J Neurophysiol 119: 413–421, 2018.
First published November 8, 2017; doi:10.1152/jn.00634.2017.

RESEARCH ARTICLE | *Control of Movement*

Control of saccadic latency in a dynamic environment: allocation of saccades in time follows the matching law

Cécile Vullings[1] **and Laurent Madelain**[1,2]

[1]*CNRS, CHU Lille, UMR 9193, SCALab Sciences Cognitives et Sciences Affectives, Université de Lille, Lille, France; and*
[2]*CNRS, Institut de Neurosciences de la Timone, UMR 7289, Aix Marseille Université, Marseille, France*

Submitted 27 August 2017; accepted in final form 27 October 2017

Vullings C, Madelain L. Control of saccadic latency in a dynamic environment: allocation of saccades in time follows the matching law. *J Neurophysiol* 119: 413–421, 2018. First published November 8, 2017; doi:10.1152/jn.00634.2017.—When exploring the visual environment, one uses saccades to shift gaze and fixation to gather spatially and temporally localized information. We propose that the temporal structure of our environment should constrain the temporal allocation of saccades. Here we probe the possibility of learning to control saccadic latencies in a choice paradigm. Six participants made saccades within 80–300 ms following a target horizontally stepping by 10° between two fixed locations. For each participant we constructed two classes of latencies, "short" and "long," using the first and last quartiles of the individual baseline distribution (e.g., [80;152] ms and [185;300] ms, respectively). We then concurrently reinforced each class in three blocked conditions across ~60 experimental sessions per participant, using different reinforcement probabilities such that the relative ratio of reinforcement rates for short vs. long latencies was 9/1, 1/9, or 1/1. Latency distributions followed the reinforcement conditions: distributions shifted toward the shorter or longer values or became strongly bimodal. Moreover, the relative rates of short over long latencies matched the relative rates of reinforcers earned for the corresponding latencies (slope up to 0.95), which reveals the ability to choose when to saccade. Our results reveal that learned contingencies considerably affect the allocation of saccades in time and are in line with recent studies on the temporal adjustment of behavior to dynamic environments. This study provides strong evidence for fine operant control of saccadic latency, supporting the hypothesis of a cost-benefit control of saccade latencies.

NEW & NOTEWORTHY Saccades may be regarded as an information-foraging behavior mostly concerned with the spatial localization of objects, yet our world is dynamic and environmental temporal regularities should also affect saccade decisions. We present behavioral data from a choice task establishing that humans can learn to choose their saccadic latencies depending on the reinforcement contingencies. This suggests a cost-benefit-based policy that takes into account the learned temporal properties of the environmental contingencies for controlling saccade triggering.

decision; latency; learning; matching law; saccade

INTRODUCTION

Resources are not evenly distributed but instead are localized, and to successfully interact with their environment animals must engage in foraging activities that result in finding food or a potential mate but also avoiding predators or aversive stimuli. Importantly, these requirements also apply to stimuli signaling the availability of resources: information is localized as well. This is particularly true for visual information, and the necessity to orient toward stimuli provides strong evolutionary pressures that have shaped the eye movement systems of animals depending on the visual neural structures and behavioral needs of each species (Krauzlis 2008; Land 2011). In primates an additional requirement arises from the presence of a restricted high-photoreceptor-density region of the retina, the fovea, and a corresponding increased functional acuity: we use saccades to shift gaze and fixation in combination to gather information when exploring our environment.

However, another critical feature of our environment is that it is dynamic and constantly changing: information is therefore also localized in time, and because the fraction of our visual surroundings perceived at any given moment is limited we must decide when as much as where to look, a fact that has been mostly overlooked so far (Hoppe and Rothkopf 2016). That visual information-foraging activity is most certainly dependent on both the spatial and temporal structure of our environment is indeed striking. To borrow an image from Hoppe and Rothkopf (2016), consider a person exploring the visual environment before crossing a busy street: the observer's behavior will be driven by environmental regularities to select informative spatial locations such as the street itself and ignore others such as the sky or the buildings. Because pedestrians, bicycles, and cars are obstacles with different temporal properties, the observer might also use these regularities to sequentially allocate his/her limited visual resources over time. One may therefore postulate that we learn and use the temporal properties of our environments to shift gaze toward potentially informative locations at the right time. In other words, the temporal organization of our environment should constrain the temporal allocation of saccades.

These environmental temporal properties might have different forms. On the one hand, environmental temporal statistics, such as event durations, might be exploited: Hoppe and Rothkopf (2016) elegantly demonstrated that observers learned to efficiently shift gaze rapidly across two locations or to maintain

---

Address for reprint requests and other correspondence: L. Madelain, Université Charles-de-Gaulle Lille III, Domaine universitaire du Pont de Bois, BP 149, 59653 Villeneuve d'Ascq Cedex, France (e-mail: laurent.madelain@univ-lille3.fr).

longer fixations depending on the durations of events occurring at each location. For instance, there is no need to gaze at a traffic light that has just turned red, as it is highly probable that it will remain red for some time. In this case, learning the temporal structure of the environment is used to gather information efficiently. On the other hand, events in the environment might signal when gaze should be shifted: it might be informative to shift gaze toward the traffic light right away if you detect that the cars to the right of the intersection stop. In this case, learning the temporal structure of the environment reinforcement contingencies, that is, the temporal relations between external events, behavior, and consequences, is used to control gaze allocations. Here we focus on this second type of temporal environmental structure. However, one might question whether these differences are of functional significance, as in both cases it appears that one might detect temporal regularities in the occurrence of events, learn them, and use them to choose when to allocate visual resources to gather information regarding the state of the environment.

Saccadic reaction times (SRTs) have become one of the main behavioral measures to quantitatively study decision processes and are conventionally viewed as reflecting the accumulation of information during decision-making process (see Gold and Shadlen 2007 for a review). A critical feature of these models is that the SRT is viewed as a by-product of an underlying decision process: one decides where but not when to move the eyes because a saccade is triggered as soon as enough evidence has been accumulated. Departing from this basic assumption of evidence-accumulation models, we asked whether SRTs might be allocated depending on the temporal structure of environmental contingencies. We considered a decision-making task in which participants were offered rewards contingent on their saccade latencies. A critical component of our experiments was that we manipulated the probability of rewarding saccades with short or long latencies such that participants were free to choose their SRT on each trial according to the actual reinforcement contingencies in force.

