Temporal Decision-making: Common Procedures and Contemporary Approaches

David Freestone and Fuat Balci

1 Introduction

Various procedures have been developed to investigate interval timing and time-based choice behavior. These methods are typically first developed and validated in non-human animals and then adapted for human testing. Several researchers in particular have contributed to the adaptation of prominent procedures to humans and, thus, to our cross-species understanding of interval timing behavior (e.g., Allan & Gibbon 1991; Rakitin et al. 1998; Wearden 2002). Each of these procedures has advantages and disadvantages in terms of their sensitivity to measuring psychological variables. Thus, the choice of the most suitable procedure for a given research question entails a deep understanding of the procedure and its analysis.

In this chapter, we describe three different interval timing procedures: the peak interval task, the switch task, and the differential reinforcement of low rates of responding (DRL) task. Our aim is to provide a wide range of data acquisition and analysis tools to our readers. The peak interval task is a temporal differentiation task, the switch task is a real-time temporal discrimination task, and the DRL task is a response inhibition task. The peak interval procedure is used most widely, whereas the switch and DRL tasks are used relatively less often. However, together, these three tasks add complementary results about human timing ability. Furthermore, performance on the switch and DRL tasks (and, maybe, the peak interval task) can be evaluated with respect to objective optimal functions to investigate the temporal risk assessment (for review see Balci et al. 2011).

We will describe each procedure separately. We will first provide background information regarding the use of each procedure, followed by the description of the analysis. Programming code written in Python to collect and analyze the data accompany this chapter. These applications include line by line descriptions and tutorials for readers who are interested in following each step of the experimental procedure through data processing. The Python code can be found

at our github page for this chapter (https://github.com/freestone-lab/timing_tasks.git) and at the book's GitHub repository.

2 Peak Interval Procedure

2.1 *Background*

Experimental psychologists have been using fixed interval (FI, operant conditioning) and fixed time (FT, classical conditioning) schedules for decades. The response patterns across species in both tasks is similar, and have led to general and important results regarding animal timing ability (Ferster & Skinner 1957; Pavlov, 1927; Schneider, 1969). Freestone, MacInnis, and Church (2013) have argued that the timing mechanisms exert their effect on both operant and classically conditioned responses in the same way. Pavlov showed that the delay to the onset of conditioned response is correlated with the interval between conditioned stimulus and unconditioned stimulus. This observation was termed as the “inhibition of delay” by Pavlov (Pavlov 1927; Drew et al. 2005). In the original version of the FI schedule of reinforcement, the participant is reinforced for the first response after the fixed interval. This cycle continues over the course of a session without any interruption by inter-trial intervals (ITIs). At steady state, the rate of responding abruptly increases a little more than half way through the trial (e.g., Schneider, 1969). When averaged over trials, the response rate gradually increases throughout the interval. The discrete version of the task, where trials are signaled by a conditioned stimulus (and separated by an ITI), shows the same pattern of results.

The peak interval procedure extends the discrete FI procedure, and is arguably one of the most widely used tasks in interval timing research. This procedure was originally developed by Catania (1970) and Roberts (1981) for animal testing and it has been used to characterize the temporal characteristics of anticipatory responses for reinforcers that are available after a fixed delay after the onset of a conditioned stimulus. Unlike discrete FI procedures, the peak interval procedure also contains test trials in which the conditioned stimulus lasts much longer than the FI and reinforcement is not given. These trials allow us to observe anticipatory responses as a function of trial time based on reinforcement expectancy without any contamination by the reinforcement delivery. The average rate of responding forms a curve with a peak (hence the name). This curve is roughly normally distributed, centered around the typical reinforcement time, which presumably reflects the animal's expectation about when food is typically delivered. The data analysis can be conducted at the

level of average response curves or individual trials. This task has been adapted to human testing in various ways (e.g., Balci et al., 2013; Rakitin et al. 1998). We use the human version of the task below because it is straightforward for readers to follow along using the code we provide with this chapter.

2.2 *Procedural Details*

2.2.1 Fixed Interval Training

In the FI procedure, a timing stimulus is presented and a response option is available (this could be a key on a keyboard or response box or a keyboard). This starts a trial. The first response following the fixed interval results in a reinforcer (e.g., monetary reward), responses prior to this time do not payoff, but are not penalized. The stimulus turns off when the reinforcer is delivered. The FI trials are separated by an intertrial interval (ITI), which should not be predictable. In the limited hold version of the task, the trial ends without reinforcement if too much time has elapsed, that is, the reinforcement is only held for a limited time. This ensures that the experienced delay to the outcome is not much longer than the scheduled delay (e.g., Stoddard et al., 1988).

