Malheiro, & Erlhagen, 2009), and the Behaviour Economic Model (Jozefowicz, Staddon, & Cerutti, 2009) have been conceived to explain performance when an animal is exposed to some form of differential reinforcement with respect to time. It is not clear whether—and how well—the models can account for the effects of the interdimensional protocols in general and the shape of the resulting generalization gradients in particular. In the closing section of the paper we identify some of the challenges posed by the interdimensional protocol to any theory of timing and discuss them in the light of Scalar Expectancy Theory and the Learning-to-Time model.

EXPERIMENTAL STUDY

Method

Subjects

The subjects were 16 pigeons (*Columba livia*) from two different colonies. The 8 pigeons of Group 20 belonged to the Laboratory of Animal Learning and Behaviour of the University of Minho and had different experimental histories unrelated to temporal discrimination. The 8 pigeons of Group 10 belonged to the Institute of Psychology of the University of São Paulo and were experimentally naïve. All pigeons were housed in individual home cages with water and grit continuously available and were maintained at 85% of their free-feeding body weights throughout the experiment. Both colonies had a 13:11 h light/dark cycle, beginning at 8:00 am for the Group 20 colony and at 7:00 am for the Group 10 colony.

Apparatus

Six standard experimental chambers for pigeons from Med Associates® were used (4 with Group 20, 2 with Group 10). The front panel of each chamber contained three keys, 2.6 cm in diameter,

centred on the wall 24 cm above the floor and 6 cm apart, centre to centre. The keys could be illuminated from behind with red, green, blue, yellow, or white lights. Directly below the centre key and 6 cm above the floor was a hopper opening measuring 6 × 5 cm. The bird had access to mixed grain when the hopper was raised and illuminated with a white light. On the opposite wall of the chamber, a 7.5-W houselight provided general illumination. An outer box equipped with a ventilating fan enclosed the experimental chamber. Two personal computers, one running the **ABET II** software and another running **MED-PC IV** controlled all experimental events and recorded the data for Groups 20 and 10, respectively.

Procedure

Pretraining. Because of their previous experimental history, the pigeons from Group 20 advanced directly to Training. The pigeons from Group 10 learned to peck the keys through an autoshaping procedure. On each trial, the centre key was illuminated with a white light. When the illuminated key was pecked or when 8 s had elapsed, whichever occurred first, the keylight was turned off, and the food hopper was activated for 6 s. After the feeding period, a 60-s dark ITI followed. Each autoshaping session lasted until 60 reinforcers had been delivered.

Next, to train the pigeons to peck different keys and keylight colours, two sessions with the following structure were conducted. On each trial, the houselight was illuminated with white light and one of the three keys also was illuminated with one of five colours (red, green, blue, yellow, or white). When the illuminated key was pecked, both the keylight and the houselight were turned off, and the food hopper was activated for 6 s. After the feeding period, a 30-s dark ITI followed. Each session lasted until 60 reinforcements had been delivered.

Training. Training sessions comprised two types of trials, Standard (S) and Non-Standard (NS). The S trials began with the illumination of the houselight for 20 s (Group 20) or 10 s (Group 10), the standard duration. Next, the houselight was turned

off, and the two side keys were illuminated, one with a red light and one with a green light. The first response to the red or the green key turned both keylights off. If the response was correct, the food hopper was activated, and then the ITI started; if the response was incorrect, the ITI started immediately. We refer to the correct responses on the S and NS trials as S and NS responses, respectively. The mapping between the green and red keylight colours and the S and NS responses was counterbalanced across pigeons. The ITI was spent in total darkness, and it lasted a minimum of 30 s plus a variable-time averaging 20 s for Group 20, or a minimum of 15 s plus a variable-time of 10 s for Group 10.

The NS trials did not have any distinctive cue signalling their onset. Instead, after a variable-time schedule of 40 s (Group 20) or 20 s (Group 10) timed out, the two side keys were illuminated with red and green lights. Again, the first response to one of the keys turned both keylights off; if the response was correct, the food hopper was raised, and then the ITI started; if the response was incorrect, the ITI started immediately. To minimize extra-session feeding, hopper duration varied across birds from 3 s to 9 s.

The variable-time values for the ITI and the NS trials were generated from Fleshler and Hoffman's (1962) distributions, with means of 20 s and 40 s, respectively, for Group 20, and with means of 10 s and 20 s, respectively, for Group 10. For each variable-time, 60 intervals were generated. These values were the same for the 8 pigeons of each group, and their order of presentation was randomized every session.

A correction procedure was employed on both trial types. Thus, during the first 5 training sessions, if the pigeons made an incorrect response, then on the following trial only the correct key was presented after the sample. After Session 6, two errors were allowed before the correct key was presented alone. Correction trials were not included in the data analyses.

Each training session comprised 60 trials, 30 S and 30 NS, excluding correction trials. The position of the green and red keylights was counterbalanced such that, on half of the S trials and on half

of the NS trials, green appeared on the left and red appeared on the right, and on the remaining trials green appeared on the right and red appeared on the left. The criterion to advance to the test phase was a minimum of 15 training sessions and at least 85% of correct choices on both S and NS trials for five consecutive sessions. Training lasted 16 sessions on average (range = 15–20) for Group 20. For Group 10, when the learning criterion was about to be met, an unplanned two-week interruption of the experiment brought the performance of some pigeons close to chance, which increased the number of training sessions (average = 35; range = 30–38).

Testing. A test session comprised 28 regular training trials (14 S and 14 NS) and 32 unreinforced test trials, with 8 different houselight durations, each presented 4 times. The test durations were 3.0, 8.1, 10.9, 14.8, 27.0, 36.5, 49.2, and 66.4 s for Group 20, and 3.0, 4.1, 5.5, 7.4, 13.5, 18.2, 24.6, and 33.2 s for Group 10, forming a logarithmic series except at the shortest duration, which served as an anchor close to 0. The session always began with 4 training trials, 2 S and 2 NS, in random order. The remaining training and test trials were intermixed randomly. To compensate for the large number of unreinforced trials, hopper duration was increased during testing in order to maintain constant the amount of food consumed per session and minimize extra-session feeding (this varied across birds between 3.5 s and 13 s).

If during any test session the percentage of correct choices following one of the two training samples fell below 80%, the bird was returned to Training for at least 2 sessions and until its percentage of correct choices exceeded 85%. Testing continued for 20 sessions for Group 20 and for 40 sessions for Group 10. We extended the number of test sessions for Group 10 because, given the results described below, we wanted to determine whether the generalization gradient changed with more extensive testing. The number of sessions (test + intercalated retraining when necessary) averaged 24 for Group 20 (range = 20–33) and 45 for Group 10 (range = 40–56).