To investigate the adaptation of SRTs to the temporal properties of the environmental contingencies we designed a task with four critical features. First, alternating between two fixed target locations isolated the temporal aspects of saccades from spatial selection and visual information other than target step. Second, we used the actual individual latencies to define “short” and “long” SRTs, to avoid imposing arbitrary constraints on SRT distributions. Third, we designed a temporal-choice task to probe the extent of control over SRTs, allowing us to quantify the allocation of latencies in response to the temporal structure of the reinforcement contingencies. Finally, we reinforced specific latencies to probe whether the law of effect (Thorndike 1927) could apply to SRTs. Our results clearly demonstrate that temporal regularities in the environmental contingencies consistently control SRTs, suggesting that observers learn to alter the triggering of saccadic eye movements in response to new environmental contingencies.

METHODS

Participants

Six adults (2 men and 4 women; mean age = 26.7 yr, age range: 18–45 yr) participated in this study. They were naive as to the purpose of the study, except for two participants (i.e., the authors, *S1* and *S2*, who experienced several pilot procedures before the actual experiment), and had normal or corrected-to-normal vision. Participants were informed that the experiment was about short and long reaction times and instructed to earn as many points as possible; no further explanation was given as to how to earn points. When the experimental conditions changed, the same instruction was given again. Naive participants received €30 for participating plus an additional sum depending on the points collected (1 point equals 2 cts); they received €94 on average at the end of the experiment. They were informed that they could earn up to 120 points per session, with a monetary bonus of €1 for each session in which their score exceeded 100. All experimental procedures received approval from the Ethical Committee in behavioral sciences of the University of Lille (Agreement No. 2015-1-S34) and conformed to the standards set by the Declaration of Helsinki. All participants gave written informed consent.

Apparatus

Stimuli were generated with the Psychophysics Toolbox extensions (Brainard 1997; Pelli 1997) for MATLAB and displayed on a video monitor (Iiyama HM204DT, 100 Hz, 22 in.). Participants were seated on an adjustable stool in a darkened quiet room, facing the center of the computer screen at a viewing distance of 60 cm. To minimize measurement errors, the participant’s head movements were restrained with a chin and forehead rest, so that the eyes in the primary gaze position were directed toward the center of the screen. Viewing was binocular, but only the right eye position was recorded and digitized in both the vertical and horizontal axes. Eye movements were measured continuously with an infrared video-based eye tracking system (EyeLink, SR Research), sampled at 2,000 Hz. Data were transferred, stored, and analyzed via programs written in MATLAB running on an Ubuntu Linux computer.

Stimuli were light gray disks (luminance = 15.99 cd/m^2), diameter 0.5°, displayed on a dark gray background (luminance = 1.78 cd/m^2). The target position alternated between two fixed locations on the screen separated by 10° horizontally around the center of the monitor.

Before each experimental session, we calibrated the eye tracker by having the participant fixate a set of 13 fixed locations distributed across the screen. Every 50 trials, participants looked at a target displayed in the center of the screen for a 1-point calibration check.

Procedure

The experiment lasted on average sixty 400-trial sessions divided among familiarization, baseline, and reinforcement sessions. Two participants also did four latency training sessions between baseline and reinforcement (see below). Three daily sessions were typically recorded, separated by 5-min breaks during which participants were free to move. On average, the experiment lasted 20 consecutive days (5 days a week, from Monday to Friday). Regardless of the actual condition, participants were asked to make saccades toward the horizontally stepping target.

Familiarization. Participants who had no previous experience in oculomotor experiments were trained to make saccades without blinking, anticipating, or looking away from the target. The 100-trial familiarization sessions lasted until at least 90% of saccades were correctly detected online (see *Acquisition and Data Analysis*) and were identical to baseline (see below) except for the number of trials.

Baseline. Seven 400-trial baseline sessions were completed in which participants made saccades toward the target horizontally stepping between the two fixed positions (Fig. 1). At the beginning of the trial, the participant looked at the target for an unpredictable period varying between 400 and 650 ms (sampled from a uniform distribution). The target then stepped horizontally by 10° to the other position. The participant made a saccade with a latency that had to range from 80 to 300 ms. If the latency was outside this range or if no

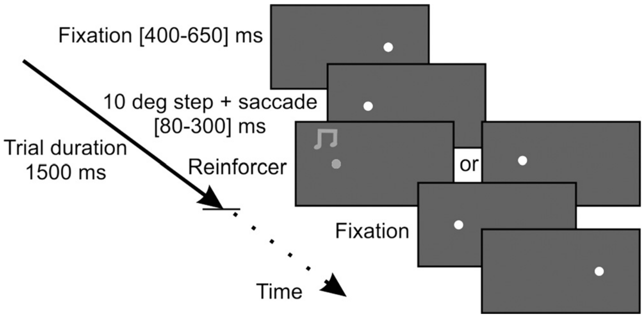

Fig. 1. Experimental design of an ongoing trial during reinforcement conditions.

saccade was detected, the target disappeared and the trial was discarded for off-line analysis. The poststep period depended on the time elapsed since the beginning of the trial, so that each trial lasted 1,500 ms; there was no additional intertrial interval. It should be noted that because the fixation period duration was randomized the timing of the target step was not predictable. The target position became the fixation position for the next trial. Saccadic latency was defined as the interval of time elapsed between the target step and the saccade onset.

After completion of the baseline, we constructed two individual latency classes—"short" and "long" saccadic latencies—using the first and last quartiles, respectively, of baseline latency distributions for each participant. These class boundaries were fixed for the remainder of the experiment. Latencies between 80 ms and the first quartile were categorized as short latencies (i.e., the shortest 25% of all baseline SRTs; e.g., 80–152 ms for *S2*), and latencies between the last quartile and 300 ms were categorized as long latencies (i.e., the longest 25% of all baseline SRTs; e.g., 185–300 ms for *S2*). The first and last quartiles were 164 and 214 ms, 152 and 185 ms, 157 and 185 ms, 171 and 207 ms, 140 and 169 ms, and 119 and 169 ms for *participants S1–S6*, respectively.

Reinforcement sessions. In all of our reinforcement conditions, short and long latencies were independently reinforced on a concurrent random interval (RI) reinforcement schedule in which the relative probabilities of reinforcing short and long latencies were manipulated. The availability of reinforcement is time-controlled with interval schedules: they require a defined elapse of time since the last reinforcer before a response produces reinforcement. Once the interval of time is elapsed, reinforcement is delivered contingent on the first correct response. The RI schedule uses probabilities for assigning reinforcement randomly in time with exponentially distributed interreinforcement intervals (Millenson 1963). Importantly, a reinforcer, once available, remains available until collected.

A concurrent reinforcement schedule is used to study choice between two alternatives by having two reinforcement contingencies operating independently and simultaneously for two responses. Typically, with interval schedules in a concurrent paradigm participants do not respond exclusively in one alternative but rather distribute their choices between the two options because the probability of being reinforced increases with time (Ferster and Skinner 1957).