2.2.2 Peak Interval Testing

Testing is composed of both discrete FI and peak interval (PI) trials. In PI trials, the stimulus is presented much longer than the fixed interval to reinforcement (typically $3 \times \text{FI}$) and responses are not reinforced. Peak trials allow the experimenters to observe the timed behavior not only prior to the typical reinforcement time (as in FI trials) but how expectation declines after the typical reinforcement time has elapsed without reinforcement. The proportion of PI trials to total number of trials often ranges from .33 to .5. Experimenters might choose first give 100% FI trials (training) before adding peak trials, or they may skip training and mix the trial types from the start. The training phase allows participants to learn the reinforcement time on every trial before being tested on peak trials, but this limits the total number of PI trials that can be collected in a single session. There is no restriction on the number of intervals participants may time. For example, two PI procedures can each be associated with a different response, running either concurrently or in separate trials in the same session. The timing stimulus can be a visual cue (e.g., change in the color of a square or the screen) or an auditory signal. And these stimuli can be the same or different for different intervals.

Participants can be either asked to emit multiple responses (Rakitin et al. 1998) or press and hold the key when they think the reinforcement time approaches, then release it when they think the reinforcement was omitted (e.g., Balci et al., 2013). In order to capture the biomechanical cost of responding

in animal studies, an explicit cost can be introduced per response (in the multiple response version) and duration of response (in the press-release version). The tutorial that accompanies this chapter (see Github repository) presents the version of the task that requires multiple responses during the stimulus. Each response is time-stamped and recorded.

2.2.3 Procedure Code

The task was written using PsychoPy (version 1.83.03), and is available on our github page. When the program loads, it reads a file “peak_session_information.csv” that specifies the task parameters like the FI and peak intervals, the reward amount, and the cost per response, along with parameters like the session duration and break duration (it is often useful to give participants a short break to reduce the effects of boredom and fatigue).

In Psychopy, procedures are written as a series of routines. Each routine performs some function, and either waits for a response or some amount of time before moving to the next routine. Typically, there will be a routine for the initial instructions, for closing the experiment, and one or more routines for running the actual trial. PsychoPy provides a graphical user interface for creating these routines. The critical routine for the PI procedure is shown in Figure 6. 1. The top image shows the procedure schematic, and the bottom image shows the primary PsychoPy routine that runs a trial. Stimuli and responses are added as components to each routine. A peak interval trial consists of a keyboard component (for responses), a text component (for the fixation point), and a code component. The code component records the time of each response, and keeps track of the score per response. Not shown in the schematic is the ITI component, but it is important. It not only waits a random time before starting a new trial (randomly drawn between 0.5 and 1.5 s), it also randomizes which trial type to show. Conditional branching in Psychopy is done by creating Routine “Loops” around each trial type that executes either zero or one times per trial. For example, on PI trials, the “Fixed Interval” loop will execute zero times, the “Peak Interval” loop will execute once, the “feedback” loop will execute once, and the “break” loop will execute only if it is time for the participant to take a break (the time between breaks is a variable specified in the “peak_session_information.csv” file).

2.3 Data Processing and Analysis

2.3.1 Average Response Curves

Data gathered from the PI trials are typically expressed in the form of average response curves. In order to build the response curves, the trial duration is divided into fixed-width bins (typically 1 s and the number of responses that fall

FIGURE 6.1 *The peak interval task. Top: A schematic of the task. Bottom: The PsychoPy builder screenshot of a fixed interval trial.*

within each bin is recorded for each trial. The average over trials is computed. The result is a scaled histogram that gives the average number of responses per second, the response rate in each one-second bin. The average response curve is roughly normally distributed with its peak around the typical reinforcement time. The reward expectancy increases smoothly as the participant approaches the reinforcement time, and decreases smoothly after. The more data collected, the smoother the curve (we show data from a single participant

FIGURE 6.2
Example average response curves for Fixed Interval (blue) and Peak Interval (red) trials for a short 30-minute session for a single participant. The target duration was 3s in this example.

in a short session in figure 6.2). Many studies also report a slight right skew in the peak response curves, which may theoretically important (e.g., Balci & Simen 2016; Simen et al. 2011; 2013).

