Published in final edited form as:

J Exp Psychol Anim Learn Cogn. 2015 January ; 41(1): 39–51. doi:10.1037/xan0000049.

Optimal response rates in humans and rats

David M. Freestone¹,

Brown University

Fuat Balci,

Koç University, İstanbul, Turkey

Patrick Simen, and

Oberlin College

Russell M. Church

Brown University

Abstract

The analysis of response rates has been highly influential in psychology, giving rise to many prominent theories of learning. There is, however, growing interest in explaining response rates, not as a global response to associations or value, but as a decision about how to space responses in time. Recently, researchers have shown that humans and mice can time a single response optimally, i.e., in a way that maximizes reward. Here, we use the well-established DRL timing task to show that humans and rats come close to optimizing reinforcement rate, but respond systematically faster than they should.

Keywords

Interval Timing; Reward Rate; Response Rate; DRL; Decision Making

Introduction

Response rates are studied heavily in experimental psychology. Often, response rates are viewed as a measure of an underlying internal variable like associative strength (see Pearce & Bouton, 2001 for an overview), subjective value (Baum, 1974; Herrnstein, 1961), or confidence (Blough, 1967; Yi, 2009), to name a few. In most theoretical frameworks, there is a directional flow in which reinforcement updates the internal variable, and the internal variable is transformed into a response rate. Often, this is where the flow stops because it represents one of the simplest models that allow us to make useful inferences about those internal variables by measuring the response rate. But the picture is somewhat more complicated because the response rates often modulate reinforcement in instrumental situations. There is a response-reinforcer feedback loop in which response rates dictate how

Correspondence: David Freestone, Center for Neural Science, New York University, 4 Washington Pl., New York, NY 10003, freestone@nyu.edu.

¹Now at New York University

much reinforcement is received, and reinforcement may influence the response rate (Baum, 1973; Rachlin, 1978).

Rachlin (1978) approached this as a maximization problem: The response rate should be chosen such that some value function (e.g., reinforcement rate) is maximized. Nowhere is this idea more evident than in the Matching Law literature from the late 1970s and early 1980s (e.g., (Baum, 1981; Herrnstein & Heyman, 1979; Heyman & Luce, 1979; Heyman & Tanz, 1995; Houston & McNamara, 1981; Staddon & Motheral, 1978). This view has not driven nearly as much enthusiasm or original research in single-operant paradigms.

One recent example, though, comes from the reinforcement learning literature (Niv, Daw, & Dayan, 2005; Niv, Daw, Joel, & Dayan, 2007; Niv, Joel, & Dayan, 2006). Normally, reinforcement learning models either assume a monotonic transformation from value (or association) to response rate (see Miller, Barnet, & Grahame, 1995), or they are silent on the transformation altogether (Sutton & Barto, 1998). But Niv and colleagues view the response rate as the outcome of a sequence of decisions about which action to perform, and for how long (which they term *vigor*). How long to perform an action is amenable to optimality analysis: Given the schedule of reinforcement, what is the optimal time between actions needed to maximize the reinforcement rate? They show that this time is proportional to the cost of making an action, and more importantly for the present discussion, but perhaps not surprisingly, inversely proportional to the average reinforcement rate. That is, it is optimal to perform all actions faster (a higher response rate) when the average reinforcement rate is higher. They derived exact results for vigor for ratio schedules, but not for interval schedules, nor for more specific reinforcement schedules.

The differential reinforcement of low rates (DRL) task was originally developed to test the hypothesis that rats could withhold responding (Wilson & Keller, 1953). In this task, rats are rewarded for spacing their responses in time. A lever press is reinforced if, and only if, the time between it and the last response (the interresponse time; IRT) is greater than some experimenter defined target duration often called the *DRL schedule*. This task is interesting because it imposes a well-defined feedback loop between the response rate and the reinforcement rate.

Wearden (1990) set the foundation for the work presented here by analyzing what the interresponse time should be in order to maximize reinforcement on the DRL task. He assumed (as we do) that the correct quantity to maximize is the reinforcement rate (see Balci, Freestone, & Gallistel (2009) for a task in which reinforcement per trial is maximized). Wearden's initial simulations suggested that the variability of the interresponse time distribution was a key component in the feedback loop and should be accounted for by any model that seeks to maximize reinforcement rate on this task (in contrast to earlier "minimize the average delay" models that did not account for variability, see Zeiler, Scott, & Hoyert, 1987). Çavdaro lu, Zeki, & Balcı (2014) laid the critical mathematical foundation by using statistical decision theory (Girshick & Blackwell, 1954; Wald, 1950) to rigorously specify how the interresponse time should depend on the variability. They showed that humans, without explicit task instructions, learned the target intervals and nearly maximized

reinforcement rate. The present article seeks to replicate these findings in humans and extend the findings to rats.

Optimal Response Rates in the DRL Task

The mathematical treatment of the interresponse times in the DRL task requires a few definitions and assumptions. We define the optimal response rate as the response rate that maximizes the reinforcement rate. The reinforcement rate R is

$$R(t, \gamma) \equiv p(r; t, \gamma) t^{-1} \quad (1)$$

where r represents reinforcement or not (1 or 0), t is the mean interresponse time, γ is the variability of the interresponse time distribution (we will formalize this in a moment), so that $p(r; t, \gamma)$ is the probability of a reinforcement given the interresponse time distribution. It is an empirical fact that animals are imperfect timers, that is, $\gamma > 0$ (Gibbon, Malapani, Dale, & Gallistel, 1997; Gibbon, 1977).

Because of this fact, there is a speed-accuracy tradeoff in this task. The probability of reinforcement (accuracy) grows with t because the longer the rat waits before responding, the more likely it is that a lever press will be reinforced. The highest possible accuracy is when the entire interresponse time distribution falls longer than the target interval T . By definition, high accuracy results in a high reinforcement rate per response. The problem is that high accuracy leads to a low rate of return per unit of time. As t grows, the time between reinforcers grows, leading to a smaller overall reinforcement rate per unit of time.

The first assumption that our model makes is that animals *do not* control their own variability. This is an oversimplification (see Schoenfeld, Harris & Farmer, 1966), but a useful one. The second assumption is that animals *do* control their own mean interresponse time. In other words, the response-reinforcer feedback loop depends on both the mean and the variability of the response distribution, but the animal can only adjust the mean in order to increase reinforcement rate. How the animal does this, either through hill-climbing the reinforcement rate gradient or through computations on representations (etc.), is outside the scope of this paper, but see Kheifets & Gallistel (2012) and Trommershäuser, Landy, & Maloney (2006) for some thoughts and data on the topic.

The response rate that maximizes the reinforcement rate is obtained by differentiating equation 1 with respect to time, setting it to zero, and solving for t (the response rate is $1/t$).

We can test for optimal response rates in two ways. The first is empirical: Shift the actual interresponse time distribution until the reinforcement rate would be maximized. The amount the data needed to be shifted gives an index of optimality (the smaller the shift, the closer to optimal the response times). The second is theoretical: Fit the interresponse time distribution with a suitable generative model and analyze the parameters of the fit. We believe the second method is more informative, but we will leave that judgment to our readers. However, we will explain that method in detail now so we can make concrete statements about equation 1.

When Wearden (1990) performed his initial simulations, he assumed the interresponse time distribution was Gaussian. Sanabria & Killeen (2008) also assumed normally distributed interresponse times. Our empirical data, and others like it, suggest that the interresponse time distribution is positively skewed and bimodal (Richards, Sabol, & Seiden, 1993). The full interresponse time distribution has both short (seemingly untimed) and long (seemingly timed) responses. We will formalize the full bimodal model in a moment, but for simplicity and clarity we focus on the long (timed) responses here. Taking our cues from Simen et al. (2013), we assume the interresponse times come from an inverse Gaussian distribution (in our Results section, we show this is superior to the Gaussian assumption). The probability density function for this is given by

$$IG(x; \mu, \lambda) = \sqrt{\frac{\lambda}{2\pi x^3}} e^{-\frac{\lambda(x-\mu)^2}{2x\mu^2}} \quad (2)$$

This equation gives the probability of obtaining x given scale and shape parameters μ and λ (respectively). This allows us to write the probability of a reinforcement given the interresponse time distribution as

$$p(r; t, \gamma) = \int_{\tau=T}^{\infty} IG(\tau; t, \lambda) d\tau \quad (3)$$

with λ given by

$$\lambda = \frac{t}{\gamma^2} \quad (4)$$

and γ is the coefficient of variation. In words, the probability of reinforcement is simply the fraction of interresponse times greater than the target interval, T , and this fraction depends on γ . Our choice of parameterizing the model with the coefficient of variation γ instead of the inverse Gaussian scale parameter λ is important. Timescale invariance (that response time distributions overlap when plotted on a normalized time-axis) is a well-established result in the timing literature, including the DRL task (e.g., Wearden, 1990), and underlies all timing models. This empirical fact allows us to normalize the observed data from different DRL schedules and then compare them against a single optimum. Equation 4 must hold in order for the inverse Gaussian to be timescale invariant across different values of t .