The reinforcer consisted of having the target turn green for the remainder of the trial (luminance = 5.3 cd/m^2), a brief auditory feedback tone (100 ms, 500 Hz), and earning a point that was exchanged for 2 cts collected at the end of the experiment. The total number of 400-trial reinforcement sessions completed in each condition depended on the time needed to reach stable reaction time distributions. Our stability criterion used the proportion of trials with short vs. long latencies for the last three sessions with the requirement that these proportions should not be different by >1.5 standard deviations from the average of the preceding seven sessions.

Schedule manipulations. Each class of latencies had a probability of being reinforced, and we manipulated the relative frequency of reinforcer availability between the two simultaneous concurrent alternatives such that the ratio of reinforcement rates for short vs. long latencies was 1/1, 9/1, or 1/9. Across conditions, the overall programmed reinforcement rate was always 12 reinforcers per minute (note that because our trial duration was fixed at 1.5 s this reinforcement rate may be discretized in reinforcement probability per trial). In the ratio 1/1, both classes of latencies were reinforced with the same probability; a RI schedule with an interval of 10 s on average was used for both short and long latencies. In the ratio 9/1, reinforcers were more often available for short latencies than for long ones (average intervals for short latencies equaled 5.56 s vs. 50 s for long latencies). In the ratio 1/9, the opposite was true. A changeover delay was used so that a reinforcer, once available, could not be collected with the first response in an alternative after a switch between alternatives. For instance, if a reinforcer for short latencies was available, the participants had to saccade twice in a row with short latencies for the reinforcement to be delivered; if they were to systematically switch between long and short latencies, they would never obtain a reinforcer. This penalty is typically used to eliminate alternation and compel choice (Herrnstein 1961).

Training. To probe whether latency discrimination training affects SRT control, we introduced a four-session training phase following the baseline sessions for the last two participants (*S5* and *S6*). For one participant, in the first session we instructed that short saccadic latencies would be reinforced and we used a continuous reinforcement schedule (i.e., the target turned yellow for every latency within the criteria). The second session was identical but reinforced long latencies. In the third session, only short latencies were intermittently reinforced (average intervals for the reinforcer availability equaled 5.56 s). The fourth session was similar, but long latencies were reinforced. The order between short and long latency reinforcement sessions was counterbalanced across the two participants.

Acquisition and Data Analysis

Eye movements were recorded and measured throughout each trial. For online saccade detection, we used the EyeLink online saccade detector to identify saccade onset and offset, using 30°/s velocity and 8,000°/s^2 acceleration thresholds. Saccade parameters were retrieved on average with a 12-ms delay after saccade offset. For off-line analyses, saccades with amplitude gain lower than 0.5 or duration longer than 100 ms were automatically excluded. A human observer then validated each saccade manually. On average, we kept 97.89%, 93.26%, 95.86%, 98.04%, 95.57%, and 89.74% of saccades for *participants S1–S6*, respectively. We observed very few anticipatory saccades (i.e., saccades with a latency shorter than 80 ms) across the experiment (0.03%, 0.44%, 0.57%, 0.22%, 0.34%, and 1.26% for *participants S1–S6*, respectively).

We used bootstrapping methods (resampling with replacement 100,000 times) to estimate all the individual statistical parameters and 98% confidence intervals (CIs; Efron 1979). Means were compared by Fisher's exact test with 100,000 permutations.

RESULTS

Figure 2*A* plots the saccadic latency distributions for all sessions for all participants; each column of pixels represents one 400-trial session, and the color of the pixel in a given row indicates the frequency of SRTs in that time bin. Importantly, the distributions changed across experimental conditions: in the 1/9 condition in which long latencies were more likely to be reinforced more long SRTs are observed, whereas in the 9/1 condition the opposite is true. Moreover, distributions were much more spread in the 1/1 condition; bimodal distributions can even be observed in some participants (e.g., *S1* or *S5*). It is noteworthy that transitions between conditions were idiosyn-

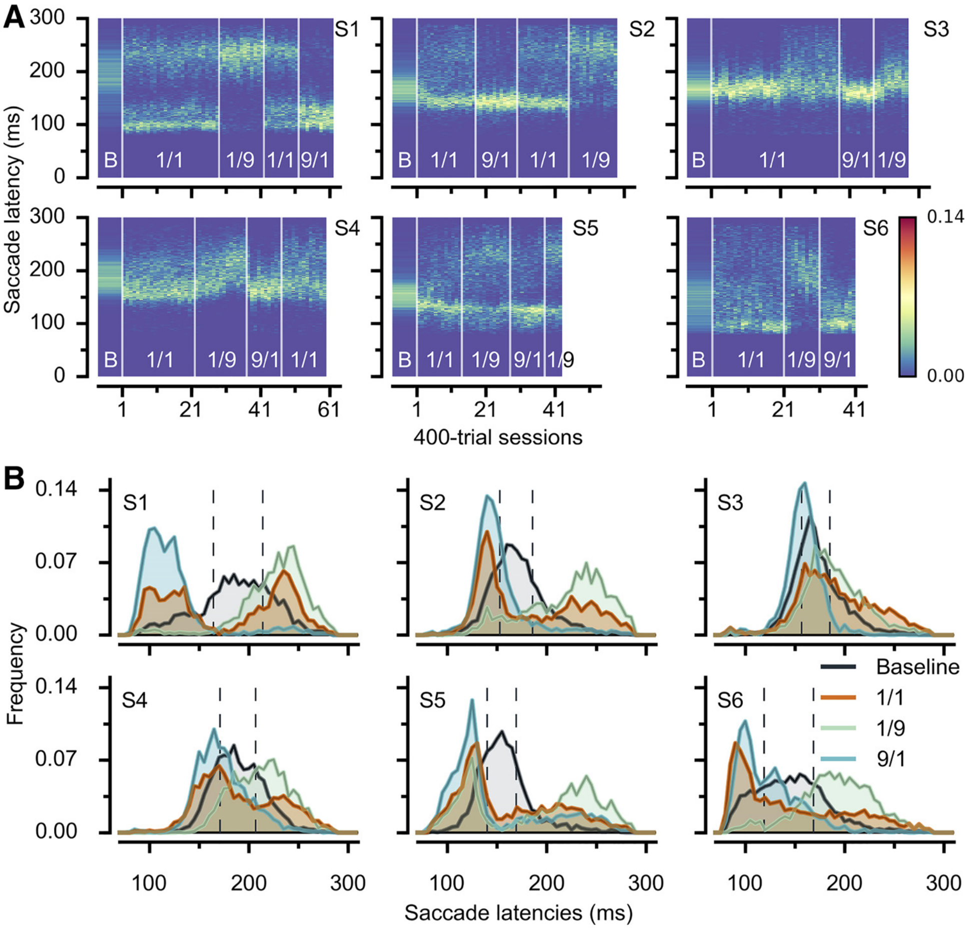

Fig. 2. *A*: saccade latency frequencies across all conditions (B, baseline; 1/1, 1/9, and 9/1, the 3 reinforcement conditions) for all participants in each session. Baseline sessions were combined in a single distribution. Color indicates the frequency of SRTs. *B*: frequency distributions of saccadic latencies for baseline and ratio 1/1, 1/9, and 9/1 experimental conditions for all participants in the last 5 sessions of each condition. Vertical dashed lines plot the first and last individual quartiles of baseline distribution, used to construct the 2 individual classes of latencies. Latencies below the first quartile were categorized as "short" latencies, and those above the last quartile were categorized as "long" latencies.