A number of key parameters can be estimated from the average response curve using parametric and nonparametric statistics. The primary parameters are: (i) the location of the peak (peak time), which indicates timing accuracy. The closer the peak time is to the typical food time, the more accurate the participant. And (ii) the width of the response curve, which indicates timing precision. The wider the distribution, the worse a timer a participant is. The first parameter can be estimated by simply locating the trial time of the maximum of the response curve, although this might prove difficult if there is substantial amount of noise in the data. In these cases, the data can be smoothened using methods that would minimize the shift in the actual curve. Alternatively, researchers can take the average of the trial times at which the response curve first exceeds and then first fell below the 90% of the curve's peak. Similarly, the second parameter can be estimated by finding the distance in time between the distributions 25th and 75th percentiles. These constitute nonparametric approaches to the estimation of the timing accuracy and precision.

The parametric approach finds the best fitting distribution function to the response curve. For instance, the best-fit mean μ parameter of the normal distribution function is a measure of the peak time (accuracy), and the best-fit standard deviation σ is a measure of spread (timing precision). The ratio of

the standard deviation to the mean is the coefficient of variation (σ / μ), a critically important measure of precision in the timing literature because it is scale-free, it does not depend on the interval timed. Before using these methods, researchers should study the shape of the peak response curves to choose the best function possible to characterize their data. Researchers should default to using all of the available data unless there is a strict outlier criterion specified in advance. A reasonable metric for the fitted distributions is how well the model captures the data (e.g., a measure like omega squared). In animal data a second gradual increase is sometimes observed toward the end of a PI trial. This increase is thought to reflect anticipatory responding to the next trial. Increasing the stimulus interval on peak trials and making the ITI less predictable can minimize this. When it does occur, fitting a second, increasing function to the unimodal distribution function (for example, a cumulative distribution function) might be necessary. Finally, the amplitude of the peak response curve (peak rate) is another meaningful parameter; it correlates with the subjective value of the anticipated reward in animal studies (Roberts, 1981).

2.3.2 Single Trial Responding

Although the average response curve suggests that participants smoothly ramp up their response rate, the responses in individual trials suggest that participants start responding abruptly at some variable time in each trial. Participants seldom respond early in the trial, then abruptly start rapidly responding. When food is not delivered, participants abruptly lower their response rate again (called the “break-run-break” pattern). This pattern is shown in Figure 6.3. In animal studies, where subjects emit many responses in a trial, experimenters need to apply change detection algorithms to estimate the trial time in which there is an abrupt increase (start time) and an abrupt decrease in response rates (stop time; see Church, Meck, & Gibbon 1994). In human studies, the first response time can often be treated as the start time and the last response time can be treated as the stop time. In the press-release version of this task, the press time is the start time, and the release is the stop time. The time between the start and stop called the spread, is a trial-by-trial index of timing precision and their midpoint can be treated as the trial-based timing accuracy. The distribution of both the start and stop times can give additional measures of timing accuracy and precision (Freestone, MacInnis, & Church 2013). The data analysis tutorial that accompanies this chapter implements these analyses.

There are an enormous number of change detection algorithms that can be used to estimate the start and stop times. Many of these are computationally

FIGURE 6.3 Example single trial raster plots for Fixed Interval (left) and Peak Interval (right) trials for a short 30-minute session for a single participant.

intensive and were not designed to detect changes in hundreds or thousands of trials quickly or efficiently. A few change detection algorithms are uniquely designed to find the start (and sometimes stop) times in FI schedules. The Church, Meck, and Gibbon (1994) method finds the set of transition times that maximize the total distance between each segment's response rate and the overall rate. In other words, it finds the response that maximally segments the data into low and high response rates. The usual implementation of this method is to run through each combination of possible start and stop times on every trial and compute the distance metric just described. The computational cost of this approach is high and depends heavily on the number of trials and number of responses per trial. Gallistel, Fairhurst, and Balsam's (2004) method finds the point that maximizes the distance between the cumulative response data and the null hypothesis line connecting the first point to the last. This method is computationally fast for detecting the start times, but when detecting multiple transitions, the usual implementation is to iterate over the responses until a significant change-point is found, then segment the data and start over. This reduces the speed of the algorithm, and requires a researcher-defined significance value for deciding when a change has been found. For finding start times, these two methods are conceptually different but mathematically identical. They both find the location, where the cumulative residuals against the mean peaks. Framed this way leads to an algorithm that is computationally very fast; the start times for thousands of trials can be estimated in less than a second. A more general method is to use the cumulative residuals against a regression line, called the OLS-CUSUM algorithm (Ploberger & Kramer 1992). These three methods differ for the PI task when both start and stop times need to be estimated. More formal methods for jointly estimating

the start and stop times exist, but they are computationally intensive, even for a relatively small number of trials.