Results

Training

All pigeons learned the task. The learning criterion (at least 85% correct choices on both S and NS trials for five consecutive sessions) was met by Group 20 by the 12th training session on average (range = 9–20) and by Group 10 by the 14th training session on average (range = 8–22). The proportion of correct choices during the last five sessions of training averaged .96 on the S trials (range = .91–.99) and .97 on the NS trials (range = .95–.99) for Group 20, and .97 on the S trials (range = .93–.99) and .98 on the NS trials (range = .96–.99) for Group 10.

Testing

Figure 1 shows the average temporal generalization gradients obtained during all test sessions. The proportion of trials in which the pigeons chose the S key is plotted as a function of the houselight duration. The pattern of responding was consistent across pigeons (see supplemental online material). For Group 20 (left panel), when no sample was presented on the NS trials, proportion S was approximately zero (see empty circle at 0 s); when the houselight was presented, proportion S was close to chance at 3 s, and above chance at all other durations. In short, proportion S increased monotonically as the houselight duration ranged from 0 s to 20 s and then remained high for longer durations.

For Group 10 (right panel), when the houselight was not present, proportion S was approximately zero; when the houselight was present, proportion S was lowest, although still above chance, for the shortest duration of 3 s, increased to a high value, and then remained high for longer durations. Similar to Group 20, proportion S increased monotonically as the houselight duration ranged from 0 to 7.4 s and then remained high for longer samples.

The individual gradients were well described by a negative exponential function with equation $f(t) = a(1 - e^{-t/b})$, where t is the houselight duration, a is the asymptotic proportion of S choices, and b is the time constant that determines inversely the rate of approach to the asymptote. The variance

accounted for, r^2 , averaged .99 for both groups. The solid lines of Figure 1 show that the average of the individual best-fitting curves fitted the average data well.

Concerning the parameters, the average of the a parameter for the two groups, \bar{a}_{20} and \bar{a}_{10} , equalled .97 (ranges = .92–.99 for a_{20} , and .94–.97 for a_{10} ; see supplemental material, Figures S1 and S2). These results show that asymptotic performance did not vary between groups. Concerning parameter b , $\bar{b}_{20} = 4.76$ s (range = 2.98–6.22 s), and $\bar{b}_{10} = 2.01$ s (range = 1.24–2.74 s), which means that proportion S reached half of its asymptote ($a/2$) at houselight durations between 2.1 s and 4.3 s in Group 20 and at houselight durations between 0.9 s and 1.9 s in Group 10 [$t = \ln(2) \times b$]. The two sets of values for Groups 20 and 10 are approximately in a 2:1 ratio, as the scalar property predicted. We return to this finding below.

The first goal of the present study was to investigate whether pigeons produced generalization gradients similar to those obtained with rats by Russell and Kirkpatrick (2007). Figure 2 compares the average results from the two studies. The gradients are plotted relative to the houselight standard duration, 20 s in Russell and Kirkpatrick (2007) and Group 20, and 10 s in Group 10. Because the maximum value of proportion S was higher in the present study than in Russell and Kirkpatrick's (2007) study, we normalized the average data of each curve by dividing each value along the curve by its maximum. The normalization better isolates the gradient shape.

The figure shows that the average gradients from the two studies were roughly similar for durations shorter than the standard—proportion S increased with houselight duration and reached a maximum at or near the standard duration—but this increase was faster in the present study. For durations longer than the standard, the gradients from the two studies differed—proportion S remained high in the present study but decreased in Russell and Kirkpatrick's (2007) study. At the individual level, only one pigeon from the present experiment (see supplemental material, Figure S2, P17) showed a gradient similar to the average




Figure 1. Average generalization gradients obtained from all test sessions of Group 20 (left panel) and Group 10 (right panel). Proportion of S choices is plotted as a function of the houselight duration. The empty circles correspond to the NS trials (no houselight). The solid lines through the average data show the average of the individual best-fitting negative exponential curves. The vertical bars show the SEM.

gradient reported by Russell and Kirkpatrick (2007).

Figure 2 also shows that the two gradients from the present study overlapped when plotted in a relative scale. This result is consistent with the previously reported finding that the a parameter values did not differ between the two groups ($\bar{a}_{20} = \bar{a}_{10}$), whereas the b parameter values for Group 20 were roughly twice the values for Group 10 ($\bar{b}_{20} \approx 2\bar{b}_{10}$). By inverting a permutation test (Ernst, 2004), we constructed the 95% confidence interval (95% CI) for the ratio $\bar{b}_{20}/\bar{b}_{10}$, and obtained [1.73, 3.06]; the interval excludes 1, but includes 2. We conclude that the observed b values are consistent with the 2:1 ratio expected from the scalar property.




Figure 2. Average generalization gradients obtained by Group 20 (filled circles), Group 10 (empty circles), and Russell and Kirkpatrick (2007) (empty triangles). Both axes were normalized.

Another goal of the present study was to assess whether non-flat generalization gradients were present from the first test sessions on or, on the contrary, emerged during testing. Figure 3 shows the average temporal generalization gradients from the first test session. The left panel displays the average gradients in absolute time; the right panel displays the average gradients in relative time. The symbols show the average data and the lines show the average of the individual (negative exponential) fits.

For both groups, the average data showed the same shape: Proportion S increased from 0 s to the standard duration and then remained approximately constant. The individual gradients were more variable, but for 6 of the 8 pigeons in each group the individual gradient was similar to the average gradient. For the other 2 pigeons in each group, the individual gradient was flat for durations equal to or greater than 3 s.

To further examine the asymmetry of the gradients on the first test session, we compared the total number of S choices following the three sample durations below the standard with the total number of S choices following the three sample durations above the standard (e.g., for Group 20, the 14.8 s, 10.9 s, and 8.1 s samples vs. the 27.0 s, 36.5 s, and 49.2 s samples). Although the three samples on each side of the standard are equidistant from the standard in log units (suggesting equal discriminability from the standard), if the gradients




Figure 3. Average generalization gradients obtained by Groups 20 and 10 during the first test session. The left panel shows the gradients as a function of houselight duration. The right panel shows the gradients in relative time. The curves show the average of the negative exponential functions fitted to the individual data.

are asymmetric, then the number of S choices on each side of the standard should differ. Moreover, because in the two groups these samples preserved the ratio with the standard, there should be no differences between the two groups in the degree of asymmetry. A mixed ANOVA with Group (20 vs 10) as the between-subjects factor and sample (below vs above the standard) as the repeated factor yielded a significant effect of sample, $F(1, 14) = 16.5$, $MSE = 2.07$, $p = .001$, but not of group, $F(1, 14) = 3.66$, $MSE = 3.76$, $p = .076$, or Group \times Sample interaction, $F(1, 14) = 1.22$, $MSE = 2.07$, $p = .29$. On average, the probability of choosing S following the three samples to the right of the standard exceeded the probability of choosing S following the three samples to the left of the standard by 0.17 (95% CI[0.08; 0.26]).