Solving for the optimal response rate analytically in this case is difficult. To give some intuition, Çavdaro lu, Zeki, Balci (2014) showed that the optimal interresponse time is the solution to the equation

$$0 = 1 - p(r|t, \gamma) - T \times IG(T; t, \lambda) \quad (5)$$

when solved for t . We will denote the solution t_o ; it is the optimal interresponse time needed to maximize reinforcement rate. The result is a set of optimal interresponse times—an *optimal performance curve*—that depends on the animal's coefficient of variation γ .

Figure 1 shows the step-by-step process toward creating the optimal performance curve. The target interval is set to 1 (dotted vertical line), reflecting the fact that the interresponse time distributions can be normalized by the target time because of timescale invariance. The rising dashed line displays the probability of reinforcement (Equation 3). The falling dashed line displays $1/t$. This line shows, holding the probability of reinforcement constant, how decreases in the response rate decrease the reinforcement rate. The product of these two curves (element-wise multiplication) gives the reinforcement rate as a function of t for a given γ , shown as the solid black curve. The maximum of this curve is the maximum reinforcement rate possible (given γ), and its location on the time-axis gives the optimal mean interresponse time t_o . Normally distributed interresponse times do not change the main result, only the details.

Panel b shows how the reinforcement rate curve depends on the animal's coefficient of variation. Each solid black curve shows the reinforcement rate curve for a different coefficient of variation. Keep in mind that the coefficient of variation is a constraint in our model; the animal cannot optimize it. The first thing of note is that as γ increases, the entire reinforcement rate function decreases. All things being equal, animals with higher variability will necessarily get less reward than animals with lower variability. This is because animals with higher variability necessarily make more errors (i.e., early responses that go unreinforced). The dashed line in the figure displays the optimal performance curve that runs through the maxima of all possible reinforcement rate curves when γ is varied. One way to think about the optimal performance curve is that it specifies the ideal relationship between an individual animal's variability and its mean interresponse time. This is one of the critical predictions we test in this article: As the variability increases, so should the mean interresponse time.¹

This prediction is tested in 57 rats and 14 humans. The take-home results are that, while both species show some critical features of optimality, their response rates are systematically faster than optimal.

Methods

Sixty rats and 15 humans were tested on this task. Rats were required to press a lever to obtain a food pellet, and humans pressed the spacebar on a computer keyboard to obtain a point that was later converted into money (between \$10 and \$20 for a session, depending on performance). Both rats and humans were tested in a between-subject design in which each participant saw only one DRL target interval. Participants were only rewarded if the time between two consecutive responses was greater than a particular target interval. For rats, the target intervals were 7, 10, 14, 28, and 56 seconds (12 rats per schedule). For humans, the

¹This is true for plausible values for the coefficient of variation γ . For larger values of γ (roughly > 0.7), the optimal performance curve turns back toward the schedule. This is because the reinforcement rate curve flattens to the point where waiting longer no longer compensates for having high variability.

schedules were 5, 8, 10, 12, and 15 seconds. Rats were tested for 41 daily sessions, and humans were tested in a single session.

Portions of the human and rat data from this study were used to construct Figure 4 in Balci et al. (2010) and briefly discussed there.

Subjects

Humans—Fifteen adults (7 males and 8 females between 18 and 30 years old) participated. They were recruited from online postings and fliers posted around the Princeton University campus. All participants provided written consent for their participation, and the Institutional Review Panel for Human subjects of Princeton University approved the experiment. One participant in the DRL 12 s group reported that s/he did not follow the instructions in the task, so the data is not included in the analysis.

Rats—Sixty male Sprague Dawley rats (Taconic Laboratories, Germantown, NY) were used in this experiment. The colony room was on a 12:12 light-dark cycle (lights off at 8:30 am). During the dark cycle, dim red lights provided light in the colony room and testing rooms. Upon arrival, the rats were roughly 4 weeks old and weighed between 75 and 100 grams. During the first week, the rats were on a free-feeding schedule. After a week, their daily food (FormuLab 5008) was rationed to 16 grams per day. During each experimental session, the rats were fed 45 mg Noyes pellets (Improved Formula A) as a reinforcer. Water was available ad libitum in both the home cage and the testing chamber. Forty-eight of the rats were previously used in a Fixed-Interval experiment with gaps that used lights and sounds as stimuli (these rats were approximately 6 months old at the start of this experiment). They were assigned to each DRL schedule randomly so that their previous training did not impact the results. Over the course of the experiment, one rat rarely pressed the lever, and two rats died. This left 57 rats for data analysis. All procedures were approved by the Brown University Institutional Animal Care and Use Committee.

Apparatus

Humans—The visual stimulus was a white square on black background. The stimuli were coded in MATLAB and displayed on a Macintosh computer using the Psychophysics Toolbox extension (Brainard, 1997). Participants indicated their responses with a standard computer keyboard.

Rats—Twenty-four experiment chambers (Med Associates, dimensions 25 × 30 × 30 cm) were situated in two separate experiment rooms (twelve in each room). Each chamber was contained in a sound-attenuating box (Med Associates, dimensions 74 × 38 × 60 cm) with a fan for ventilation. Each experimental chamber was equipped with a pellet dispenser (Med Associates, ENV-203) on the front wall that delivered the reward into a food cup. A head entry into this cup interrupted a photo beam (Med Associates, ENV-254). On both sides of the food cup, there were two retractable levers. On the opposite wall, a water bottle protruded into the chamber allowing ad libitum access to water during the session. A lick on the spout of the water bottle completed an electric circuit. Four Gateway Pentium III/500 computers running Med-PC for Windows (version 1.15) controlled the experiments and

recorded the data. The interruption of the photo beam and the completion of the lick and lever circuits were recorded in time-event format with 2-ms accuracy.

Procedure

Humans—Participants attended single session experiments. One of the following DRL schedules was assigned to each session: 5s, 8s, 10s, 12s, or 15s. Participants were told that they would earn money for their responses if, and only if, they waited for a minimum duration since their previous response (but they were not explicitly told the duration). They were also told that each response would reset the trial clock. Participants were told how much they would earn for each correct response, that the session time was fixed, and they were encouraged to make as much money as possible. Participants could earn at most around \$20 per session and the reward magnitude per correct response was adjusted for different DRL schedules (higher magnitude for longer DRL schedules). This aimed to equate the reinforcement rate for different schedules. Participants were also asked not to count, tap, or adopt any rhythmic activity in order to time the intervals.

Participants were first presented with the minimum wait time (the DRL schedule) three times. This duration was signaled by a white square presented on a black background. Participants then reproduced this interval roughly 50 times over two blocks. During this phase, a white square appeared in the middle of the screen to signal the onset of the interval and participants were instructed to press the space key when they thought the target interval elapsed.

The square disappeared upon pressing the space key and participants were provided feedback on a fixed-length horizontal line. The parametric feedback graphically indicated how far their reproduction time was from the target interval in that trial. A white vertical line represented the target (white) interval and the red vertical line represented the reproduction interval in that particular trial. The length of the horizontal line was normalized by the DRL schedule in order to ensure that the same number of pixels referred to the same proportion of discrepancy for every schedule. The point of the feedback was to give the participants time to learn their own variability, if they did not already have an estimate (the rats had 41 sessions in which to do so).