Because the Church, Meck, and Gibbon (1994) implementation is still the most heavily used, we include it in the accompanying Python code (Just-In-Time compiled to be faster; “see the Github repository”). We also include an algorithm that combines aspects of the above discussion. We fit a regression line to the cumulative inter-response times (not the cumulative record), and compute the residuals. The minimum residual is the stop time. The start time is the maximum residual in the data record up to the stop time. That is, start times are conditioned on stop times. This method is computationally efficient, and seems to do a robust job on our data. It should be noted that any algorithm has pros, cons, and edge-cases in which the algorithm fails to find reasonable change-points. The experimenter should examine their particular results to assess the degree to which the algorithm provides a good description of the data.

2.3.3 Analysis Code

The analysis code uses the Python programming language with the scientific stack, most notably Pandas. Python is used because it is free, open-source, cross-platform, and includes libraries that allow for quick prototyping of the task (PsychoPy). The Pandas package contains a rich set of functions that operate on the split-apply-combine framework (Wickham 2011). This allows one to write code for a single instance and then apply it to groups of data in a single line of code, both maximizing readability and minimizing errors. The Jupyter Notebook allows for figures, code, and markdown/latex text in the same document, effectively creating a human readable analysis script. All of the necessary software was downloaded using the Python Anaconda distribution from Continuum Analytics (<https://www.continuum.io/>). The analysis uses Python (version 3.5), Jupyter Notebook (v4.2), Pandas (v0.18), Numpy (v1.10.4), and Matplotlib (v1.5.1). The R programming language with the “tidyverse” ecosystem also implements this analysis philosophy, and RStudio includes the ability to integrate figures, code, and text into a notebook.

The code that analyzes the PI task is called “Peak_Analysis.ipynb.” It renders as HTML in the browser, but can be used as an analysis script when downloaded. It shows the stages of analysis from start to finish: from loading in the data, to data wrangling (cleaning) to analysis. The outputs are figures and tables. It shows both the averaged and single trial analysis methods discussed in this chapter.

3 Switch Procedure

3.1 *Background*

The switch procedure, rooted in the free-operant temporal discrimination task Platt and Davis (1983), is closely related to the PI procedure described above and the temporal bisection task (Chapter 5, this volume). The temporal bisection task is arguably one of the most commonly used procedures in the study of temporal judgments of humans and non-human animals (Church & Deluty 1977; Allan & Gibbon 1991; Wearden, 1991). In this temporal bisection task, participants are initially trained to discriminate two reference intervals as short and long (e.g., 200 and 800 ms). Once the participants learn to discriminate accurately (e.g., 85% correct), intermediate durations are presented intermixed with the reference durations. Participants are asked to classify these durations as short or long depending on their subjective similarity to the reference durations. No feedback is provided for the categorization of intermediate intervals to avoid explicitly training them on intermediate intervals.

Experimenters can estimate both timing precision and the point of subjective equality (the interval that participants judge as equally distant from the referents) by fitting a logistic function to the participant's trial-by-trial choices to predict the proportion of long responses as a function of test durations. (see Chapter 5, this volume). Although the responses are emitted after the termination of the timing stimulus in the temporal bisection task, recent work showed that both human and non-human decision processes actually evolve dynamically during the timing stimulus (e.g., Balci & Simen, 2014; Machado & Keen 2003).

The switch task was developed specifically to capture this dynamically evolving belief state (Balci et al. 2008, 2009; Kheifets & Gallistel 2012; Kheifets, Freestone, & Gallistel, 2017). This Switch task is the prospective analogue of the temporal bisection task in which participants behaviorally invest in the "short" and "long" latency options freely over the course of the trial (during the stimulus). In this task, only the referent durations are presented. The participants are reinforced for "catching" the reinforcement at the right location at the right time. For example, a mouse will be reinforced for poking its nose into the left port at or after short trials, and reinforced for poking its nose into the right port at or after long trials. The mice do not know in advance whether it is a short or long trial, they learn to switch from the short to the long port between the intervals in order to catch the reinforcer no matter the trial type. Humans may be asked to hold down the left key and then switch to the right when they

believe the short duration has elapsed without reinforcement. This allows the experimenter to observe, on an individual trial, when the belief state of the participant switches from short to long. In other words, instead of estimating the PSE from binary response data, the criterion is directly measured trial-by-trial via the switch times. Below, we present the procedural implementation and analysis for this task.