cratic: progressive transitions were observed for some participants in specific cases (*S4*, *S5*, and *S6*), while other participants tended to show more abrupt transitions (*S1*, *S2*, and *S3*). However, these transition patterns were not systematically observed for a given subject. We now further describe these data considering the steady state (i.e., the last 5 sessions of each experimental condition). Figure 2*B* plots the SRT distributions for the steady state. In comparison to baseline, the distributions strongly drifted toward shorter values when short latencies were more often reinforced (i.e., ratio 9/1). When long latencies were favored (i.e., ratio 1/9), the distributions shifted toward longer values. In addition, for several participants (*S1*, *S2*, *S4*, and *S5*), we obtained a bimodal distribution when both short and long latencies were under identical reinforcement schedules (i.e., ratio 1/1). We observed a large amount of very short latencies (between 80 and 110 ms) for three participants (*S1*, *S5*, and *S6*).

To quantify the changes in the distributions, we computed the respective proportions of long and short latencies with respect to all latencies for each experimental condition in the steady-state sessions (Fig. 3). The gray dashed lines in Fig. 3 show the baseline proportions for short and long latencies (25% each, by construction): data points falling away from these lines indicate a change in the proportion of short or long SRTs with respect to baseline levels. In the 9/1 ratio condition, instead of the baseline 25% the distributions averaged 62.30% (SD = 16.69) of short latencies and 11.34% (SD = 8.98) of long latencies. Conversely, in the 1/9 ratio condition, the distributions averaged 11.85% (SD = 11.46) of short latencies and 65.75% (SD = 11.14) of long latencies. Finally, in the 1/1 ratio condition, the average proportions were 39.88% (SD = 12.43) for short latencies and 39.91% (SD = 6.59) for long latencies. For all participants, we observed significant differences with respect to baseline values: in all three conditions, all

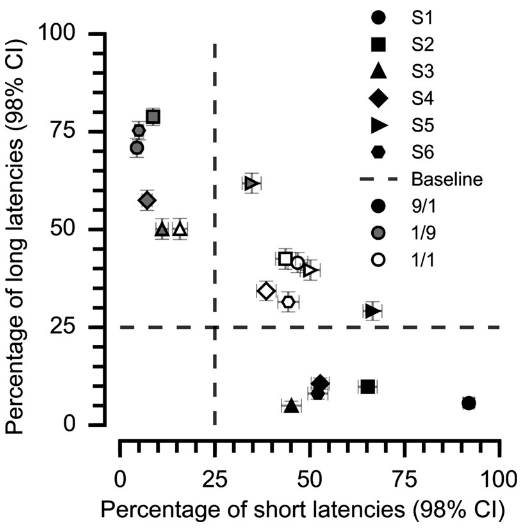

Fig. 3. Percentage of long latencies as a function of percentage of short latencies (with respect to all latencies) for each of the last 5 sessions of the 3 reinforcement conditions for all participants. The bootstrap 98% confidence intervals are shown for both the short and long latency percentages.

the points were different from baseline levels and 25% was never within the CIs. This adjustment to reinforcement contingencies was also apparent in the proportions of intermediate-latency saccades (the central 50% of baseline distributions, i.e., latencies comprised between the first and last quartiles, which were never reinforced; e.g., 153–184 ms for *S2*). Those proportions considerably decreased for the nonnaive *participants S1* and *S2* (13% and 17% on average across conditions, respectively). For naive participants these proportions were 41%, 33%, 6%, and 28% for *participants S3–S6*, respectively. Interestingly, these proportions were lowest for *S5* and *S6*, who received an additional latency-discrimination training phase. This is apparent when comparing the distributions in Fig. 2*B* (e.g., *S4* vs. *S5*).

To estimate how choices of saccadic latencies depended on the actual reinforcement conditions, we used the matching law (Baum 1974; Herrnstein 1961), which states that the relative rate of choices made to an option matches the relative rate of reinforcers earned from that option. This law describes the behavior allocation across reinforcement conditions: the frequency of a behavior occurrence depends on its own reinforcement schedule and on reinforcement frequencies from other activities performed by the individual. Expressed logarithmically, the equation, as proposed by Baum (1974) to account for deviation from strict matching, is

$$\log\left(\frac{B1}{B2}\right) = s \times \log\left(\frac{R1}{R2}\right) + \log b$$

B1 and *B2* represent the number of responses made for each alternative, and *R1* and *R2* correspond to the number of reinforcers respectively obtained for those options; *s* is a measure of the sensitivity of behavior to changes in the relative rate of reinforcement; log *b* is called bias, as it measures a constant proportional preference for one alternative over the other. Figure 4 illustrates the individual data for all participants fitted by the method of least squares (solid line). The relative rates of latencies approximately matched the relative rates of reinforcement, albeit not perfectly ($s = 1$ in case of perfect matching), as measured by the sensitivity of latency choice to changes in reinforcement ratios, which was 0.95, 0.87, 0.58, 0.57, 0.35, and 0.67 for *participants S1–S6*, respectively. The two nonnaive participants (*S1* and *S2*) showed almost a perfect matching, while sensitivity was lower for *S3*, *S4*, and *S6*. This range of deviation is classically observed and has been termed undermatching (see Baum 1979 for a review). The bias remained small, indicating a lack of preference for one of the latency classes (log *b* values were 0.02, 0, −0.12, −0.06, 0.06, and 0.02 for *participants S1–S6*, respectively). Observations were overall well fitted by the linear regression computed over 15 points (r^2 values were 0.98, 0.96, 0.81, 0.96, 0.95, and 0.96 for *participants S1–S6*, respectively). Interestingly, we observed the lowest slope ($s = 0.35$) for *S5* even though she had the lowest overall proportion of intermediate-latency saccades (6%) and the highest reinforcement rate across conditions compared with the other naive participants (8.90 reinforcers/min for *S5* vs. 6.35, 7.40, and 7.95 reinforcers/min for *S3*, *S4*, and *S6*, respectively). One explanation for the large undermatching in *S5* would be that the changeover delay was not costly enough (see Baum 1974 for a review). It is noteworthy that this participant received the additional training phase.