After the reproduction blocks, participants were tested in 8 five-minute long blocks of DRL testing. Test blocks were initiated with the appearance of a white square in the middle of the screen. Participants could respond at any time and rate they wished during test blocks.

If the interresponse time was greater than or equal to the target interval, the square turned green, a brief beep sounded, and the amount of money earned that trial was displayed on the screen. If the interresponse time was earlier than the target interval, the square turned red and was accompanied by a brief buzzer sound. The square turned white again after the feedback and a new interval began. Cumulative earning was always present on top of the screen during the test blocks.

In order to prevent explicit counting, a secondary task was used during both familiarization and DRL testing. Participants were presented with a four-digit number at the beginning of

each block and a single-digit number was presented at the end of that block. Participants were asked if that single digit was one of the four digits observed at the beginning of that block. Participants were told at the beginning of the experiment that they would walk away with the total earnings from the timing tests multiplied by the proportion of correct responses in the working memory task. For example, if they earned \$20 from the timing portion of the experiment, but only got 50% of the memory questions correct, they would earn a total of \$10.

Rats—Forty-eight of the 60 rats were tested daily in one-hour sessions. Twelve of the sixty rats were run nightly in three one-hour sessions per night (with a three-hour break in between sessions). The rats were placed in the operant chamber and a lever was inserted. Rats were rewarded for spacing their lever presses (interresponse time) by at least the target interval. If a rat pressed the lever before the target interval had elapsed, food was not delivered and the interval started over. There were five groups (12 rats per group), each with a different target interval: 7, 10, 14, 28, and 56 s (the 10 s group was run several months later, and overnight). Unlike the human task, the reinforcement rate was not held constant across groups. The amount of reinforcement per session for an optimal animal ranged from about 60 (target interval of 56) to about 515 (target interval of 7) in a one-hour session. The rats that were tested daily were fed 16 grams of food in their home cages after the session, while the rats that were tested nightly earned all their food in the testing chambers.

Results

As with many data sets on the DRL task, our data showed a bimodal interresponse time distribution (one near 0 and the other near the target interval). Given this, there are a few approaches toward analysis (e.g., Richards et al., 1993; Sanabria & Killeen, 2008). We employ two methods below. The first method follows Sanabria & Killen (2008) and is in spirit of building a probabilistic model. We model the data as coming from two processes: one fast and untimed (e.g., impulsive), and a slower and timed. We model the fast process as an exponential distribution; animals emit untimed responses as a Poisson process with some rate κ . These make up a fraction p of all responses. The slower, timed process is modeled as an inverse Gaussian distribution as described above. The full model for describing the interresponse time distribution is

$$p(irt=x)=p\kappa e^{-\kappa x}+(1-p)IG(x;t,\lambda) \quad (6)$$

where $IG(x; t, \lambda)$ was given above. The four parameters κ , p , t , and λ were fit with maximum likelihood estimation using *fminsearch* in Matlab, which implements the Nelder-Mead simplex method. We bounded the parameters within plausible limits ($p \in (0,1)$; $\kappa \in (0,100)$; $t \in (0,100)$; $\lambda \in (0,4000)$) and ran the algorithm with 50 seeds uniformly distributed over these bounds. The average fits are displayed in Table 1 (mean $\pm$ sem; note that γ is shown instead of λ), and two examples of individual data (one rat, one human, same DRL target interval) is shown in Figure 2, and will be described in more detail below.

The second method for dealing with bimodal interresponse times is to use the empirical data exactly to compute the shift necessary to obtain optimal performance. We do not dwell heavily on this method, but we include it for completeness. Here, the actual interresponse times (normalized by T) are shifted by t until the reinforcement rate is maximized. The value of the shift gives an index of optimality. Smaller shifts indicate closer to optimal (reinforcement maximizing) performance. The shift t for each animal was fit using *fminsearch* (50 seeds). In general, the fitting procedure used (e.g., Nelder-Mead simplex fitting vs. BFGS quasi-newton fitting) does change the exact fitted values somewhat, but it does not change any of the conclusions drawn from the fits. We chose the Nelder-Mead algorithm here because it does not require approximating the gradient of the likelihood function.

The untimed responses are not optimal in the sense described by our model. They may be adaptive in some other sense, but we do not explore the issue here. The probabilistic model analysis (Equation 6) factors out the untimed responses so that the statistics of the timed responses can be isolated. The shift analysis uses both the timed and untimed responses. This allows us to use the entire data set to assess optimality (all the interresponse times, not just the timed responses). Using more than one analysis technique helps to show that the results we obtained are robust (and not specific to our particular probabilistic model assumptions). And since our model considers the untimed responses suboptimal, the shift analysis helps determine how heavily these untimed responses impact performance.

For most of the analyses shown here, the last 10 sessions from each rat were pooled. This resulted in 1897.04 ± 107.87 (mean $\pm$ sem) responses per rat. Each human was only run for a single session. This resulted in 266.07 ± 27.05 responses. For the analyses during learning (rats only), each session was fit individually. Occasionally, the rats made no lever presses for long periods of time (e.g., they were drinking or sleeping). These outliers do not represent task performance and can heavily skew the fits. They were removed using a cutoff for responses greater than $q_3 + w(q_3 - q_1)$, where q_1 and q_3 are the first and third quartiles, respectively. For normally distributed data, $w=1.5$ is often used, which results in a criterion about 2.7σ higher than the mean. For our data, we tripled this to $w=4.5$ to exclude on the most extreme interresponse times.

For completeness, we also fit the mixture model assuming Gaussian distributed timed responses. Examples of the fits are shown in Figure 2. The inverse Gaussian fit the data better in 47/57 (82.5%) rats; the log-likelihood ratios significantly favored the inverse Gaussian assumption [$t(56)=6.04$, $p < 0.001$, $M = 44.76 \pm 7.4$, $d=0.81$]. In humans, the inverse Gaussian fit the data better in 14/14 (100%), and the log-likelihood ratios significantly favored the inverse Gaussian assumption [$t(13)=3.86$, $p < 0.001$, $M = 9.38 \pm 2.4$, $d=1.07$].

Descriptive Results

Raw interresponse time data, and the resulting estimated probability distributions, are shown in Figure 2. The first column (Figure 2a–b) shows a single rat for which the target interval was 10 s. Panel b (bottom) shows every interresponse time from the last 10 sessions as a gray dot (schedule shown in black). This was the data used in the maximum likelihood fits described above (each subject was fit individually). Panel a (top) shows the same data

plotted as a probability distribution, along with the IG fit (black) and Gaussian fit (dark gray dashed). A single human participant is shown in panels c and d (second column). Both the human and rat examples were drawn randomly from the same target interval (10 s). These random examples show a strong similarity between the IG and Gaussian fits, although this is not true for all subjects (the IG fits significantly better, as described above). In some of the analysis that follows, the normalized interresponse time is used. This refers to t/T , where t is the estimate of the mean of the IG component of the distribution, and T is the target interval. It is reasonable to think of this as simply scaling the entire raw interresponse times by T . Note that it is inappropriate to use a z-scores to normalize the data because it changes the variability of the distribution—the exact thing that needs to vary across subjects to test the hypothesis that the mean of the distribution precisely correlated with the variability.

Learning—We assessed learning over sessions by evaluating power law regressions of the form $\log y = \beta_0 + \beta_1 \log x$, where y is the parameter of interest (p , γ , and t), and x is session number. The value of β_0 is the constant (in logs) of the regression, and a β_1 significantly different from 0 indicates a change in $\log y$ over sessions. We used a mixed-effects model letting β_0 include both fixed and random effects (β_1 was fixed for a group), although this assumption is not necessary for reproducing our main results.

The coefficient of variation γ significantly decreased over sessions for all groups of rats (min $\beta_1 = 0.07 \pm 0.02$, $t(408) = -3.26$, $p = 0.005$). That is, over sessions, the width of the timed responses became narrower. Colloquially, the rats became better at timing their responses. The mean $\pm$ sem over sessions for each group is shown in Figure 3a.