The right panel of Figure 3 shows that the average data from the first session did not superimpose as strongly as the average data from all test sessions (cf. Figure 2). To better understand this result, we compared the parameters of the negative exponential fits. For one pigeon of Group 10 (P14), the fit failed because the pigeon chose S on all trials except the 0-s NS trials. For the remaining pigeons, and despite the fact that each data point was based on four test trials only, the function fit the data reasonably well (for Group 20, average $r^2 = .83$, range = .73–.97; for Group 10, $r^2 = .84$, range = .75–.98). The a parameter averaged $\bar{a}_{20} = .94$ and $\bar{a}_{10} = .97$, and the 95% CI for the difference

between the two averages, $[-0.044, +0.045]$, included 0. The b parameter averaged $\bar{b}_{20} = 7.8$ s and $\bar{b}_{10} = 2.5$ s, and the 95% CI for their ratio, $[1.8, 5.41]$, excluded 1 but included 2. We conclude that most gradients from the first test session were asymmetric around the standard and followed approximately a negative exponential curve; the gradients were consistent with the scalar property.

Figure 4 shows the generalization gradients obtained in different blocks of 5 test sessions. The top panels show that the average proportion of S choices did not change appreciably during testing for both Group 20 (left) and Group 10 (right). For Group 20, the negative exponential function fitted the data well (r^2 averaged .97 and .96 in the two blocks of sessions) with no reliable differences in the parameter averages ($\bar{a}_{20} = .96$ and .97; $\bar{b}_{20} = 6.4$ s and 4.1 s, for the two blocks). In particular, the 95% CI for the ratio $\bar{b}_{20, \text{First}}/\bar{b}_{20, \text{Last}}$, $[0.96, 2.47]$, included 1. For Group 10, the negative exponential fit was equally good ($r^2 = .96$, .96, and .94 in the three blocks) and again there were no reliable differences in the parameter averages ($\bar{a}_{10} = .95$, .95, and .98 and $\bar{b}_{10} = 2.1$, 1.9, and 2.6 s across blocks). The 95% CI for the ratio $\bar{b}_{10, \text{First}}/\bar{b}_{10, \text{Last}}$ in the first and last blocks, $[0.49, 1.38]$, also included 1. We conclude that the negative exponential shape was preserved across testing, and that the best-fitting parameters did not change significantly across testing.

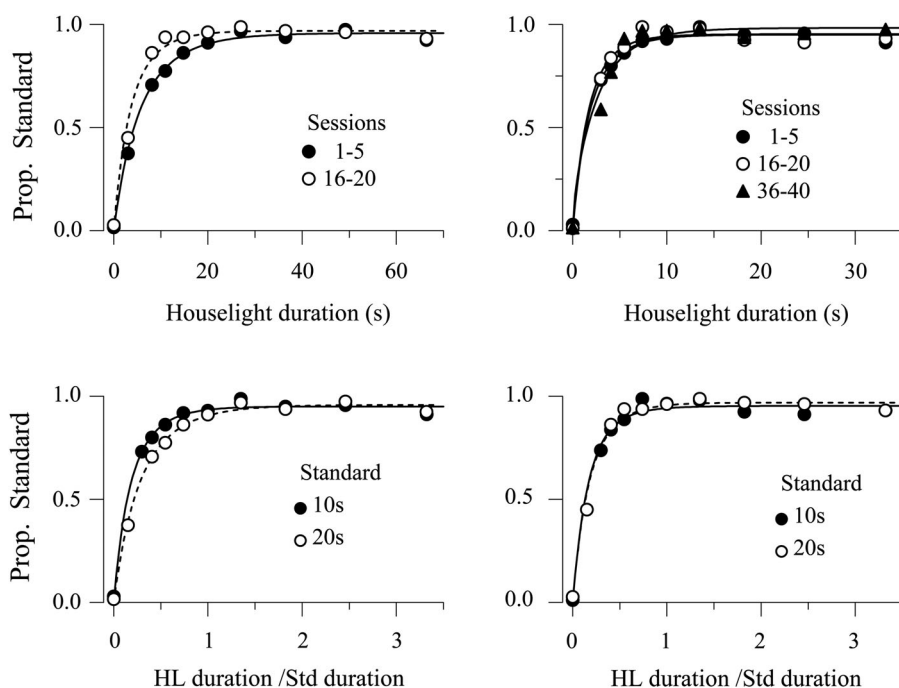


Figure 4. Average generalization gradients obtained in 5-session blocks during testing. Top left: Gradients from Group 20 obtained during the first and last 5 test sessions. Top right: Gradients from Group 10 obtained during the first, middle, and last 5 test sessions. Bottom left: Gradients from Groups 20 and 10 during the first 5 test sessions plotted in relative time. Bottom right: Gradients from Groups 20 and 10 during Sessions 16–20 plotted in relative time. The curves show the average of the negative exponential functions fitted to the individual data.

The bottom panels examine superimposition across testing. Both the average data and the average of the individual fitted functions show that, during corresponding session blocks, the gradients for the two groups superimposed to large extent. Because parameter a remained approximately constant, the scalar property predicts that $\bar{b}_{20}/\bar{b}_{10} \approx 2$. For Sessions 1–5 (left panel), $\bar{b}_{20} = 6.4$, $\bar{b}_{10} = 2.1$, and the 95% CI for the ratio $\bar{b}_{20}/\bar{b}_{10}$ equalled [2.01, 4.66]; for Sessions 16–20 (right panel), $\bar{b}_{20} = 4.1$, $\bar{b}_{10} = 1.9$, and the 95% CI for the ratio equalled [1.36, 3.22]. Since the two intervals excluded 1, we conclude that the gradients from the two groups differed. And though the first interval excluded 2, given that (a) the second interval included 2, (b) the intervals from the first test session and from all test sessions also included 2, and (c) there were no systematic trends during testing, we conclude that the gradients were consistent with the scalar property.

During the test sessions, to maintain overall reinforcement rate and minimize extra-session feeding, different pigeons experienced different reinforcement durations. These differences could explain the differences in the time constant parameter b . However, scatterplots of b_{20} (based on all test sessions) and reinforcement duration showed no obvious relation, and linear regression analyses for Groups 20 and 10 yielded no significant slopes (for β_{20} , 95% CI[−0.41, 0.34]; for β_{10} , 95% CI[−0.38, 0.08]).