3.2 *Procedural Details*

Participants are trained to anticipate the reinforcement at two different locations associated with two different intervals (short and long). These are often two different feeding hoppers located at two different sides of the operant chamber in animals and visual targets presented at two different sides of the computer screen in humans. In a given trial, only one of these options is active. The active option is not signaled to the participants; they can only rely on the elapsed time to guide their responding. After experiencing the task parameters in the first trials/sessions, participants often begin at the short location and switch to the long location once they believe the short duration has elapsed without reinforcement. This switch time is the main unit of analysis. In human experiments, participants indicate their choices by pressing one key to indicate their preference for the short option and a second key to indicate their preference for the long option.

3.3 *Procedure Code*

As before, the task was written using PsychoPy (version 1.83.03), and is available on our Github page. When the program loads, it reads a file “Switch_session_information.csv” that specifies the task parameters like the short and long durations, payoffs, and probabilities, along with parameters like the session duration and break duration.

The critical PsychoPy routine is shown in the bottom panel of Figure 6.2, and is described by the procedure schematic in the top panel. A trial consists of a keyboard component (for responses), two shape components (to give visual feedback about which response is currently being recorded), two text components (one that controls the fixation cue that starts the trial and one that keeps track of the participant’s total score), interstimulus interval (ISI; duration without any stimuli drawn from a uniform interval between 0.5 and 1.5 seconds), and a code component. The code component draws a trial type (short or long), and records whether or not the final response was correct.

3.4 *Data Processing and Analysis*

The trial time at which the participant leaves the short latency option for the long latency option is calculated for each trial. The switch latencies aggregated

FIGURE 6.4 The switch task. Top: A schematic of the task. Bottom: The PsychoPy builder screenshot of the trial.

across multiple long-latency trials are nearly normally distributed. The mean (or median) switch time can be treated as the PSE, whereas the coefficient of variation (or interquartile range) of the switch times can be treated as an index of precision in temporal judgments. Note that only the data from the long-latency trials are used in the analysis since participants often do not switch on short trials (they shouldn't). The primary advantage of the switch task over the temporal bisection task is that the belief state of the participant can be evaluated in real time rather than being evaluated at arbitrarily chosen decision

points (i.e., test durations). Figure 6.5 shows about 40 switch times for a single participant.

Although it applies more to the data gathered from animals, another advantage of the switch task over temporal bisection task is that the data gathered from the switch task allows the experimenter to treat responses that originate from timed vs. non-timing processes separately. It is fairly common to observe an exponential (impulsive) component in the switch latencies in addition to the normally distributed data in animals (Balci et al. 2009; Kheifets & Gallistel, 2012). Fitting mixture distributions (e.g., exponential normal mixture distribution) to switch latencies allows the experimenter to work only with the data that come from trials with temporal control over responding. This is simply not achievable with the temporal bisection task. The tutorial that accompanies this chapter allows the readers to easily conduct these analyses.

3.5 Optimality Analysis

Statistical decision theory gives an optimal reinforcement-maximizing strategy on this task. There are three important factors that determine the expected reinforcement attained in the switch task: (1) the probability of four different outcomes (correctly and incorrectly switching on short and long trials), (2) the payoffs associated with those four different outcomes and (3) the probability that a short or long trial will occur. The expected gain in this task is the dot (scalar) product of these vectors; namely the sum of payoffs associated with four different outcomes weighted by the joint probability of the corresponding outcomes. This expected value is computed at every moment in time. The switch time that maximizes the expected reinforcement depends on the level of timing precision (often measured by the coefficient of variation). The optimal switch time depends on timing precision because timing variability determines the probability of switching between the short and the long intervals. Worse timers should switch earlier. The tutorial that accompanies this

FIGURE 6.5

Example switch times from a participant on the switch procedure. The short duration was 2 s and the long duration was 3 s.

chapter allows the readers to conduct the optimality analysis of the data gathered from the switch task.

3.6 Analysis Code

The analysis code is contained in “Switch_Analysis.ipynb” on our Github repository. It walks the user through the stages of analysis, from reading in the data file to assessing optimality in the participants. A detailed mathematical treatment of the optimal solution accompanies its implementation. Briefly, it is possible to construct performance curves that specify the expected reinforcement for any given response mean switch time. Three example curves are given in Figure 6.6. Each curve shows a different level of timing precision. A coefficient of variation of 0.20 (blue curve), for example, means that most of the switches on long trials will be within about 20% of the mean switch time. The optimal time to switch from the short to the long location is when this curve peaks. Notice that as timing precision grows (from the blue to the orange to the red curve), the peak of the curve shifts earlier. The black line shows the optimal performance curve –the curve that specifies what the optimal switch time should be for any given level of timing precision.