The proportion of untimed responses p decreased over sessions in three of the five groups of rats (7s: $\beta_1 = -0.25 \pm 0.04$, $t(490) = -6.40$, $p < 10^{-10}$; 14s: $\beta_1 = -0.41 \pm 0.07$, $t(449) = -5.67$, $p < 10^{-7}$; 28s: $\beta_1 = -0.23 \pm 0.08$, $t(408) = -3.4$, $p < 10^{-3}$). But p did not significantly change for DRL 10 and 56s groups (14s: $\beta_1 = 0.006 \pm 0.13$, $t(538) = 0.04$, $p = 0.96$; 56s: $\beta_1 = -0.09 \pm 0.07$, $t(490) = -1.20$, $p = 0.23$). The mean $\pm$ sem over sessions for each group is shown in Figure 3b.

The mean of the timed response times t significantly increased over sessions. For the shorter DRL intervals (7s, 10s, and 14s) the sessions coefficient hovers around 0.27 (7s: $\beta_1 = 0.25 \pm 0.03$; 10s: $\beta_1 = 0.30 \pm 0.03$; 14s: $\beta_1 = 0.27 \pm 0.03$, all $p < 10^{-17}$). For the two longer DRL intervals, the sessions coefficient is much higher with the DRL 28s group at $\beta_1 = 0.81 \pm 0.06$ ($p = 10^{-32}$) and the DRL 56s group at $\beta_1 = 1.12 \pm 0.06$ ($p = 10^{-52}$). The mean $\pm$ sem over sessions for each group is shown in Figure 3c. The figure and the analysis suggests that it takes significant time for rats to learn longer intervals. It is possible that the rats in the DRL 28 and 56s groups did not fully reach steady state.

Steady State—The proportion of untimed responses p did not significantly differ by DRL schedule in either rats or humans (Rats: $F(4,56) = 2.15$, $p = 0.09$, $\eta^2 = 0.14$; Humans: $F(4,13) = 1.27$, $p = 0.35$, $\eta^2 = 0.36$).

There were no significant differences in the rate parameter of untimed responses κ in either rats or humans. (Rats: $F(4,56) = 2.37$, $p = 0.06$, $\eta^2 = 0.15$; Humans: $F(4,13) = 0.90$, $p = 0.50$,

$\eta^2=0.29$). Note, though, that the κ estimates in humans have large standard errors, probably because it is difficult to properly identify the κ parameter when the proportion of untimed responses is very small.

For humans, there were no significant differences in the coefficient of variation γ ($F(4,13)=0.52$, $p=0.73$, $\eta^2=0.18$). But there were differences in rats ($F(4,56)=15.5$, $p<2\times 10^{-8}$, $\eta^2=0.54$). Post-hoc tests revealed that this was because the rats with longer DRL intervals (DRL 28 and 56s) differed significantly from the DRL 7s and DRL 10s groups. The DRL 10s group also differed from the DRL 14s group. The means $\pm$ sem are shown in Table 1.

Of course, the fact that the t parameter differs across groups is trivial because both rats and humans were grouped by the target interval. An important prerequisite for optimal performance is that the interresponse times be greater than the target interval. This is true for the rats in the DRL 7s, 10s, and 14s groups (smallest effect in the DRL 14s group, $t(10)=5.63$, $p=0.001$, $d=1.78$), and every rat in these three groups had a mean interresponse time t greater than the schedule. The DRL 28s group showed a marginally significant result ($t(9)=2.38$, $p=0.041$, $d=0.79$), and 20% (2/10) of the rats had a mean interresponse time less than 28s. The interresponse times in the DRL 56s group did not significantly differ from 56s ($t(11)=-1.27$, $p=0.23$, $d=-0.38$), and 75% (9/12) of the rats had a mean interresponse time less than 56s. Figure 4b shows that the normalized interresponse time (t/T) tends to decrease as the interval increases (black dots; the figure will be described in more detail below). That is, the longer the target interval, the more unlikely it is that rats withhold responding for the entire interval. This may be because of the incomplete learning in those groups, as described above, or a natural feature of longer DRL intervals (e.g., Richardson & Loughhead, 1974).

Every human in our study had mean interresponse times t greater than their target interval. The sample size in two of the groups (12s and 15s) is too small for group level statistics ($n=1$). The participants in the 5 and 10s groups had mean interresponse times greater than the schedule (minimum in 5s group: $t(3)=4.32$, $p=0.02$, $d=2.49$). In the 8s group, one of the subjects had an interresponse times much higher than the others, which resulted in an insignificant t-test ($t(3)=2.0$, $p=0.14$, $d=1.15$). The mean interresponse times for the four participants in this DRL 8s group are 8.12, 9.39, 9.43, and 12.68.

Because the rats in the DRL 28 and 56s groups (1) tended to have a higher coefficient of variation γ , (2) tended to respond earlier than their schedule, and (3) likely failed to asymptote (as indicated by the large session coefficient in the power regression), we do not expect them to be optimal in the sense described above. Indeed, Figure 4 shows them deviating substantially from the optimal interresponse times (black dots compared to the gray dots). For completeness, we will continue to analyze their data along side the others, but because of these caveats, we will show the single subject data from these rats in open circles rather than closed. We will discuss this further in the discussion section.

Optimality Results

Figure 4 shows the mean interresponse times for each group (black dots) and the associated optimal interresponse time (gray dots). The top panels show the raw (non-normalized)

means, and the bottom panels show the normalized means (t/T). The dashed lines show the least squares linear fit. For humans (panels c,d) the interresponse times closely track the optimal interresponse times. For rats (panels a,b), this is true for shorter intervals, but the result falls apart for longer intervals, where the normalized interresponse time deviates heavily from optimality (In Figure 4a, the slopes are different, and in Figure 4b the slope of the black line is negative.). Humans were not tested at these longer intervals, so it is an open question whether their performance would also break down in the same way. In all groups, the mean interresponse time is shorter than optimal; that is, the response rate is faster than optimal. We will unpack these results below.

The critical feature of the optimality analysis presented above is that, in order to maximize reinforcement rate, animals should account for the variability of their responses. Figure 5 shows the normalized interresponse time plotted against the coefficient of variation for every rat (Figure 5a) and human (Figure 5b) in our experiment. The optimal performance curve is shown as a black line. The dashed line shows how much the optimal performance curve has to shift in order to obtain the least squares fit to the data (Rats: -0.15 ± 0.01 , $R^2 = 0.59$; Humans: -0.07 ± 0.02 , $R^2 = 0.66$). The fit has a single parameter, τ . Note that this is similar but not identical to t . Here, we shift the optimal curve until it fits the data, while the analysis that produced t shifts the raw data (both timed and untimed responses) until the reinforcement rate is maximized. Across human participants, the mean interresponse time is highly correlated with the coefficient of variation, $F(1,12) = 54.55$, $p = 8.5 \times 10^{-6}$, $R^2 = 0.82$. This is also true for shorter DRL intervals in rats ($F(1,33) = 57.31$, $p = 1.1 \times 10^{-8}$, $R^2 = 0.63$; pooled over the 7s, 10s, and 14 s groups). When longer DRL intervals are included, there is no significant main effect of schedule, $F(1,55) = 2.11$, $p = 0.15$, $R^2 = 0.04$. This result is clearly seen in the figure. The rats in the longer DRL schedules cannot be optimal in the sense we described above, even in principle, because there is no reliable relationship between the mean and variability of their interresponse time distribution. But for the humans and the rats with shorter DRL schedules, the mean and variability are correlated as prescribed by the model.

From a practical point of view, it is probably more important whether the mean interresponse times tracked the optimal interresponse time across animals (Figure 6). Across human participants, the mean interresponse times were highly correlated with the optimal interresponse times, $F(1,12) = 28.78$, $p = 2.0 \times 10^{-4}$, $R^2 = 0.71$. In rats, this was also true in the shorter DRL schedules, $F(1,33) = 50.02$, $p = 4.3 \times 10^{-8}$, $R^2 = 0.60$, but not the longer ones ($p = 0.23$). Of course, because the optimal interresponse time depends on the variability, finding a significant regression coefficient for both γ and t_o does not give completely independent pieces of evidence that animals track optimality, but they are not completely redundant measures either since their relationship is nonlinear.