Discussion

The major question examined in the present study was whether training to discriminate the presence (S^+) from the absence (S^-) of a stimulus is sufficient to establish control by its duration. In the only study that used an interdimensional protocol, Russell and Kirkpatrick (2007) found a non-flat,

bitonic temporal generalization gradient. Although this result may suggest that temporal control was acquired during training, the reported gradient was based on all test sessions, and therefore we do not know whether it was present from the first test sessions on, or developed only during testing.

In addition, a generalization gradient obtained with a single standard duration could be due to processes other than timing. Consider a simple alternative process. Let $f(t)$ be the probability of choosing the S key following a t -s houselight duration. Then $f(0)$ is the probability of choosing the S key when no houselight is present—that is, on NS trials. Because on NS trials the subject faces the comparison keys immediately after the ITI, it is reasonable to assume that during training the ITI becomes strongly associated with NS pecks; hence, we assume that $f(0) \approx 0$. This assumption is consistent with the fact that the proportion of S choices on NS trials was always close to 0 (see Figure 1). When the houselight is illuminated on S trials, we assume that the tendency to choose the S key increases, or, equivalently, that the tendency to choose the NS key decreases continuously, perhaps due to forgetting of the ITI with the passage of time, interference by the houselight, or both. If $f(t)$ increases to an asymptote (a) slightly below 1 and at a rate proportional to how much it can increase, then the following differential equation holds,

$$\frac{d}{dt}f(t) = b[a - f(t)].$$

Its solution is the negative exponential function $f(t) = a(1 - e^{-t/b})$ that fitted the individual data well. Strictly speaking, the foregoing account involves no timing, because the critical temporal variable, the standard houselight duration, does not function as a discriminative cue. The forgetting or interference process occurs in time, and therefore the probability $f(t)$ varies with time, but no specific timing process is at play.

One way to ascertain whether timing is involved is to vary the relevant standard duration (e.g., change the houselight duration from 20 s to 10 s) and see whether the two generalization gradients

superimpose when plotted in relative time, the scalar property. If the gradients express only the (exponential) decay of the ITI influence on NS choice, then the gradients will not superimpose.

In what follows, we discuss two main issues: the shape of the generalization gradient in its relation to the training protocol, and the implications of our findings to theories and models of timing.

Generalization gradients following interdimensional training

The data from all test sessions showed that all but one pigeon produced negative-exponential-like temporal generalization gradients—proportion S increased as the houselight duration ranged from 0 s to the standard duration, and then remained high for longer samples (see Figure 1; also supplemental material, Figures S1 and S2). Hence, when the data from all pigeons and test sessions were averaged, the results of the present study partially replicated Russell and Kirkpatrick (2007; cf. Figure 2). For durations shorter than the standard, our gradients reproduced theirs, but for durations longer than the standard, they did not, for whereas in Russell and Kirkpatrick's study proportion S decreased past the standard, in the present study it remained high.

Our results also showed that, for most pigeons, non-flat, negative-exponential-like gradients were present in the first test session (Figure 3) and retained their shape for a large number of sessions (Figure 4). These gradients clearly reveal temporal control. However, two pigeons from each group showed flat gradients at the beginning of testing—a result that suggests no control from the houselight duration. The reasons for the individual differences remain unclear.

A two-parameter negative exponential function described well most generalization gradients. The asymptote parameter (a) was always close to 1, and its average did not change significantly between groups. The time constant parameter (b) was larger for Group 20 than for Group 10. Moreover, the confidence intervals for the ratio $\bar{b}_{20}/\bar{b}_{10}$ on the first test session and on all test sessions failed to include 1, but included 2, the value consistent with the scalar property. When we

combine these results with (a) the degree of superimposition revealed in Figures 2, 3 (right panel), and 4 (bottom panels)—a degree equal to that in most reports of the scalar property; (b) the fact that the average b values did not change systematically from the first 5 to the last 5 sessions of each group (Figure 4, top panels); and (c) the ubiquitous presence of Weber's law in timing studies, we conclude that the generalization gradients were consistent with scalar invariance: Durations in the same ratio relative to the standard yielded approximately the same proportion of S choices. Hence, to the list of timing procedures in which the scalar property is observed, a list that includes FI schedules, the peak procedure, mixed FI-FI schedules, DRL schedules, temporal generalization, and temporal bisection tasks (e.g., Church, 2003; Church & Gibbon, 1982; Gibbon, 1977; Lejeune & Wearden, 2006; Richelle & Lejeune, 1980; Vieira de Castro et al., 2013), the present study adds the interdimensional protocol.

The stability of the generalization gradients across test sessions may be surprising because these sessions included a relatively large number of extinction trials (640 in Group 20 and 1280 in Group 10) following houselight durations around the standard. In particular, we expected that the almost exclusive initial preference for the S key following houselight durations longer than the standard would decrease with successive extinction trials. Such decrease could account for the bitonic rat data reported by Russell and Kirkpatrick (2007; see Figure 2). In fact, the reinforcement contingencies for the S response during testing—reinforcement following the standard duration, extinction following all other durations—define a temporal generalization procedure (Church & Gibbon, 1982; Weisman et al., 1999), a form of intradimensional protocol, and this procedure yields bitonic gradients. However, the preference for the S key following the longest samples remained strong during testing, and it was observed in all but one bird.

The shape of the generalization gradient is similar to the response rate gradient observed when an animal that is exposed to a fixed-interval (FI) schedule, say an FI 20 s, is transferred to an

equivalent peak procedure schedule (Catania, 1970) in which 20 s food trials alternate with significantly longer empty trials. On first exposure to the empty trials, response rate increases from 0 to 20 s, and then it remains high until the end of the trial (e.g., Balci et al., 2009; Kaiser, 2008; Monteiro & Machado, 2008). However, as the animal experiences extinction past the food time, response rate after the food time decreases, and the typical bitonic gradient around the food time emerges. Why, then, did preference for the S key following the longest samples not decrease across test sessions in the present study? Perhaps the reason is that whereas on a typical empty trial of a peak procedure the animal produces a large number of unreinforced responses, in the current study only one (choice) unreinforced response was emitted per test trial. The number of extinction responses, rather than the number of extinction trials, may be the critical difference. Support for this idea can be found in Rescorla's (2001) thorough analysis of experimental extinction, which shows compelling evidence for a positive relationship between the number of non-reinforced responses and the magnitude of extinction.