From here, experimenters can assess the degree to which their participants are optimal: find the optimal switch time for each participant (because their timing precision varies), and then compare it to their actual mean switch time. Figure 6.7 shows two such comparisons. First, how closely the participants match the optimal switch time as the ratio of the actual to the optimal switch time. (left), and second, how much reinforcement they earn compared to what they would earn if they were optimal, again as a ratio (Balcı, Freestone, & Gallistel 2009 see also Freestone et al. 2015).

FIGURE 6.6 Example performance curves for the switch task. The black optimal performance curve shows that the optimal switch time decreases as the coefficient of variation increases.

FIGURE 6.7 Example ways of assessing optimality in participants. Notice the y-limits start at .90, humans are often very close to optimal on this task.

4 The Differential Reinforcement of Low Rates Procedure

4.1 Background

The DRL procedure has been widely used in psychopharmacological animal studies because it is sensitive to anti-depressant pharmacological agents (e.g., Paterson et al. 2011). This task requires participants to wait some amount of time before a response will be reinforced, early responses reset the clock. That is, the DRL task requires participants to withhold responding for at least a fixed interval. Responses after this fixed interval are reinforced, responses earlier than this fixed interval are not reinforced. Both types of responses reset clock. For instance, in a DRL-20s schedule, the minimum wait time since the previous response is set to 20 s. Typically, the average inter-response time in this task is longer than the DRL schedule, and roughly positioned to maximize reinforcement rate (Balci et al. 2011; Çavdaroglu et al. 2014; Freestone, Balci, Simen, & Church 2015; Wearden 1990).

4.2 Procedural Details

The implementation of this task in humans is fairly straightforward. Participants are presented with a stimulus in the middle of the computer screen. They can be instructed to wait for a minimum interval since their previous response before responding again. The minimum interval can be presented at the beginning of the test session for brief training. In other versions of the task, instructions are not provided. If the participant does not wait long enough before

responding, the stimulus can turn red briefly, whereas if they have waited for the minimum interval, the stimulus can turn green briefly, both to indicate the outcome of their response.

4.3 *Procedure Code*

When the experiment starts, it reads the file “DRL_session_information.csv,” which contains the task parameters like the DRL interval, the magnitude of the reward, the cost of a response, as well as information about the experiment, like the session and break durations. The top panel of Figure 6.8 depicts a schematic of the task, the bottom panel shows a screenshot of the PsychoPy builder for the experiment, focused on the “DRL_task” routine. It contains fixation and feedback text components, along with a keyboard component for recording responses. A code component keeps track of each inter-response time and controls the reinforcement.

4.4 *Data Processing and Analysis*

The primary unit of analysis in the DRL task is the inter-response time (IRT). Similar to the Peak procedure, the distributional of IRTs provide information accuracy and precision of timing. In animal data, the IRTs typically originate from two different generative processes leading to a mixture distribution (e.g., exponential and an inverse Gaussian). In these cases, the exponential portion of the IRTs is assumed to originate from those responses for which there was no temporal control over behavior, and the inverse Gaussian portion of the IRTs is assumed to originate from trials in which the responses were under temporal control (the inverse Gaussian distribution). The central tendency relative to the DRL schedule gives timing accuracy, and the spread gives precision. The proportion of inverse Gaussian IRTs can also be used as an index of the degree of temporal control. These parameters are sensitive to both motivational and pharmacological manipulations (Paterson et al., 2011; Doughty & Richards 2002). In the human data, the proportion of exponentially distributed IRTs is virtually zero after training, showing that humans have stronger temporal control over their waiting behavior. Consequently, often fitting a single distribution to human IRTs is sufficient to estimate the parameters of performance. Readers can use the tutorial that accompanies this chapter to fit the DRL data.

4.5 *Optimality Analysis*

The mean IRTs are typically longer than the DRL schedule. This is an adaptive strategy to maximize reinforcement rate given timing imprecision. Like the switch task, it is possible to mathematically describe optimal IRTs depending

FIGURE 6.8 *The DRL Task. Top: A schematic of the DRL task. Bottom: A screenshot of the PsychoPy builder used to run the task.*

on the level of the timing precision (e.g., Çavdaroglu et al. 2014; Freestone, Balci, Simen, & Church, 2015). To a first approximation, the worse the timer, the later they should aim.