The analyses above tell us that both humans and rats (at shorter intervals) have mean interresponse times that depend on (are highly correlated with) the variability of those interresponse times. But that does not mean they are maximizing reinforcement rate, it only suggests they could be. To check for optimality, we have to assess how close the mean interresponse times are to the optimal, and how much the animals earn (their reward rate) against the optimal reward rate they could earn given their variability (recall that different

levels of variability lead to different reinforcement rates even for rate maximizing animals). We define the temporal distance ratio as the mean interresponse time divided by the optimal interresponse time: t/t_o . Note this can take on any positive value. Numbers less than one indicate the animal's interresponse time is shorter than it should be, and their response rate is too high. Numbers higher than one indicate the opposite.

Nearly all rats and humans had faster-than-optimal interresponse times. This was significant in every group of rats (minimum effect size in DRL 14s, $t(3) = -6.48$, $p = 7.1 \times 10^{-5}$, $d = -0.65$). This was true for all groups of humans (minimum effect size in DRL 5s, $t(10) = -4.47$, $p = 0.02$, $d = -1.49$) except for the DRL 8s ($t(3) = -0.31$, $p = 0.78$, $d = -0.10$). Recall that the DRL 8s group had a participant with much longer interresponse times than the others. In humans, the temporal distance ratio did not significantly depend on the schedule [$F(1,12) = 0.19$, $p = 0.67$, $R^2 = 0.02$]. For shorter intervals in rats, the ratio didn't significantly depend on the schedule, either [$F(1,33) = 2.87$, $p = 0.1$, $R^2 = 0.08$]. But when the longer intervals were included for rats, the result was highly significant [$F(1,55) = 71.77$, $p = 1.5 \times 10^{-11}$, $R^2 = 0.57$]. This again points to a key difference between the rats in the shorter and longer intervals. In fact, across all animals (human and rats) only three of 71 animals (4.2%) had a mean interresponse time greater than or equal to the optimal. For rats, the temporal distance ratios (mean $\pm$ sem, median $\pm$ iqr in parenthesis) were **7**: 0.91 ± 0.01 (0.91 ± 0.05), **10**: 0.86 ± 0.008 (0.86 ± 0.03), **14**: 0.88 ± 0.02 (0.88 ± 0.09), and **28**: 0.77 ± 0.03 (0.80 ± 0.14), **56**: 0.64 ± 0.04 (0.58 ± 0.14). For humans, the temporal distance ratios were **5**: 0.94 ± 0.01 , (0.95 ± 0.03) **8**: 0.99 ± 0.04 (0.96 ± 0.12), **10**: 0.91 ± 0.008 (0.92 ± 0.02), **12**: 0.90 (0.90), **15**: 0.97 (0.97).

Animals that have close to optimal interresponse times will have a higher reinforcement rate, and gain more food. But how much they gain depends on the curvature of the gain surface (how drastically changes in interresponse times change the reinforcement rate). Animals with higher variability have shallower reinforcement rate functions (see Figure 1); relatively large changes in their interresponse time may result in relatively small changes in their overall reinforcement rate. We defined the gain ratio as the reinforcement rate earned (based on the fits of the timed interresponse time distribution) divided by the maximum reinforcement rate for that particular subject, g/g_o . Unlike the temporal distance ratio, the gain ratio is bound by 0 and 1 (subjects cannot do worse than nothing or better than the best).

In general, the gain ratio does not significantly depend on the schedule in humans or rats (Humans: $F(1,12) = 0.15$, $p = 0.70$, $R^2 = 0.01$; Rats: $F(1,33) = 0.53$, $p = 0.47$, $R^2 = 0.01$), except, like before, when the longer interval schedules are included for rats ($F(1,55) = 43.14$, $p = 1.92 \times 10^{-8}$, $R^2 = 0.44$). For rats, the gain ratios were **7**: 0.95 ± 0.01 (0.95 ± 0.06), **10**: 0.88 ± 0.02 (0.89 ± 0.08), **14**: 0.93 ± 0.02 (0.96 ± 0.07), and **28**: 0.86 ± 0.03 (0.87 ± 0.15), **56**: 0.69 ± 0.05 (0.62 ± 0.26). For humans, the gain ratios were **5**: 0.96 ± 0.02 (0.96 ± 0.06), **8**: 0.92 ± 0.046 (0.98 ± 0.14), **10**: 0.96 ± 0.01 (0.97 ± 0.03), **12**: 0.96 (0.96), **15**: 0.995 (0.995). It is perhaps interesting to note is that, even though the performance in the DRL 28s group tended to be quite different from the performance with shorter intervals (as described at length above), the rats still obtained $86 \pm 3\%$ of the best they could have gained. There is no rigorous way to say what a good gain ratio is, partly because there are no suitable null-hypotheses, so we

leave it to our reader to decide whether ~90% in rats and ~95% in humans is suitably close to the optimal gain of 100%.

But to give some intuition, we can compare the gain ratio the subjects received with the gain ratio they would have received if they had responded at the schedule. The DRL task requires that subjects respond at or after the target interval in order to be reinforced. But it does not require that the subjects actually do so (as shown by the DRL 56s group) or even that a substantial proportion of responses are longer than the target interval (both the DRL 28s and 56s group). The gain ratio at the schedule is the baseline we would expect from a subject who learned the interval, and chose an optimal mean interresponse time without regard to the variability around it. We define the marginal gain ratio as how much better the subjects do by having a mean interresponse time that depends on (is correlated with) their variability as $(g - g_T)/(g_0 - g_T)$, where g_T is the gain the subject would have obtained given their actual coefficient of variation, but with a mean at the schedule. This value must be smaller than the gain ratio because g_T is positive, but how much smaller depends on both how close the subject is to the optimal interresponse time and how shallow the curvature of the reinforcement rate function is (which depends on how large the coefficient of variation is). For rats, the marginal gain ratios were **7**: 0.86 ± 0.03 (0.86 ± 0.19), **10**: 0.64 ± 0.05 (0.68 ± 0.24), **14**: 0.75 ± 0.07 (0.85 ± 0.27), and **28**: 0.34 ± 0.17 (0.48 ± 0.69), **56**: -0.53 ± 0.24 (-0.85 ± 1.08). For humans, the marginal gain ratios were **5**: 0.89 ± 0.05 (0.92 ± 0.15), **8**: 0.81 ± 0.14 (0.93 ± 0.32), **10**: 0.87 ± 0.04 (0.91 ± 0.08), **12**: 0.86 (0.86), **15**: 0.98 (0.98). There are two conclusions. The trivial conclusion is that overall both the rats and humans still maintain a relatively high ratio indicating closeness to optimality. The more interesting conclusions are that both the DRL 10s and 28 s group's marginal gain is quite a bit lower than their gain ratio. This indicates those groups either had relatively earlier mean interresponse times (Figure 4 suggests they did) or the curvature of the reinforcement rate function is shallow, or both (the likely culprit for the DRL 28s group). On the whole, these results suggest that the rats were about 75% better and humans about 85% better than if they had optimized their mean response time without regard to their variability.

Lastly, instead of using the parameters of the fit to assess optimality, we shifted the raw (normalized) interresponse times until the resulting reinforcement rate was maximized. This analysis includes all the interresponse times, including the untimed responses. For rats, the shifts were **7**: 0.13 ± 0.01 (0.12 ± 0.06), **10**: 0.16 ± 0.01 (0.16 ± 0.05), **14**: 0.22 ± 0.03 (0.22 ± 0.12), and **28**: 0.36 ± 0.08 (0.31 ± 0.21), **56**: 0.66 ± 0.10 (0.67 ± 0.59). For humans, the shifts were **5**: 0.05 ± 0.01 (0.05 ± 0.03), **8**: 0.05 ± 0.01 (0.05 ± 0.04), **10**: 0.09 ± 0.01 (0.08 ± 0.02), **12**: 0.14 (0.14), **15**: 0.015 (0.015). To give some intuition about what a shift of, for example, 0.15 means, note that the optimal shift is zero and the maximal positive shift is one (since at that point, every single data point is greater than the normalized schedule of one). For a shift of 0.15, this amounts to needing to shift the interresponse times by 15% in order to maximize the reinforcement rate. In humans, the shift did not significantly depend on the schedule, $F(1,12)=0.62$, $p=0.45$, $R^2=0.05$. But it did in rats ($F(1,55)=64.86$, $p=8 \times 10^{-11}$, $R^2=0.54$), even when the longer schedules were not included (although with a smaller effect size, $F(1,33)=10.6$, $p=0.01$, $R^2=0.24$). The shift is significantly positive in all groups in both rats and humans (smallest effect size in Rats (DRL 56s): $t(11)=-3.41$, $p=0.01$, $d=$

–0.31; Humans (DRL 8s): $t(3) = -84.09$, $p = 4 \times 10^{-6}$, $d = -28.03$). These results show what we already showed above in various analyses: that the longer the schedule, the more suboptimal the interresponse times. In fact, this is just another way to measure the bias optimal performance curve shown in Figure 5. But in Figure 5 we used the parameter fits to compute it, and here we used the raw data.