Naturally, our and Russell and Kirkpatrick's (2007) experiments offer an interesting opportunity to discuss cross-species similarities and differences. Nonetheless, considering the differences in experimental protocols, any strong inference as to species identity being responsible for the observed results would be unjustified. Even though the majority of the gradients from the present study were negative-exponential-like and, thus, different from Russell and Kirkpatrick's, one of our pigeons displayed a bitonic gradient similar to theirs. This finding shows that some pigeons may, like rats, display bitonic gradients when exposed to an interdimensional protocol. Because Russell and Kirkpatrick reported only average data, we do not know how their rats behaved individually. It is conceivable that in their study some rats also displayed negative-exponential-like gradients and others showed bitonic gradients.

One methodological feature may question the interpretation of our findings. Because the two groups of pigeons were studied in different

laboratories (i.e., equipment, light–dark cycles) and had different experimental histories, variables other than the difference in the standard stimulus duration (20 s vs. 10 s) could explain the differences in the generalization gradients of Groups 20 and 10. Although we cannot rule out the hypothesis, we find it unlikely, because the gradients from the two groups exhibited the scalar property. Given the ubiquity of that property when temporal learning occurs, it seems more plausible to assign the results to different standard durations than to, say, different light–dark cycles. The latter hypothesis is also less parsimonious because we would be assuming, without evidence, not only that different light–dark cycles engender different gradients in absolute time, but also that they engender gradients that overlap in relative time. Moreover, the results from our groups were similar to those reported by Russell and Kirkpatrick (2007) using rats with different experimental histories and environmental conditions.

Implications for models and theories of timing

The generalization gradients obtained with the interdimensional protocol challenge our theories and models of timing. Perhaps the most important challenge is to explain why timing occurs at all if the contingencies of reinforcement of the interdimensional protocol do not seem to require it. Obviously, any stimulus to be detected and attended to needs to have a duration greater than the subject's absolute threshold, but that requirement does not entail the timing of the *specific* duration of the stimulus, particularly when that duration, as the houselight in the present study, is significantly greater than the threshold.

The interdimensional protocol is analogous to the present/absent sample matching task in that both involve training to discriminate the presence from the absence of a stimulus. Typically, in present/absent sample matching tasks, divergent retention functions are observed with increasing retention intervals—matching accuracy on sample-absent trials remains at a high level, whereas on sample-present trials it declines rapidly (e.g., Grant, 1991; Sherburne & Zentall, 1993a, 1993b). To explain this finding, researchers

(e.g., Colwill, 1984; Grant, 1991; Sherburne & Zentall, 1993a, 1993b; Wilson & Boakes, 1985; but see Dougherty & Wixted, 1996; Wixted, 1993) have hypothesized that animals use a single-code/default strategy, whereby on sample-present trials animals attend to the sample, store it in memory, but may then lose the memory during the retention interval. This strategy requires attention to the sample, but not to the sample duration.

Perhaps more subtle contingencies embedded in the interdimensional protocol lead to timing the sample. The choice response may be the last element of a sequence of behaviours that includes waiting or engaging in competing responses, orienting, approaching, and finally pecking a choice key. Failing to time the sample may mean approaching the keys too early, which in turn may prevent the animal from engaging in other responses and collecting alternative reinforcers, or approaching the keys too late, which delays food. Hence, during training the animal may learn to maximize reinforcement by temporally regulating its approach and preparation to choose according to the sample duration. The expression of that regulation could be the timescale-invariant generalization gradients.

Another challenge concerns the shape of the generalization gradients, in particular the fact that most gradients raised quickly from about 0 to about 1 and were well described by a negative exponential function. The rapid increase yielded broad generalization gradients which may suggest that the more subtle reinforcement contingencies mentioned above may be rather weak, variable, or unsystematic.

Most models of timing have been designed to account for data from intradimensional protocols. Hence, they may need to be modified to meet the challenges posed by the data from interdimensional protocols. Consider the highly influential Scalar Expectancy Theory or SET model (e.g., Church, 2003; Gibbon, 1991, 1981). To account for our findings, SET would need to postulate a single memory store for the S stimulus (for the sample on NS trials has zero “duration”) and a single threshold. The account could run as follows. On

the NS trials, the animal learns to choose the NS key, and no timing is involved. On the S trials, the animal forms a memory of the (scalar) houselight standard duration S . To choose which key to peck on trials with the houselight on, the animal compares the current duration, t , with a sample, T , extracted from its memory store and decides to peck the S key when the ratio $|T - t|/T$ falls below a threshold.

This version of SET may account for the shape of the generalization gradient in the interdimensional protocol, but it needs additional work to provide a principled account of the data—in particular, to specify the conditions required to activate the threshold and the learning processes engaged on the NS and S trials. Until then the model remains limited to a steady-state account and somewhat ad-hoc, for it invokes processes needed to fit the data but does not explain their use (see Machado et al., 2009, for a discussion of this issue).

In what follows, we consider a second model, LeT (Machado, 1997; Machado et al., 2009), that attempts to explain acquisition and steady-state temporal performance. We describe a first and only partly successful attempt to extend it to the interdimensional case. According to LeT—see Figure 5, left panel—timing involves three processes: the serial *activation* of a set of behavioural states (the top circles), the coupling or *association* of each state with the choice responses (the middle lines connecting the states with the choice responses), and the *decision rule* determining which choice is made (the two bottom arrows). We summarize each process below (for more details, see Machado et al., 2009).

Activation. At the onset of a sample (e.g., the houselight), only the first state is active; after some time, say, τ s, the first state becomes inactive, and the second state becomes active; the second state remains active for another τ s, and the process continues. The activation of successive states proceeds with speed λ ($= 1/\tau$) states per second, with λ sampled at the onset of each sample stimulus from a Gaussian distribution with mean μ_λ and standard deviation σ_λ .

Decision rule. At the end of the sample, one state is active, say, state $n \geq 1$, and that state is linked with the two choice responses. The strength of each link ranges between 0 and 1. If we let $WS(n)$ and $WNS(n)$ represent the strength of the link between state n and the S and NS keys, respectively, the probability of choosing S equals the relative strength $WS(n)/[WS(n) + WNS(n)]$.

Association. If the animal chooses S, say, and receives food, then $WS(n)$ increases according to a Bush and Mosteller (1955) linear operator model, $\Delta WS(n) = \beta \times (1 - WS(n))$, with $0 < \beta < 1$ a reinforcement parameter, and $WNS(n)$ decreases also according to a linear operator model, $\Delta WNS(n) = -\beta \times WNS(n)$. Similarly, if the animal chooses S and does not receive food, $WS(n)$ decreases by the amount $\Delta WS(n) = -\alpha \times WS(n)$, with $0 < \alpha < 1$ an extinction parameter, and $WNS(n)$ increases by the amount $\Delta WNS(n) = \alpha \times [1 - WNS(n)]$. To simplify the model, we assume that the reinforcement and extinction parameters are equal ($\beta = \alpha$).