The reinforcement rate in this task is the the probability of reinforcement divided by the time to reinforcement. That is, it's the probability that an inter-response time is later than the DRL schedule divided by the average inter-response time. $(p(r; \mu, \gamma\mu) / \mu)$. As the participants aim to wait longer, the numerator of this ratio – the probability of a reinforcement – increases. The denominator – the time cost – also increases. The longer they wait, the more likely they are to be reinforced, but the longer they have to wait for reinforcement. The reinforcement rate maximizing (i.e., optimal) inter-response time finds the balance between these two time varying quantities. The tutorial accompanying this chapter allows the reader to conduct the optimality analysis of data gathered from the DRL task.

FIGURE 6.9 Example inter-response times for a single human participant on the DRL task. The left panel shows every inter-response time in a scatter plot, and the right panel shows the histogram.

4.6 Analysis Code

The analysis code on our Github page is called “DRL_Analysis.ipynb.” It walks the user through the analysis up to assessing optimality in the participants, and generating figures. A detailed mathematical treatment of the optimal solution accompanies its implementation.

5 Conclusions

In this chapter, we have introduced three different timing procedures that are equally applicable to humans and animals. The PI procedure is a widely used timing task that has provided valuable information regarding the nature of the generative psychological (e.g., Church, Meck & Gibbon, 1994; Gibbon & Church 1990), and neurobiological processes (e.g., Meck 2006) that underlie interval timing. For instance, the “break-run-break” pattern of responding in individual trials (as well as other derived quantities such as the middle time and spread) provide insights regarding the possible sources of noise in the timing behavior. The switch procedure allows the experimenters to track the evolving temporal belief state of the individuals in each trial. This procedure is particularly well-suited to study how quantities like probabilities and payoffs change timing behavior. When it is coupled with the participant’s level of timing precision, one can apply statistical decision theory to obtain the optimal response against which the actual data can be compared. Finally, the DRL task is ideal for the study of timed inhibitory control. Performance on this task can also be compared against the optimal response times, and reinforcement-rate maximizing wait time is determined purely by timing precision.

References

- Allan, L.G., & J. Gibbon (1991). Human bisection at the geometric mean. *Learning and Motivation*, 22, 39–58.
- Balci, F., & Simen P. (2014). Decision processes in temporal discrimination. *Acta psychologica*, 149, 157–168.
- Balci, F., & P. Simen (2016). A decision model of timing. *Current Opinion in Behavioral Sciences*, 8, 94–101.
- Balci, F., E.B. Papachristos, C.R. Gallistel, D. Brunner, J. Gibson, & G.P. Shumyatsky (2008). Interval-timing in the genetically modified mouse: A simple paradigm. *Genes, Brain, and Behavior*, 7(3), 373–384.
- Balci, F., D. Freestone, & C.R. Gallistel (2009). Risk assessment in man and mouse. *Proceedings of the National Academy of Sciences*, 106(7), 2459–2463.
- Balci, F., D. Freestone, P. Simen, L. deSouza, J.D. Cohen, & P. Holmes (2011). Optimal temporal risk assessment. *Frontiers in Integrative Neuroscience*, 5, 1–15.
- Balci, F., & Simen, P. (2014). Decision processes in temporal discrimination. *Acta psychologica*, 149, 157–168.
- Balci, F., M. Wiener, B. Cavdaroglu, & B.H. Coslett (2013). Epistasis effects of dopamine genes on interval timing and reward magnitude in humans. *Neuropsychologia*, 51(2), 293–308.
- Catania, A.C. (1970). Reinforcement schedules and psychophysical judgments: A study of some temporal properties of behavior. In Schoenfeld, W.N. (ed.), *The theory of reinforcement schedules*. New York: Appleton-Century-Crofts.
- Çavdaroğlu, B., M. Zeki, & F. Balci (2014). Time-based reward maximization. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 369(1637), 20120461.
- Church, R.M., & M.Z. Deluty (1977). Bisection of temporal intervals. *Journal of Experimental Psychology: Animal Behavior Processes*, 3(3), 216.
- Church, R.M., W.H. Meck, & J. Gibbon (1994). Application of scalar timing theory to individual trials. *Journal of Experimental Psychology: Animal Behavior Processes*, 20, 135–155.
- Doughty, A.H., & J.B. Richards (2002). Effects of reinforcer magnitude on responding under differential-reinforcement-of-low-rate schedules of rats and pigeons. *J Exp Anal Behav.* 78,(1) 17–30.
- Drew, M.R., B. Zupan, A. Cooke, P.A. Couvillon, & P.D. Balsam (2005). Temporal control of conditioned responding in goldfish. *Journal of Experimental Psychology: Animal Behavior Processes*, 31, 31–39.
- Ferster, C.B., & B.F. Skinner (1957). *Schedules of reinforcement*. New York: Appleton-Century-Crofts.