Discussion

Our model was designed to be as simple as possible. The model is a set of three key assumptions. The first is that animals try to maximize reinforcement rate. This imposed the restriction that our model use the equation for the reinforcement rate (Equation 1). The second assumption is that animals do not control their variability. This imposed the restriction that γ could not be optimized. The last assumption is that animals do control their mean interresponse time. Together, these three assumptions imposed the final form of the model: maximize the reinforcement rate equation with respect to t constrained by γ .

We also assumed, for simplicity, that the subjective value of a reinforcement does not change over time. This restriction can be relaxed—we do not see it as a key assumption of our model—but we do not do so here because it requires adding assumptions about how subjective value changes over time, and would likely add free parameters to our model.

Nevertheless, our model captures some key aspects of performance on the task. The model gives the mean interresponse time that maximizes the reinforcement rate given the coefficient of variation. This specifies an ideal (reinforcement-maximizing) relationship between the mean and the variability. This relationship was robust in our data (except for the longer intervals in rats). And the optimal interresponse time (t_o) predicted the actual interresponse times quite well in a simple regression ($R^2 > 0.60$). The model generates an optimal benchmark for both the mean interresponse time and the reinforcement earned, and the animals in our experiment, both humans and rats, were roughly ~90% or more toward that optimal benchmark.

To see if the same results were obtained in different labs, we performed simple secondary data analysis on DRL data sets in rats from three different labs (Orduña, Lourdes, & Bouzas, 2009; Sanabria & Killeen, 2008; Taylor & Balsam, unpublished). In the combined data set, the target intervals were 5s, 10s, and 16s. And in each case, the performance was close to optimal; the rats median gain ratios were 93%, 94%, and 92%, respectively.

What does a generative model give us that an empirical description cannot?

We noted throughout this article that there were at least two ways to go from the raw interresponse times to measures that can be compared to the model's predictions. The first way is to use the raw interresponse time distribution and shift it until reinforcement rate is maximized. This method is computationally efficient and has the nice property of working directly with the actual interresponse times at all times. The second, and our focus, was to build a generative model of the interresponse time distribution, fit it to the raw interresponse times, and analyze the parameters of the fit. The drawback of this approach is that parameter estimates are analyzed, rather than the raw data itself.

But what we gain are summary measures of the data that give rise to theoretical questions. Our data, and others like it, suggested two processes generated the data. A fast untimed one, which we modeled as an exponential distribution, and a slow timed one, which we modeled as an inverse Gaussian distribution. The inverse Gaussian distribution fit our data better than the Gaussian assumption. This is in line with the recent push to bring the drift-diffusion model of perceptual decision making into the timing domain (Patrick Simen, Balci, de Souza, Cohen, & Holmes, 2011; see also Brunton, Botvinick, & Brody, 2013 for a different model in a related context). Here, the decision to respond is based on accumulating evidence toward a single response threshold (similar to the way SET is conceptualized; Gibbon, Church, & Meck, 1984) where the accumulation rate is related to the target interval (similar to the way BET is conceptualized; Killeen & Fetterman, 1988; P Simen et al., 2013). The drift-diffusion model predicts inverse Gaussian response times (Luce, 1986). The inverse Gaussian assumption is also used as the interresponse time distribution in Modular Theory (Guilhardi, Yi, & Church, 2007).

Fitting a generative model to our interresponse time data also allowed us to consider the fast, seemingly untimed responses in a principled way (as others had; e.g., Richards et al., 1993; Sanabria & Killeen, 2008). We modeled these as an exponential distribution which embodies the notion that animals emit operant responses at a random rate. Informally, we did not find anything that predicted the random rate parameter at a subject specific level, and it did not significantly change based on the target interval. These responses are not optimal in the sense described by our model, since all they do is reduce the probability of reinforcement and skew the total mean interresponse times toward even faster-than-optimal response rates. But that does not mean they are completely without purpose. The simplest explanation may be that they are exploratory responses. A few, quick responses may be cheap compared to the potential loss incurred by failing to detect a change in the environment (Cohen, McClure, & Yu, 2007). Of course, our data cannot rule out the hypothesis that they are impulsive and maladaptive, nor can we rule out the hypothesis that animals just have a tendency to respond in stereotypical bouts in many task settings (Kirkpatrick, 2002), which is simply too engrained to easily overcome.

Deviations from Our Optimal Model

There are two glaring deviations from optimality in our data set. The first is that the rats with longer target intervals show no signs of optimality. Their mean interresponse time was not significantly correlated with their variability, and they did not seem to track the optimal interresponse time. For many of the measures we discussed, pooling the rats in the longer DRL schedules in with the shorter target schedules resulted in a skewed picture of the performance in shorter schedules. Our results suggest the rats should be split into two groups: those with shorter target intervals that showed the necessary characteristics of optimality, and those with longer target schedules that showed no signs of optimality. Why this seemingly arbitrary split? We speculate that the culprit is buried in the learning results: It took the rats in the DRL 28s and 56s group far longer to learn to withhold responding. Indeed, it is unclear if they ever reached asymptote in the 40 days of training in our experiment. Given that the rats in those groups were rewarded several hundred times during the experiment (**28**: 607 ± 23 , **56**: 222 ± 19), and that learning temporal intervals often

requires only a few trials (Drew, Zupan, Cooke, Couvillon, & Balsam, 2005), we find it unlikely that those rats did not adequately learn the fixed DRL requirement; although, we cannot rule it out. Instead, we believe the issue is that it takes considerable time to learn to withhold a response that results in reinforcement, and there may be schedules for which this is extremely difficult (e.g., Richardson & Loughhead, 1974; Doughty & Richards, 2002). Machado & Rodrigues (2007) present a process model, designed for a similar task, in which two of the parameters together determine how much a reinforcement will change behavior. According to this view, the longer schedules in our study may result in reinforcement that is less effective at driving behavioral change.

The second glaring deviation of our data set is that 68 of the 71 subjects responded at a faster than optimal rate. Figures 4–6 suggest that simply subtracting a constant from the optimal time predicts performance relatively well. The constant we fit depended on species (Rats: -0.15 ± 0.01 ; Humans: -0.07 ± 0.02) but was a strikingly reliable result in our data. This is in contrast to our recent paper that showed no systematic deviation from optimality in humans (Çavdarolu et al., 2014). There are a few methodological differences between the studies. For instance, the participants in that study were not given any temporal information about the task structure, nor were they pretrained with any intervals. Further, the participants in that study were trained for 8 sessions with the same target interval, while the participants in the present study were only trained for a single session. We cannot rule out the possibility that the bias would have decreased or disappeared in this data as well. Further studies are needed to determine the sources of these systematic deviations from optimality.

Conclusion

The take-home message is that animals choose a response rate that largely maximizes reinforcement rate on this task, except that they are consistently faster than optimal. Our simplistic model is a good starting place for more in-depth models that can account for these deviations. For example, it may be that animals put a premium on speed over accuracy in this task that simply did not go away with training, unlike in other tasks (Balci et al., 2011; Patrick Simen et al., 2009).