The LeT model has been applied successfully to the intradimensional protocol (see Machado et al., 2009 for a review). To extend it to the interdimensional protocol, we make three additional assumptions. First, we assume that the stimulus conditions preceding the onset of each trial—in the present experiments, the ITI—can become a discriminative stimulus and bias choice. In fact, on the NS trials, the ITI may be the most important discriminative stimulus, strongly biasing choice towards the NS key. Second, the discriminative function of the ITI is learned through reinforcement and extinction. And, third, when the houselight is turned on during the S trials, the influence of the ITI on choice decreases exponentially with time.

The right panel of Figure 5 shows the new model architecture. The square represents the ITI, and the two lines linking the square to the choice responses represent the discriminative function of the ITI. We denote the strength of the links between the ITI and the choice responses by the variables ZS and ZNS . Similar to the variables

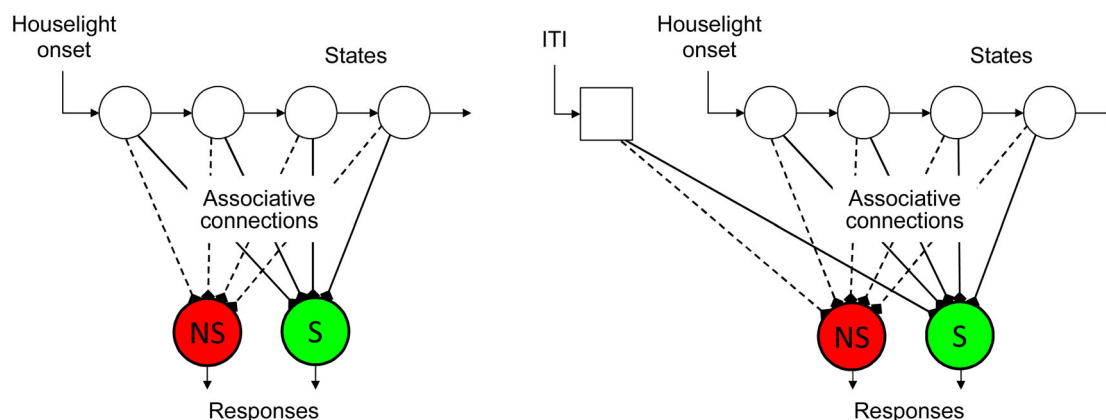


Figure 5. The left panel shows the architecture of the Learning-to-Time (LeT) model for intradimensional protocols, with a series of states, two choice responses, and the associative links connecting the states to the responses. The right panel shows the new architecture for the interdimensional discrimination protocol. To view this figure in colour, please visit the online version of this Journal.

$WS(n)$ and $WNS(n)$, the values of ZS and ZNS affect choice and change with training.

The model works as follows: At the beginning of each trial (i.e., at the end of the ITI), the square is active. If the trial is an NS trial, then the discriminative function of the ITI is in effect, and the animal chooses the S key with probability $ZS/(ZS + ZNS)$. Then, depending on the choice effectively made and the trial outcome—reinforcement if NS, extinction if S—the values of ZS and ZNS change according to the same Bush–Mosteller operators described above for $WS(n)$ and $WNS(n)$. If the trial is an S trial, then the serial activation of the states (top circles in Figure 5) proceeds for the duration of the houselight while, at the same time, the discriminative function of the ITI is increasingly likely to cease (i.e., the square becomes inactive) according to an exponentially distributed random interval with mean θ s.

At the end of the S trials, choice depends on whether or not the ITI is still remembered (the square remains active). In the former case, the probability of choosing S equals $[WS(n) + ZS]/[WS(n) + ZS + WNS(n) + ZNS]$. If the ITI is not remembered, the probability of choosing S reduces to $WS(n)/[WS(n) + WNS(n)]$.

We ran the extended LeT model for the interdimensional protocols of both groups. Figure 6 shows the results for the $S = 20$ s (left) and $S = 10$ s

(right) cases. The top two panels show the link strengths at the end of training. The link strengths of the ITI (ZS and ZNS) are conveniently plotted at state $n = 0$; the link strengths of the states activated by the houselight $WS(n)$ and $WNS(n)$ are plotted against $n \geq 1$. The smooth lines show the model average predictions for a large number of simulations; the symbols show one simulation run.

At the end of training, the ITI is strongly associated with the NS choice (triangle at $n = 0$, $ZNS \approx 1.0$) and weakly with the S choice (circle at $n = 0$, $ZS \approx 0$). The link strength of the remaining states with the S choice (circles at $n \geq 1$) show an excitatory generalization gradient centred at the states most likely to be active at the end of the S signal (left panel: $n \approx 20$ for a 20-s houselight; right panel: $n \approx 10$ for a 10-s houselight). Moreover, the link strengths with the NS choice (triangles at $n \geq 1$) show an inhibitory-like generalization gradient centred also at the states most likely to be active at the end of the S signal. In both cases, states unlikely to be active at the end of the S signal retain their initial strength.

The middle panels show the generalization gradients. The symbols and the curves increase from a value close to 0 to a value close to 1 as the houselight duration increases from 0 (NS trials) to S. However, as the houselight test duration continues to increase past the S value, the gradients