Finally, we asked whether the variations in latencies we have observed were accompanied by systematic changes in other saccade metrics. One could postulate that these variations were correlated with changes in the accuracy of the saccade such that amplitudes might change with latencies. To assess this relation we first probed whether there was any difference in amplitudes and peak velocities as a function of saccade direction, as left-right asymmetries have been reported (e.g., Collewijn et al. 1988; Vergilino-Perez et al. 2012). All participants made saccades with slightly shorter amplitudes toward the left (mean amplitude difference ranging from −0.07° to −0.46°; all values greater than the null hypothesis 98% CIs) and had faster peak velocities toward the right (mean peak velocity difference ranging from −39.99°/s to −70.28°/s; all values greater than the null hypothesis 98% CIs). We found no

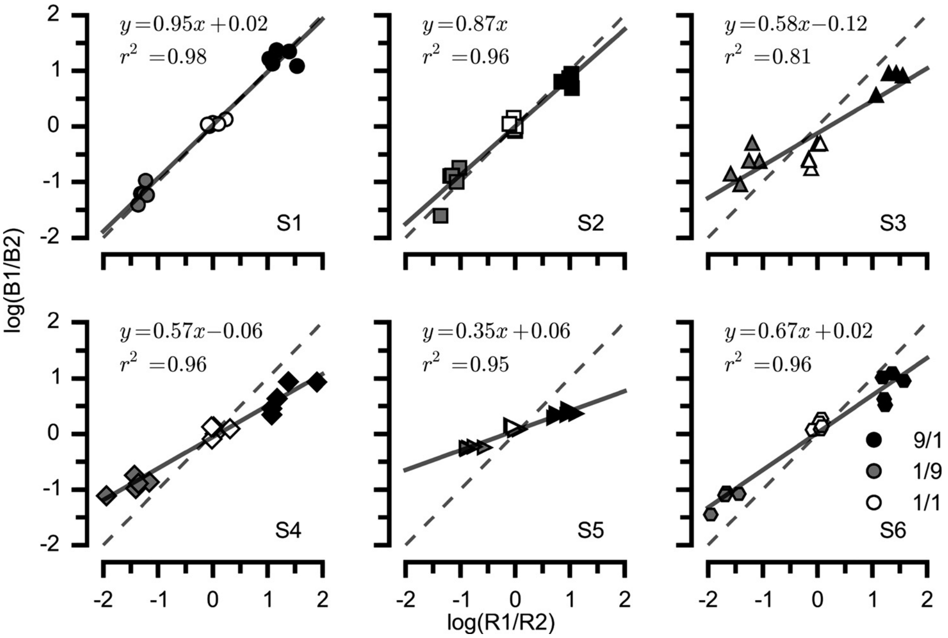

Fig. 4. Relative rates of short and long latencies (log unit) as a function of relative rates of obtained reinforcers for short and long latencies (log unit) for the last 5 sessions of the 3 reinforcement conditions for all participants. *B1* and *B2* represent the total number of trials with short and long latencies, respectively; *R1* and *R2* represent the total number of reinforcers obtained for these 2 classes of latencies, respectively. Each point plots the log relative rate of latencies for 1 session. Symbols are as in Fig. 3. Solid lines were fitted to the data with the method of least squares. The equations of the regression lines are shown at *top* for each participant. The sensitivities (i.e., slopes of the linear fits) ranged from 0.35 to 0.95, and the log bias (i.e., intercept) ranged from −0.12 to 0.06. Dashed lines show the locus of perfect matching ($s = 1.0$).

systematic change in amplitude between short- and long-latency saccades, regardless of saccade directions. However, as shown in Fig. 5, the peak velocities were systematically faster for short-latency saccades (except for *S6*): the mean difference in peak velocities between short and long classes was 20.35°/s (ranging from −5.79°/s to 44.29°/s; all values greater than the null hypothesis 98% CIs). Surprisingly, *participant S6* had slightly faster peak velocities for long latencies for leftward saccades but not for rightward saccades.

DISCUSSION

The present study is the first to specifically manipulate SRTs in a latency-choice paradigm using reinforcement contingencies. SRT distributions considerably changed and choices between short and long latencies strongly matched reinforcement contingencies, revealing a fine control of SRTs. Factors such as visual saliency or reward expectancy (see Sumner 2011) cannot account for the observed changes in SRT distributions. We discuss these findings in the context of previous studies on the temporal organization of behavior below.

SRTs and the Matching Law

Participants matched their SRT relative rates to the obtained reinforcer relative rates across all reinforcement contingencies (with the exception of *S5*, who showed weak matching at best; see Fig. 4). This matching reveals proper choice adjustment: our study falls within a long history of research on matching behavior, i.e., the tendency of participants to match their relative choices to the relative incomes derived from them (reviewed in Davison and McCarthy 2016; McDowell 2013). Matching has been found in a variety of species, behaviors, and reinforcers (see Baum 1979; de Villiers and Herrnstein 1976 for reviews) such that it has been proposed that it constitutes an innate policy (Gallistel 2005). From an evolutionary perspective, matching reveals an effective foraging strategy, as it results in an equilibrium in which returns from two competing behaviors are equalized (Sugrue et al. 2004). The systematic matching relation we observed reveals that our participants adjusted their SRTs according to the reinforcement contingencies in force: they learned to choose when to saccade with short or long latency in our temporal foraging task. To our best knowledge, our study is the first to establish that SRT allocation follows the same rule as that found for other choice situations, revealing a common policy of the saccadic system for spatial (Sugrue et al. 2004) and temporal foraging.

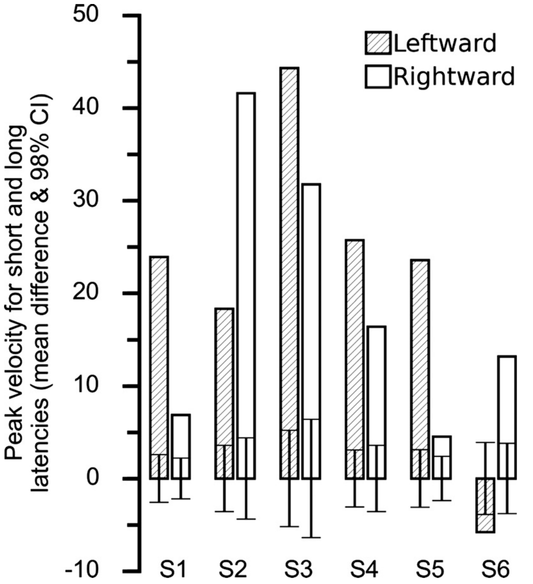

Fig. 5. Differences of the means in saccadic peak velocity between the short and long latency trials across the last 5 sessions of the 3 reinforcement conditions (15 sessions in total) for leftward and rightward saccades for each participant. The corresponding 98% confidence intervals of rejecting the null hypothesis (i.e., the 2 groups of trials have identical probability distribution) are shown. All mean differences are outside the null hypothesis confidence intervals.