We chose to remain silent about the underlying mechanism(s) that might give rise to the behavior. In Marr's (1982) levels of analysis, our model is firmly in the computational camp. Our model specifies a set of goals and constraints for an organism (maximizing the reinforcement rate given timing variability). These assumptions give rise to precise predictions about what we should see in the averaged data. There are likely to be many algorithmic (e.g., process) models that could give rise to both the optimal behavior and the actual behavior. For example, a reinforcement learning approach would posit that the animal's interresponse times are adjusted by hill-climbing a reinforcement rate gradient (Sutton & Barto, 1998). A representational approach would posit that the animal keeps track of the task relevant parameters (i.e., the target interval, and the statistics of the interresponse time distribution) and computes the optimal aim point (Trommershäuser et al., 2006). In each of these approaches, there are several ways to formulate the problem. And, of course, there are other approaches altogether. In turn, these algorithmic models constrain and are constrained by the underlying neural circuitry.

In contrast, our goal was to further explore the notion that both rats and humans choose a response rate that maximizes the reinforcement rate (regardless of the mechanism). The DRL task is particularly well designed to tackle this problem because the task is relatively simple and the interresponse time statistics are well defined, making it straightforward to write out the equations that relate the response rate to the reinforcement rate. Here, we showed that both rats and humans show some of the key characteristics of optimality. Expanding this approach to other tasks, as well as focusing on the deviations from optimality in this task, may give insight into the underlying mechanisms that give rise to response rates.

Acknowledgments

We thank Federico Sanabria, Vladimir Orduna, and Peter Balsam for sharing data with us. The results from the analysis of this data are presented in the discussion. We also thank Laura deSouza for aiding us in data collection, and Andra Geana for invaluable discussions on earlier versions of the manuscript. Lastly, we thank two anonymous reviewers for their helpful suggestions. Parts of this work was supported by an FP7 Marie Curie PIRG08-GA-2010-277015, TÜB TAK 1001 #111K402, and a BAGEP Grant from Bilim Akademisi - The Science Academy, Turkey to FB.

References

- Balci F, Freestone D, Gallistel C. Risk assessment in man and mouse. *Proceedings of the National Academy of Sciences of the United States of America*. 2009; 106(7):2459–2463. [PubMed: 19188592]
- Balci F, Freestone D, Simen P, Desouza L, Cohen JD, Holmes P. Optimal temporal risk assessment. *Frontiers in Integrative Neuroscience*. 2010;5. [PubMed: 20428494]
- Balci F, Simen P, Niyogi R, Saxe A, Hughes J, Holmes P, Cohen J. Acquisition of decision making criteria: reward rate ultimately beats accuracy. *Attention, Perception & Psychophysics*. 2011; 73(2):640–657.
- Baum WM. The correlation-based law of effect. *Journal of the Experimental Analysis of Behavior*. 1973; 20(1):137–153. [PubMed: 16811687]
- Baum WM. On two types of deviation from the matching law: bias and undermatching. *Journal of the Experimental Analysis of Behavior*. 1974; 22:231–242. [PubMed: 16811782]
- Baum WM. Optimization and the matching law as accounts of instrumental behavior. *Journal of the Experimental Analysis of Behavior*. 1981; 36(3):387–403. [PubMed: 16812255]
- Blough DS. Stimulus generalization as signal detection in pigeons. *Science*. 1967; 158:940–941. [PubMed: 6054168]
- Brainard DH. The Psychophysics Toolbox. *Spatial Vision*. 1997; 10:433–436. [PubMed: 9176952]
- Brunton B, Botvinick M, Brody C. Rats and humans can optimally accumulate evidence for decision-making. *Science*. 2013; 340(6128):95–98. [PubMed: 23559254]
- Çavdarolu B, Zeki M, Balci F. Time-based reward maximization. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 2014; 369(1637):20120461.
- Cohen J, McClure S, Yu A. Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*. 2007; 362(1481):933–942.
- Doughty AH, Richards JB. Effects of reinforcer magnitude on responding under differential-reinforcement-of-low-rate schedules of rats and pigeons. *Journal of the Experimental Analysis of Behavior*. 2002; 78(1):17–30. [PubMed: 12144310]
- Drew MR, Zupan B, Cooke A, Couvillon PA, Balsam PD. Temporal control of conditioned responding in goldfish. *Journal of Experimental Psychology Animal Behavior Processes*. 2005; 31:31–39. [PubMed: 15656725]
- Gibbon J. Scalar expectancy theory and Weber's law in animal timing. *Psychological Review*. 1977; 84(3):279.

- Gibbon J, Church R, Meck W. Scalar timing in memory. *Annals of the New York Academy of Sciences*. 1984; 423:52–77. [PubMed: 6588812]
- Gibbon J, Malapani C, Dale CL, Gallistel C. Toward a neurobiology of temporal cognition: advances and challenges. *Current Opinion in Neurobiology*. 1997; 7(2):170–184. [PubMed: 9142762]
- Girshick, MA.; Blackwell, D. *Theory of games and statistical decisions*. London: John Wiley & Sons, Inc; 1954.
- Guilhardi P, Yi L, Church RM. A modular theory of learning and performance. *Psychonomic Bulletin & Review*. 2007; 14(4):543–559. [PubMed: 17972717]
- Herrnstein RJ. Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*. 1961; 4(3):267–272. [PubMed: 13713775]
- Herrnstein RJ, Heyman GM. Is matching compatible with reinforcement maximization on concurrent variable interval variable ratio? *Journal of the Experimental Analysis of Behavior*. 1979; 31(2): 209–223. [PubMed: 16812126]
- Heyman GM, Luce RD. Operant matching is not a logical consequence of maximizing reinforcement rate. *Animal Learning & Behavior*. 1979; 7(2):133–140.
- Heyman GM, Tanz L. How to teach a pigeon to maximize overall reinforcement rate. *Journal of the Experimental Analysis of Behavior*. 1995; 64(3):277–297. [PubMed: 16812773]
- Houston AI, McNamara J. How to maximize reward rate on two variable-interval paradigms. *Journal of the Experimental Analysis of Behavior*. 1981; 35(3):367–396. [PubMed: 16812223]
- Kheifets A, Gallistel CR. Mice take calculated risks. *Proceedings of the National Academy of Sciences of the United States of America*. 2012; 109(22):8776–8779. [PubMed: 22592792]
- Killeen PR, Fetterman JG. A behavioral theory of timing. *Psychological Review*. 1988; 95(2):274. [PubMed: 3375401]
- Kirkpatrick K. Packet theory of conditioning and timing. *Behavioural Processes*. 2002; 57(2):89–106. [PubMed: 11947991]
- Luce, RD. *Response Times: Their Role in Inferring Elementary Mental Organization*. Oxford University Press; 1986.
- Machado A, Rodrigues P. The differentiation of response numerosities in the pigeon. *Journal of the experimental analysis of behavior*. 2007; 88(2):153–178. [PubMed: 17970413]
- Marr, D. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. New York: Freeman; 1982.
- Miller RR, Barnet RC, Grahame NJ. Assessment of the Rescorla-Wagner model. *Psychological Bulletin*. 1995; 117:363–386. [PubMed: 7777644]
- Niv Y, Daw ND, Dayan P. How fast to work: Response vigor, motivation and tonic dopamine. *NIPS Proceedings*. 2005; 18:1019–1026.
- Niv Y, Daw N, Joel D, Dayan P. Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology*. 2007; 191(3):507–520. [PubMed: 17031711]
- Niv Y, Joel D, Dayan P. A normative perspective on motivation. *Trends in Cognitive Sciences*. 2006; 10(8):375–381. [PubMed: 16843041]
- Orduña V, Lourdes VT, Bouzas A. DRL performance of spontaneously hypertensive rats: dissociation of timing and inhibition of responses. *Behavioural Brain Research*. 2009; 201(1):158–165. [PubMed: 19428629]
- Pearce JM, Bouton ME. Theories of associative learning in animals. *Annual review of psychology*. 2001; 52(1):111–139.
- Rachlin H. A molar theory of reinforcement schedules. *Journal of the Experimental Analysis of Behavior*. 1978; 30:345–360. [PubMed: 16812114]
- Richards J, Sabol K, Seiden L. DRL interresponse-time distributions: quantification by peak deviation analysis. *Journal of the Experimental Analysis of Behavior*. 1993; 60(2):361–385.10.1901/jeab.1993.60-361 [PubMed: 8409824]
- Richardson WK, Loughhead TE. Behavior under large values of the differential-reinforcement-of-low-rate schedule. *Journal of the experimental analysis of behavior*. 1974; 22(1):121–129. [PubMed: 16811770]