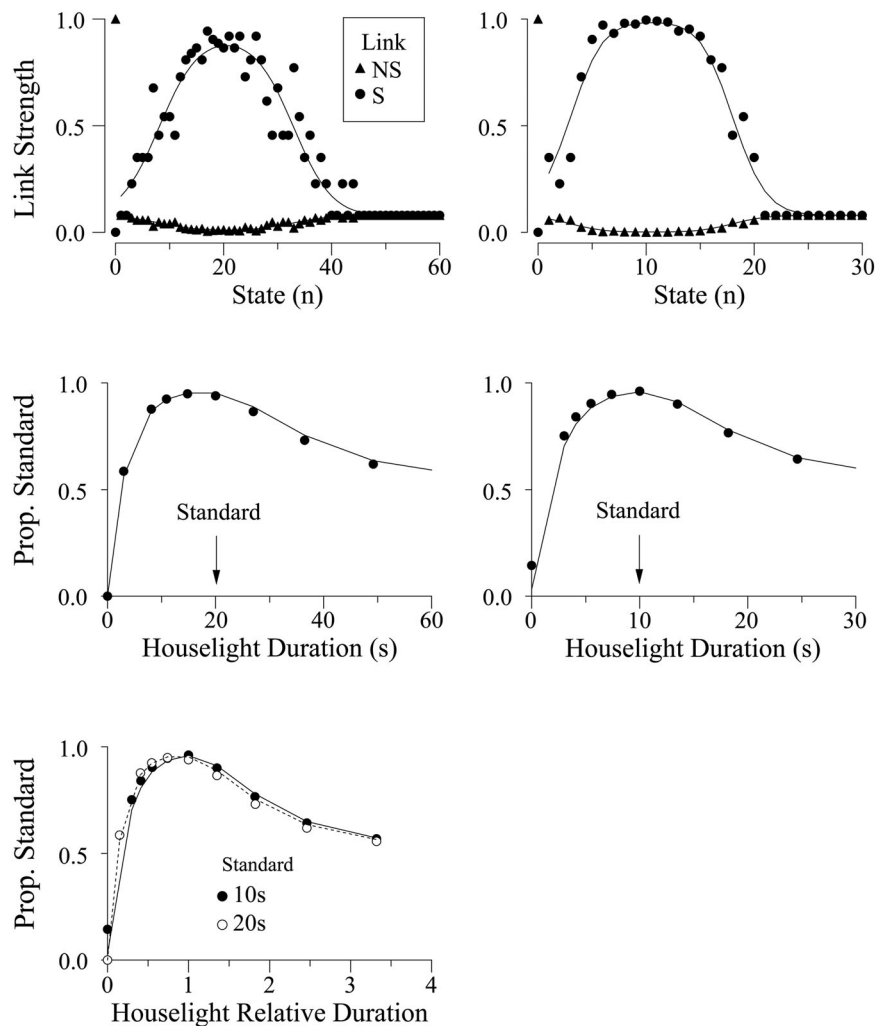


Figure 6. Results from the simulations of the extended LeT model, each comprising 500 trials. The parameter values used were $\mu_\lambda = 1$, $\sigma_\lambda = 0.4$, $\beta = \alpha = 0.16$; all links were initially set at 0.08. These parameter values are similar to those used in previous studies with the intradimensional protocol (e.g., Vieira de Castro & Machado, 2012). Parameter θ was set at 3 s. Top panels: Strength of the links at the end of training with a 20-s houselight (left panel) and a 10-s houselight (right panel). Middle panels: Generalization gradients following training with a 20-s houselight (left panel) and a 10-s houselight (right panel). Bottom panel: Simulated gradients plotted in a relative scale.

decrease and approach 0.5. The decrease in the gradients past the S value stems from the fact that after houselight test durations significantly greater than S, the active states at the moment of choice were rarely or never active during training and therefore retained their initial link strengths. These initial strengths are unbiased and predict indifference. Finally, as the bottom panel shows, the model

predicts gradients that overlap considerably when the houselight test duration is scaled by S (the scalar property).

In summary, the extended LeT model reproduces some but not all features of the data from an interdimensional protocol: (a) like pigeons and rats, the model learns the contingencies, to choose the NS alternative on most NS trials and

the S alternative on most S trials; (b) it produces generalization gradients that, in the range from 0 to S, approximate those produced by pigeons and rats; and (c) similar to the pigeon gradients, the model gradients overlap to a considerable extent. However, for durations greater than S, the model does not reproduce the data, because its gradients return to indifference, whereas the pigeon gradients remain close to 1 and the rat gradients remain significantly above 0.5.

The foregoing analysis raises a key issue: what processes are responsible for the asymmetric gradients in the interdimensional protocol? The answer to this question may involve categorical decision rules such as, "If the houselight duration relative to the standard duration is greater than a threshold, choose S; otherwise, choose NS." This rule is similar to the single-code/default hypothesis and to the ad-hoc SET version discussed above. If this proves to be the case, then the challenge to theories of timing will be to explain how learning to use, activate, or set the threshold occurs during training. Alternatively, the answer may involve a more complex model architecture—for example, one with an extra node representing the non-temporal elements of the houselight stimulus (similar to the square representing the ITI in Figure 5, right panel). Along with such non-temporal elements of the houselight, the ITI node itself may need to include temporal elements, as they seem to have a measurable impact on temporal discriminations, at least in Pavlovian preparations (Bouton & García-Gutiérrez, 2006; Bouton & Hendrix, 2011). These alternatives are likely to entail changes in the learning rules to accommodate the conditioning of the non-temporal and temporal elements of the S stimulus, and perhaps the ITI. Independently of how such issues may be resolved, the present results and theoretical explorations attest to the importance of the interdimensional protocol to our understanding of timing.

Supplemental material

Supplemental content is available via the "Supplemental" tab on the article's online page (<http://dx.doi.org/10.1080/17470218.2015.1091851>).

REFERENCES

- ABET (II) [Computer software]. Retrieved from www.lafayetteneuroscience.com
- Balci, F., Gallistel, C. R., Allen, B. D., Frank, K. M., Gibson, J. M., & Brunner, D. (2009). Acquisition of peak responding: What is learned? *Behavioural Processes*, 80, 67–75.
- Bouton, M. E., & García-Gutiérrez, A. (2006). Intertrial interval as a contextual stimulus. *Behavioural Processes*, 71, 307–317.
- Bouton, M. E., & Hendrix, M. C. (2011). Intertrial interval as a contextual stimulus: Further analysis of a novel asymmetry in temporal discrimination learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 37(1), 79–93.
- Bush, R., & Mosteller, F. (1955). *Stochastic models of learning*. New York: John Wiley & Son.
- Catania, A. C. (1970). Reinforcement schedules and psychophysical judgment: A study of some temporal properties of behavior. In W. N. Schoenfeld (Ed.), *The theory of reinforcement schedules* (pp. 1–42). New York: Appleton-Century-Crofts.
- Church, R. M. (2003). A concise introduction to scalar timing theory. In W. H. Meck (Ed.), *Functional and neural mechanisms of interval timing* (pp. 3–22). Boca Raton, FL: CRC Press.
- Church, R. M., & Deluty, M. Z. (1977). Bisection of temporal intervals. *Journal of Experimental Psychology: Animal Behavior Processes*, 3, 216–228.
- Church, R. M., & Gibbon, J. (1982). Temporal generalization. *Journal of Experimental Psychology: Animal Behavior Processes*, 8, 165–186.
- Colwill, R. M. (1984). Disruption of short-term memory for reinforcement by ambient illumination. *The Quarterly Journal of Experimental Psychology Section B*, 36(3), 235–258.
- Dougherty, D. H., & Wixted, J. T. (1996). Detecting a nonevent: Delayed presence-versus-absence discrimination in pigeons. *Journal of the Experimental Analysis of Behavior*, 65(1), 81–92.
- Elsmore, T. F. (1971). Control of responding by stimulus duration. *Journal of the Experimental Analysis of Behavior*, 16, 81–87.
- Ernst, M. D. (2004). Permutation methods: A basis for exact inference. *Statistical Science*, 19, 676–685.
- Fleshler, M., & Hoffman, H. S. (1962). A progression for generating variable interval schedules. *Journal of the Experimental Analysis of Behavior*, 5, 529–530.
- Gallistel, C. R. (1990). *The organization of learning*. Cambridge, MA: Bradford Books/MIT Press.