SRTs and Reinforcement

That SRT allocations followed the matching law provides further evidence that saccades are instrumental behavior (Madelain et al. 2011), i.e., are controlled by their functional consequences (Skinner 1981), in line with previous studies on saccade latencies (e.g., Ikeda and Hikosaka 2007; Lauwereyns et al. 2002; Madelain et al. 2007; Montagnini and Chelazzi 2005; Nakamura and Hikosaka 2006; Watanabe et al. 2003), saccade peak velocities (Montagnini and Chelazzi 2005; Reppert et al. 2015; Takikawa et al. 2002), saccade amplitudes (e.g., Madelain et al. 2008; Paeye and Madelain 2011, 2014), or target choice (Sugrue et al. 2004). Our results further establish that saccades follow the law of effect, which states that behaviors are shaped by their consequences: SRT allocation changed when functional consequences changed, revealing that the temporal structure of the environmental contingencies affects the temporal structure of behaviors. Future research should further examine these findings with a more ecological paradigm as well as probing its maintenance with the removal of reinforcement. Because our study implied a learning procedure requiring a large number of reinforced trials to observe steady-state SRT allocations, we suggest that this sensitivity of the saccadic system to environmental contingencies relies on the accumulation of learning across time. Whether this learning concerns the temporal properties of the reinforcement contingencies, the temporal control of saccades, or both remains an open question at this stage.

It is noteworthy that this fine control of SRTs raises the issue of perceiving our own reaction times. In our paradigm, one can distinguish three candidate sources of uncertainty: the ongoing contingency (Gallistel et al. 2001; Killeen and Smith 1984), the perception of time (Wearden 2016), and the SRT production (Sumner 2011). Indeed, the actual contingency was unsignaled: participants were not explicitly told what the short or long latency classes were, nor did they have feedback on their actual latency. Thus when a participant received a reinforcer he/she could view the behavior and the consequence as either related (contingent) or not (noncontingent). Our SRT distributions changed, implying that the credit assignment problem (Staddon 2001) must have been solved, i.e., causes and effects were connected. Therefore, despite the uncertainties and noise regarding the motor execution, their own latency perception, and the reinforcement contingency, participants were surprisingly good at adjusting their relative rate of latencies to the relative rate of reinforcement.

SRTs Are Not a Function of Reward Expectancy

One might interpret our results regarding SRTs as a function of reward expectancy. Indeed, predicted reward outcome is known to influence saccadic eye movements, with mean SRT being shorter in the rewarded condition (e.g., Dunne et al. 2015; Glaser et al. 2016; Takikawa et al. 2002; Watanabe et al. 2003). Reward expectancy is also thought to have an effect on saccade vigor, i.e., a saccade velocity greater than the expected velocity given its amplitude (e.g., Reppert et al. 2015). Interestingly, Glaser et al. (2016) observed a negative correlation between vigor and latency. Thus it is commonplace to associate reward expectation, short latencies, and saccade vigor. However, in our study we cannot conclude that the observed changes in latencies are simply a function of reward expectancy. First, reward expectancy and shorter latencies were disentangled, since in our 1/9 condition reinforcement rate—and probably reward expectancy—was higher for long latencies, driving a higher proportion of long latencies. Second, higher peak velocities were observed for short latencies regardless of reinforcement expectancy (see Fig. 5). Therefore, even if we have found, as others, greater vigor associated with shorter latencies, both vigor and shorter latencies are independent of reward expectancy.

Effects of a Dynamic Environment on Saccades

To our knowledge, the article by Hoppe and Rothkopf (2016) reports the only previous study on the learning of temporal eye movement strategies in a dynamic task. They demonstrated that humans could efficiently learn to adapt the time spent within a particular region to temporal regularities depending on the target spatial location on the screen. Their study differed from our own in that *1*) we measured SRTs instead of the time spent in a specific region, 2) we used a concurrent paradigm with simultaneous alternatives, *3*) we used shorter temporalities (between 80 and 300 ms in our experiments vs. 150–1,500 ms), and *4*) those temporal events were produced by the participant in the absence of external cues other than the target step. While these authors evidenced the effects of event durations on fixation durations, our study demonstrated the influence of the temporal organization of the environmental contingencies on the temporal organization of saccades. Taken together, their results and ours provide strong evidence of the possibility of finely controlling the temporal allocation of gaze by reinforcement contingencies. This is strongly compatible with observations from other preparations such as manual interception tasks, which demonstrated that prior information regarding the temporal features of the task (de la Malla et al. 2012) affects the choice of the interceptive point favoring either spatial or temporal precision (de la Malla and López-Moliner 2015), or reward harvesting in complex search tasks in which humans successfully seek multiple targets under time pressure (Navalpakkam et al. 2010). Altogether, these results provide strong evidence that humans may learn both the spatial and temporal regularities of the environmental contingencies to regulate their motor responses.

Costs and Benefits of Saccades

Most models of decision making, such as the LATER model (Carpenter and Williams 1995; Tatler et al. 2017) or diffusion models (Ratcliff and Rouder 1998), are based on the hypothesis of some noisy accumulation of information to the decision criterion. A critical feature of these models is that a saccade is triggered as soon as enough information regarding the target location has been accumulated: SRT reveals the time needed to reach a decision. Departing from the idea of a long-lasting information accumulation, it has been proposed that saccadic decisions might be driven by the sensory information present within the first 100 ms (Ludwig 2009; Ludwig et al. 2005). Why then should SRTs be longer than 100 ms? This procrastination might be a way of prioritizing actions (Harwood et al. 2008) resulting from a trade-off between the benefit of a saccade, which enhances the visual information from a newly fixated location, and associated costs such as the temporary impairment of vision and commitment cost due to the saccade refractory period (Saslow 1967). Such a conception has also been proposed to account for switching between smooth pursuit and catch-up saccades when tracking an object (de Brouwer et al. 2002; Orban de Xivry and Lefèvre 2007) or changes in SRTs due to explicit instructions (encouraging either urgency or accuracy; Reddi and Carpenter 2000). In our case, saccades were explicitly reinforced and did not give access to visual information, but these consequences might be regarded as functionally identical (Davison and Baum 2006; Shahan and Cunningham 2015). Indeed, it is established that SRTs may be experimentally affected by benefit manipulation (e.g., De Vries et al. 2016; Harwood et al. 2008; Madelain et al. 2005, 2007; Montagnini and Chelazzi 2005). Our present results further confirm that observers may learn to alter their information-foraging behavior depending on the actual benefit of making a saccade.