- Sanabria F, Killeen P. Evidence for impulsivity in the Spontaneously Hypertensive Rat drawn from complementary response-withholding tasks. *Behavioral and Brain Functions: BBF*. 2008;4. [PubMed: 18234079]
- Schoenfeld WN, Harris AH, Farmer J. Conditioning response variability. *Psychological Reports*. 1966; 19(2):551–557. [PubMed: 5912818]
- Simen P, Balci F, de Souza L, Cohen J, Holmes P. A model of interval timing by neural integration. *The Journal of Neuroscience*. 2011; 31(25):9238–9253. [PubMed: 21697374]
- Simen P, Contreras D, Buck C, Hu P, Holmes P, Cohen J. Reward rate optimization in two-alternative decision making: empirical tests of theoretical predictions. *Journal of Experimental Psychology Human Perception and Performance*. 2009; 35(6):1865–1897. [PubMed: 19968441]
- Simen P, Rivest F, Ludvig EA, Balci F, Killeen P. Timescale invariance in the pacemaker-accumulator family of timing models. *Timing and time perception*. 2013; 1(2):159–188.
- Staddon JE, Motheral S. On matching and maximizing in operant choice experiments. *Psychological Review*. 1978; 85(5):436–444.
- Sutton, RS.; Barto, AG. *Introduction to reinforcement learning*. MIT Press; 1998.
- Trommershäuser J, Landy MS, Maloney LT. Humans rapidly estimate expected gain in movement planning. *Psychological Science*. 2006; 17(11):981–988. [PubMed: 17176431]
- Taylor K, Balsam P. unpublished raw data.
- Wald, A. *Statistical decision functions*. Oxford: John Wiley & Sons; 1950.
- Wearden JH. Maximizing reinforcement rate on spaced-responding schedules under conditions of temporal uncertainty. *Behavioural processes*. 1990; 22(1):47–59. [PubMed: 24896922]
- Wilson MP, Keller FS. On the selective reinforcement of spaced responses. *Journal of Comparative and Physiological Psychology*. 1953; 46(3):190–193. [PubMed: 13061646]
- Yi L. Do rats represent time logarithmically or linearly? *Behavioural Processes*. 2009; 81:274–279. [PubMed: 19007867]
- Zeiler MD, Scott GK, Hoyert MS. Optimal temporal differentiation. *Journal of the Experimental Analysis of Behavior*. 1987; 47:191–200. [PubMed: 16812476]

Figure 1. The model

Panel a (top) shows how the optimal model is constructed. The gray vertical line at 1 shows the target interval. The rising dashed line shows the probability of a reward, given the statistics of the interresponse time distribution (the mean, t , and coefficient of variation, γ). The falling dashed line shows how much the reinforcement rate declines simply because of the passage of time ($1/t$). The solid black curve is the element-wise multiplication of the two dashed lines, and gives the reinforcement rate for a single coefficient of variation for any interresponse time ranging from 0 to 3 (i.e., 0 to 3 times the target interval). The light dashed horizontal and vertical lines intersect at the maximum of the reinforcement rate curve,

showing the optimal interresponse time. Panel b shows 3 more reinforcement rate curves ($\gamma=0.1, 0.25, \text{ and } 0.50$). The dashed line goes through the maximum of all possible reinforcement rate curves when γ is varied. It describes the optimal relationship between the mean and variability of the interresponse time distribution.

Figure 2. Example subjects

Panels a and b (first column) show the interresponse times of a single randomly chosen rat from the DRL 10 second group (indicated by the vertical dashed line at 10). Panel b (bottom left) shows every raw interresponse time (x-axis) as a function of response number (y-axis) for the last 10 sessions of this rat. Panel a (top left) shows the probability density computed from the raw interresponse times shown in panel b. The solid curve through the data is the best fitting exponential-inverse Gaussian mixture distribution, and the dashed curve through the data is the best fitting exponential-Gaussian mixture distribution. Panels c and d (right column) show a randomly chosen human from the DRL 10s group.

Figure 3. Learning in rats

Panel a (top) shows the mean coefficient of variation for each of the five groups of rats. Each data point was computed by fitting the exponential-inverse Gaussian model every session for every rat, and then averaging over rats in a group for a given session with standard error bars. Panel b (middle) shows the proportion of untimed responses declined over sessions. And panel c (bottom) shows the mean interresponse time from the inverse Gaussian fit over sessions. For shorter schedules (DRL 7s, 10s, and 14s) the rats learn to withhold responding within a few sessions. For the longer schedules (DRL 28s and 56s), the rats take considerably longer to learn to withhold responding and it is unclear that they ever reached steady state (see text for details).

Figure 4. Interresponse times by schedule

The black dots in top row (panels a and c) show the mean interresponse times, from the inverse Gaussian fit, averaged over rats in a group (mean $\pm$ sem). The black dashed line is the best fitting line. The gray dots and dashed line show the optimal interresponse time for that group (the average of the optimal for each rat in a group). The bottom row (panels b and d) shows the same thing, except normalized by the schedule. The left column shows the rats, the right column shows the humans. Note that both the x- and y-axes are different on the four plots.

Figure 5. Interresponse time as a function of the coefficient of variation

Panel a (top) shows the normalized interresponse time from the fit against the coefficient of variation (from the model fit) for every rat in the experiment. The dashed horizontal line shows the normalized target interval. The solid curve is the optimal performance curve (i.e., it is the same curve as the dashed optimal performance curve in figure 1b, just displayed on different axes). The dashed curve through the data shows how far the optimal performance curve has to shift in order to fit the data (it is an index of optimality such that values closer to zero show more optimality). The open circles show the rats from the DRL 28s and 56 s groups, since their performance differs from the shorter intervals in many ways (see text for details). Panel b (bottom) shows the results for humans.

Figure 6. Interresponse time as a function of the optimal interresponse time

Panel a (top) shows the normalized interresponse time as a function of the optimal interresponse time for every rat in the experiment. The vertical and horizontal dashed lines show the target interval at 1 (a.u). The diagonal line is the identity line. Points below this line show faster than optimal responses, and points above this line show slower than optimal responses (points on the line show optimal responses). The rats with longer target intervals (DRL 28 and 56s) are still shown as open circles. Panel b shows the results for humans. The normalized interresponse times increases as the optimal increases. But nearly all data points are below the line.

Table 1

Parameter estimates of Equation 6

The rows are the parameters (the proportion of untimed responses, p , the rate parameter of the untimed responses, κ , the mean of the timed responses, t , and the coefficient of variation, γ . The first five columns are the groups (DRL schedule) and the elements in those columns are the mean (+ sem) parameter estimates. The sixth column is the F statistic of the ANOVA comparing the estimate over groups, and the last column is the p value of the statistics (effect sizes are given in the text). The table is split by Rats (top) and Humans (bottom).

Rats	n=12	n=12	n=11	n=10	n=12
schedule	7	10	14	28	56
p	0.29 (0.04)	0.21 (0.02)	0.17 (0.04)	0.17 (0.02)	0.22 (0.03)
κ	10.75 (0.88)	9.31 (1.05)	7.56 (1.90)	2.57 (1.20)	9.18 (3.44)
t	8.54 (0.19)	11.2 (0.16)	17.0 (0.53)	30.9 (1.22)	51.8 (3.31)
γ	0.23 (0.02)	0.19 (0.01)	0.32 (0.04)	0.40 (0.03)	0.42 (0.02)
Humans	n=4	n=4	n=4	n=1	n=1
schedule	5	8	10	12	15
p	0.02 (0.01)	0.003 (0.003)	0.008 (0.004)	3.4×10^{-11} (n/a)	0.007 (n/a)
κ	26.5 (24.5)	20.9 (13.6)	16.84 (14.75)	87.0 (n/a)	0.001 (n/a)
t	5.78 (0.18)	9.92 (0.96)	11.92 (0.21)	14.37 (n/a)	19.07 (n/a)
γ	0.13 (0.02)	0.16 (0.06)	0.19 (0.01)	0.22 (n/a)	0.20 (n/a)