- Ghirlanda, S., & Enquist, M. (2003). A century of generalization. *Animal Behaviour*, 66, 15–36.
- Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review*, 84, 279–325.
- Gibbon, J. (1981). On the form and location of the psychometric bisection function for time. *Journal of Mathematical Psychology*, 24, 58–87.
- Gibbon, J. (1991). Origins of scalar timing theory. *Learning and Motivation*, 22, 3–38.
- Gibbon, J., Church, R. M., & Meck, W. H. (1984). Scalar timing in memory. In L. Allan & J. Gibbon (Eds.), *Timing and time perception* (pp. 52–77). New York: Annals of the New York Academy of Sciences.
- Grant, D. S. (1991). Symmetrical and asymmetrical coding of food and no-food samples in delayed matching in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 17(2), 186–193.
- Guttman, N., & Kalish, H. (1956). Discriminability and stimulus generalization. *Journal of Experimental Psychology*, 51, 79–88.
- Honig, W. K., & Urcuioli, P. J. (1981). The legacy of Guttman and Kalish (1956): 25 years of research on stimulus generalization. *Journal of the Experimental Analysis of Behavior*, 36, 405–445.
- Jenkins, H. M., & Harrison, R. H. (1960). Effect of discrimination training on auditory generalization. *Journal of Experimental Psychology*, 59, 246–253.
- Jozefowicz, J., Staddon, J. E. R., Cerutti, D. T. (2009). The behavioral economics of choice and interval timing. *Psychological Review*, 116, 519–539.
- Kaiser, D. H. (2008). The proportion of fixed interval trials to probe trials affects acquisition of the peak procedure fixed interval timing task. *Behavioural Processes*, 77, 100–108.
- Killeen, P. R., & Fetterman, J. G. (1988). A behavioral theory of timing. *Psychological Review*, 95, 274–295.
- Lejeune, H., & Wearden, J. H. (2006). Scalar properties in animal timing: Conformity and violations. *The Quarterly Journal of Experimental Psychology*, 59, 1875–1908.
- Machado, A. (1997). Learning the temporal dynamics of behavior. *Psychological Review*, 104, 241–265.
- Machado, A., Malheiro, M. T., & Erilagen, W. (2009). Learning to time: A perspective. *Journal of the Experimental Analysis of Behavior*, 92(3), 423–458.
- MED-PC (IV) [Computer software]. Retrieved from www.med-associates.com
- Mellgren, R. L., Mays, M. Z., & Haddad, N. F. (1983). Discrimination and generalization by rats of temporal stimuli lasting for minutes. *Learning and Motivation*, 14, 75–91.
- Monteiro, T., & Machado, A. (2008). Oscillations following periodic reinforcement. *Behavioural Processes*, 81, 170–188.
- Rescorla, R. A. (2001). Experimental extinction. In R. R. Mowrer & S. B. Klein (Eds.), *Handbook of contemporary learning theories* (pp. 119–154). Mahwah, NJ: Lawrence Erlbaum.
- Richelle, M., & Lejeune, H. (1980). *Time in animal behavior*. New York: Pergamon.
- Riley, D. A. (1968). *Discrimination learning*. Boston: Allyn & Bacon.
- Rilling, M. (1977). Stimulus control and inhibitory processes. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 432–480). Englewood Cliffs, NJ: Prentice Hall.
- Russell, R., & Kirkpatrick, K. (2007). The role of temporal generalization in a temporal discrimination task. *Behavioural Processes*, 74, 115–125.
- Sherburne, L. M., & Zentall, T. R. (1993a). Asymmetrical coding of food and no-food events by pigeons: Sample pecking versus food as the basis of the sample code. *Learning and Motivation*, 24(2), 141–155.
- Sherburne, L. M., & Zentall, T. R. (1993b). Coding of feature and no-feature events by pigeons performing a delayed conditional discrimination. *Animal Learning & Behavior*, 21(2), 92–100.
- Siegel, S. F. (1986). A test of the similarity rule model of temporal bisection. *Learning and Motivation*, 17, 59–75.
- Spetch, M. L., & Cheng, K. (1998). A step function in pigeons' temporal generalization in the peak shift task. *Animal Learning and Behavior*, 26, 103–118.
- Stubbs, D. A. (1968). The discrimination of stimulus duration by pigeons. *Journal of the Experimental Analysis of Behavior*, 11, 223–238.
- Switalski, R. W., Lyons, J., & Thomas, D. R. (1966). Effects of interdimensional training on stimulus generalization. *Journal of Experimental Psychology*, 72, 661–666.
- Terrace, H. S. (1966). Stimulus control. In W. K. Honig (Ed.), *Operant behavior: Areas of research and applications* (pp. 271–344). Englewood Cliffs, NJ: Prentice Hall.
- Thomas, D. R. (1969). The use of operant conditioning techniques to investigate perceptual processes in

- animals. In R. M. Gilber & N. S. Sutherland (Eds.), *Animal discrimination learning* (pp. 1–33). New York: Academic Press.
- Vieira de Castro, A. C., & Machado, A. (2012). The interaction of temporal generalization gradients predicts the context effect. *Journal of the Experimental Analysis of Behavior*, 97, 263–279.
- Vieira de Castro, A. C., Machado, A., & Tomanari, G. Y. (2013). The context effect as interaction of temporal generalization gradients: Testing the fundamental assumptions of the Learning-to-Time model. *Behavioural Processes*, 95, 18–30.
- Weisman, R., Brownlie, L., Olthof, A., Njegovan, M., Sturdy, C., & Mewhort, D. (1999). Timing and classifying brief acoustic stimuli by songbirds and humans. *Journal of Experimental Psychology: Animal Behavior Processes*, 25, 139–152.
- Wilson, B., & Boakes, R. A. (1985). A comparison of the short-term memory performances of pigeons and jackdaws. *Animal Learning & Behavior*, 13(3), 285–290.
- Wixted, J. T. (1993). A signal detection analysis of memory for nonoccurrence in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 19(4), 400–411.