Conclusions

About three times per second a decision is made regarding when and where to move the eyes, providing many opportunities to learn about environment regularities and the benefits of adjusting the information-foraging policy to these properties. In line with pioneer works based on the notions *1*) that acquisition of information is reinforcing (Wyckoff 1952) and *2*) that saccades are controlled by reinforcement (Schroeder and Holland 1968, 1969), we propose that the ability to adapt the allocation of saccades in response to specific temporal organization of information further reveals the exquisite plasticity of the saccadic system and its sensitivity to the current state of the observer and the environment.

ACKNOWLEDGMENTS

We thank Jeremie Jozefowiez, Mark R. Harwood, and Anna Montagnini for helpful comments on an earlier version of the manuscript.

GRANTS

This research was supported by Agence Nationale de Recherche Grant ANR-13-APPR-008 (L. Madelain) and a scholarship from the French Ministry of Research (C. Vullings).

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the authors.

AUTHOR CONTRIBUTIONS

C.V. and L.M. conceived and designed research; C.V. and L.M. performed experiments; C.V. and L.M. analyzed data; C.V. and L.M. interpreted results of experiments; C.V. and L.M. prepared figures; C.V. and L.M. drafted manuscript; C.V. and L.M. edited and revised manuscript; C.V. and L.M. approved final version of manuscript.

REFERENCES

Baum WM. On two types of deviations from the matching law: bias and undermatching. *J Exp Anal Behav* 22: 231–242, 1974. doi:10.1901/jeab.1974.22-231.

Baum WM. Matching, undermatching, and overmatching in studies of choice. *J Exp Anal Behav* 32: 269–281, 1979. doi:10.1901/jeab.1979.32-269.

Brainard DH. The Psychophysics Toolbox. *Spat Vis* 10: 433–436, 1997. doi:10.1163/156856897X00357.

Carpenter RH, Williams ML. Neural computation of log likelihood in control of saccadic eye movements. *Nature* 377: 59–62, 1995. doi:10.1038/377059a0.

Collewijn H, Erkelens CJ, Steinman RM. Binocular co-ordination of human horizontal saccadic eye movements. *J Physiol* 404: 157–182, 1988. doi:10.1113/jphysiol.1988.sp017284.

Davison M, Baum WM. Do conditional reinforcers count? *J Exp Anal Behav* 86: 269–283, 2006. doi:10.1901/jeab.2006.56-05.

Davison M, McCarthy D. *The Matching Law: A Research Review*. New York: Routledge, 2016.

de Brouwer S, Yuksel D, Blohm G, Missal M, Lefèvre P. What triggers catch-up saccades during visual tracking? *J Neurophysiol* 87: 1646–1650, 2002. doi:10.1152/jn.00432.2001.

de la Malla C, López-Moliner J. Hitting moving targets with a continuously changing temporal window. *Exp Brain Res* 233: 2507–2515, 2015. doi:10.1007/s00221-015-4321-x.

de la Malla C, López-Moliner J, Brenner E. Seeing the last part of a hitting movement is enough to adapt to a temporal delay. *J Vis* 12: 4, 2012. doi:10.1167/12.10.4.

de Villiers PA, Herrnstein RJ. Toward a law of response strength. *Psychol Bull* 83: 1131–1153, 1976. doi:10.1037/0033-2909.83.6.1131.

De Vries JP, Azadi R, Harwood MR. The saccadic size-latency phenomenon explored: proximal target size is a determining factor in the saccade latency. *Vision Res* 129: 87–97, 2016. doi:10.1016/j.visres.2016.09.006.

Dunne S, Ellison A, Smith DT. Rewards modulate saccade latency but not exogenous spatial attention. *Front Psychol* 6: 1080, 2015. doi:10.3389/fpsyg.2015.01080.

Efron B. Bootstrap methods: another look at the jackknife. *Ann Stat* 7: 1–26, 1979. doi:10.1214/aos/1176344552.

Ferster CB, Skinner BF. *Schedules of Reinforcement*. East Norwalk, CT: Appleton-Century-Crofts, 1957. doi:10.1037/10627-000.

Gallistel CR. Deconstructing the law of effect. *Games Econ Behav* 52: 410–423, 2005. doi:10.1016/j.geb.2004.06.012.

Gallistel CR, Mark TA, King AP, Latham PE. The rat approximates an ideal detector of changes in rates of reward: implications for the law of effect. *J Exp Psychol Anim Behav Process* 27: 354–372, 2001. doi:10.1037/0097-7403.27.4.354.

Glaser JI, Wood DK, Lawlor PN, Ramkumar P, Kording KP, Segraves MA. Role of expected reward in frontal eye field during natural scene search. *J Neurophysiol* 116: 645–657, 2016. doi:10.1152/jn.00119.2016.

Gold JI, Shadlen MN. The neural basis of decision making. *Annu Rev Neurosci* 30: 535–574, 2007. doi:10.1146/annurev.neuro.29.051605.113038.

Harwood MR, Madelain L, Krauzlis RJ, Wallman J. The spatial scale of attention strongly modulates saccade latencies. *J Neurophysiol* 99: 1743–1757, 2008. doi:10.1152/jn.00589.2007.

Herrnstein RJ. Relative and absolute strength of response as a function of frequency of reinforcement. *J Exp Anal Behav* 4: 267–272, 1961. doi:10.1901/jeab.1961.4-267.

Hoppe D, Rothkopf CA. Learning rational temporal eye movement strategies. *Proc Natl Acad Sci USA* 113: 8332–8337, 2016. doi:10.1073/pnas.1601305113.

Ikeda T, Hikosaka O. Positive and negative modulation of motor response in primate superior colliculus by reward expectation. *J Neurophysiol* 98: 3163–3170, 2007. doi:10.1152/jn.00975.2007.

Killeen PR, Smith JP. Perception of contingency in conditioning: scalar timing, response bias, and erasure of memory by reinforcement. *J Exp Psychol Anim Behav Process* 10: 333–345, 1984. doi:10.1037/0097-7403.10.3.333.

Krauzlis RJ. Eye movements. In: *Fundamental Neuroscience*, edited by Squire LR, Berg D. Amsterdam: Elsevier, 2008, p. 775–792.