- Freestone, D.M., F. Balci, P. Simen, & R.M. Church (2015). Optimal response rates in humans and rats. *Journal of Experimental Psychology: Animal Learning and Cognition*, 41(1), 39.
- Freestone, D.M., M.L. MacInnis, & R.M. Church (2013). Response rates are governed more by time cues than contingency. *Timing & Time Perception*, 1(1), 3–20.
- Gallistel, C.R., S. Fairhurst, & P. Balsam (2004). The learning curve: implications of a quantitative analysis. *Proceedings of the national academy of Sciences of the United States of America*, 101(36), 13124–13131.
- Gibbon, J., & R.M. Church (1990). Representation of time. *Cognition*, 37, 23–54.
- Kheifets, A., & C.R. Gallistel (2012). Mice take calculated risks. *Proceedings of the National Academy of Sciences*, 109(22), 8776–8779.
- Kheifets, A., D. Freestone, & C.R. Gallistel (2017). Theoretical implications of quantitative properties of interval timing and probability estimation in mouse and rat. *Journal of the Experimental Analysis of Behavior*, 108(1), 39–72.
- Machado, A., & R. Keen (2003). Temporal discrimination in a long operant chamber. *Behavioural Processes*, 62, 157–182.
- Meck, W.H. (2006). Neuroanatomical localization of an internal clock: A functional link between mesolimbic, nigrostriatal, and mesocortical dopaminergic systems. *Brain research*, 109(1), 93–107.
- Paterson, N.E., F. Balci, U. Campbell, B. Olivier, & T. Hanania (2011). The triple reuptake inhibitor DOV216,303 exhibits limited antidepressant-like properties in the differential reinforcement of low-rate 72-sec responding assay, likely due to dopamine reuptake inhibition. *Journal of Psychopharmacology*, 25(10), 1357–1364.
- Pavlov, I. P. (1927). Conditional reflexes: An investigation of the physiological activity of the cerebral cortex (V Anrep, Trans). Martino Fine Books.
- Platt, J.R., & E.R. Davis (1983). Bisection of temporal intervals by pigeons. *Journal of Experimental Psychology: Animal Behavioral Processes*, 9, 160–170.
- Ploberger, W., & W. Krämer (1992). The CUSUM test with OLS residuals. *Econometrica: Journal of the Econometric Society*, 271–285.
- Rakitin, B.C., J. Gibbon, T.B. Penney, C. Malapani, S.C. Hinton, & W.H. Meck (1998). Scalar expectancy theory and peak-interval timing in humans. *Journal of Experimental Psychology: Animal Behavior Processes*, 24, 15–33.
- Roberts, S. (1981). Isolation of an internal clock. *Journal of Experimental Psychology: Animal Behavior Processes*, 7(3), 242.
- Schneider, B. A. (1969). A two-state analysis of fixed interval responding in the pigeon. *Journal of the Experimental Analysis of Behavior*, 12(5), 677–687.
- Simen, P., F. Balci, L. deSouza, P. Holmes, & J.D. Cohen (2011). A model of interval timing by neural integration. *Journal of Neuroscience*, 31(25), 9238–9253.

- Simen, P., F. Rivest, E.A. Ludvig, F. Balci, & P. Killeen (2013). Timescale invariance in the pacemaker-accumulator family of timing models. *Timing & Time Perception*, 30, 159–188.
- Stoddard, L. T., Sidman, M., & Brady, J. V. (1988). Fixed-interval and fixed-ratio reinforcement schedules with human subjects. *The Analysis of verbal behavior*, 6(1), 33–44.
- Wearden, J.H. (1990). Maximizing reinforcement rate on spaced-responding schedules under conditions of temporal uncertainty. *Behavioural processes*, 22(1), 47–59.
- Wearden, J. H. (1991). Human performance on an analogue of an interval bisection task. *The Quarterly Journal of Experimental Psychology*, 43(1), 59–81.
- Wearden, J.H. (2002). Travelling in time: A time-left analogue for humans. *Journal of Experimental Psychology: Animal Behaviour Processes*, 28, 200–208.
- Wickham, H. (2011). The split-apply-combine strategy for data analysis. *Journal of Statistical Software*, 40(1), 1–29.