Land MF. Oculomotor behavior in vertebrates and invertebrates. In: *The Oxford Handbook of Eye Movements*, edited by Liversedge SP, Gilchrist ID, Everling S. New York: Oxford Univ. Press, 2011, p. 3–16.

Lauwereyns J, Watanabe K, Coe B, Hikosaka O. A neural correlate of response bias in monkey caudate nucleus. *Nature* 418: 413–417, 2002. doi:10.1038/nature00892.

Ludwig CJ. Temporal integration of sensory evidence for saccade target selection. *Vision Res* 49: 2764–2773, 2009. doi:10.1016/j.visres.2009.08.012.

Ludwig CJ, Gilchrist ID, McSorley E, Baddeley RJ. The temporal impulse response underlying saccadic decisions. *J Neurosci* 25: 9907–9912, 2005. doi:10.1523/JNEUROSCI.2197-05.2005.

Madelain L, Champrenaut L, Chauvin A. Control of sensorimotor variability by consequences. *J Neurophysiol* 98: 2255–2265, 2007. doi:10.1152/jn.01286.2006.

Madelain L, Krauzlis RJ, Wallman J. Spatial deployment of attention influences both saccadic and pursuit tracking. *Vision Res* 45: 2685–2703, 2005. doi:10.1016/j.visres.2005.05.009.

Madelain L, Paeye C, Darcheville JC. Operant control of human eye movements. *Behav Processes* 87: 142–148, 2011. doi:10.1016/j.beproc.2011.02.009.

Madelain L, Paeye C, Wallman J. Saccadic adaptation: reinforcement can drive motor adaptation. *J Vis* 8: 919, 2008. doi:10.1167/8.6.919.

McDowell JJ. On the theoretical and empirical status of the matching law and matching theory. *Psychol Bull* 139: 1000–1028, 2013. doi:10.1037/a0029924.

Millenson JR. Random interval schedules of reinforcement. *J Exp Anal Behav* 6: 437–443, 1963. doi:10.1901/jeab.1963.6-437.

Montagnini A, Chelazzi L. The urgency to look: prompt saccades to the benefit of perception. *Vision Res* 45: 3391–3401, 2005. doi:10.1016/j.visres.2005.07.013.

Nakamura K, Hikosaka O. Role of dopamine in the primate caudate nucleus in reward modulation of saccades. *J Neurosci* 26: 5360–5369, 2006. doi:10.1523/JNEUROSCI.4853-05.2006.

Navalpakkam V, Koch C, Rangel A, Perona P. Optimal reward harvesting in complex perceptual environments. *Proc Natl Acad Sci USA* 107: 5232–5237, 2010. doi:10.1073/pnas.0911972107.

Orban de Xivry JJ, Lefèvre P. Saccades and pursuit: two outcomes of a single sensorimotor process. *J Physiol* 584: 11–23, 2007. doi:10.1113/jphysiol.2007.139881.

Paeye C, Madelain L. Reinforcing saccadic amplitude variability. *J Exp Anal Behav* 95: 149–162, 2011. doi:10.1901/jeab.2011.95-149.

Paeye C, Madelain L. Reinforcing saccadic amplitude variability in a visual search task. *J Vis* 14: 20, 2014. doi:10.1167/14.13.20.

Pelli DG. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat Vis* 10: 437–442, 1997. doi:10.1163/156856897X00366.

Ratcliff R, Rouder JN. Modeling response times for two-choice decisions. *Psychol Sci* 9: 347–356, 1998. doi:10.1111/1467-9280.00067.

Reddi BA, Carpenter RH. The influence of urgency on decision time. *Nat Neurosci* 3: 827–830, 2000. doi:10.1038/77739.

Reppert TR, Lempert KM, Glimcher PW, Shadmehr R. Modulation of saccade vigor during value-based decision making. *J Neurosci* 35: 15369–15378, 2015. doi:10.1523/JNEUROSCI.2621-15.2015.

Saslow MG. Effects of components of displacement-step stimuli upon latency for saccadic eye movement. *J Opt Soc Am* 57: 1024–1029, 1967. doi:10.1364/JOSA.57.001024.

Schroeder SR, Holland JG. Operant control of eye movements. *J Appl Behav Anal* 1: 161–166, 1968. doi:10.1901/jaba.1968.1-161.

Schroeder SR, Holland JG. Reinforcement of eye movement with concurrent schedules. *J Exp Anal Behav* 12: 897–903, 1969. doi:10.1901/jeab.1969.12-897.

Shahan TA, Cunningham P. Conditioned reinforcement and information theory reconsidered. *J Exp Anal Behav* 103: 405–418, 2015. doi:10.1002/jeab.142.

Skinner BF. Selection by consequences. *Science* 213: 501–504, 1981. doi:10.1126/science.7244649.

Staddon JE. *Adaptive Dynamics: The Theoretical Analysis of Behavior*. Cambridge, MA: MIT Press, 2001.

Sugrue LP, Corrado GS, Newsome WT. Matching behavior and the representation of value in the parietal cortex. *Science* 304: 1782–1787, 2004. doi:10.1126/science.1094765.

Sumner P. Determinants of saccade latency. In: *The Oxford Handbook of Eye Movements*, edited by Liversedge SP, Gilchrist ID, Everling S. New York: Oxford Univ Press, 2011, p. 3–16.

Takikawa Y, Kawagoe R, Itoh H, Nakahara H, Hikosaka O. Modulation of saccadic eye movements by predicted reward outcome. *Exp Brain Res* 142: 284–291, 2002. doi:10.1007/s00221-001-0928-1.

Tatler BW, Brockmole JR, Carpenter RH. LATEST: a model of saccadic decisions in space and time. *Psychol Rev* 124: 267–300, 2017. doi:10.1037/rev0000054.

Thorndike EL. The law of effect. *Am J Psychol* 39: 212–222, 1927. doi:10.2307/1415413.

Vergilino-Perez D, Fayel A, Lemoine C, Senot P, Vergne J, Doré-Mazars K. Are there any left-right asymmetries in saccade parameters? Examination of latency, gain, and peak velocity. *Invest Ophthalmol Vis Sci* 53: 3340–3348, 2012. doi:10.1167/iovs.11-9273.

Watanabe K, Lauwereyns J, Hikosaka O. Neural correlates of rewarded and unrewarded eye movements in the primate caudate nucleus. *J Neurosci* 23: 10052–10057, 2003.

Wearden JH. *The Psychology of Time Perception*. London: Palgrave Macmillan, 2016. doi:10.1057/978-1-137-40883-9.

Wyckoff LB Jr. The role of observing responses in discrimination learning. *Psychol Rev* 59: 431–442, 1952. doi:10.1037/h0053932.