Simulating bout-and-pause patterns with reinforcement learning

Kota Yamada

Keio University & LeapMind Inc.

Atsunori Kanemura

LeapMind Inc.

Author Note

Correspondence should be addressed to K. Yamada, Department of Psychology,

Keio University, Tokyo 108-8345, Japan (Email: haroldthebarrel.yk@gmail.com).

Abstract

Animal responses often occur according to a specific temporal structure composed of
two states, where a bout (or a burst of engaged responses in a short period) is followed
by a long pause (or disengagement) period until the next bout. Such a bout-and-pause
pattern has three components: the bout length, the within-bout response rate, and the
bout-initiation rate. Previous studies have investigated how these three components are
affected by experimental manipulations. However, it remains unknown what underlying
mechanisms cause bout-and-pause patterns. In this article, we propose a computational
model that is characterized by two mechanisms. The first mechanism is choice—an
agent makes a choice between operant and other behaviors. The second mechanism is
cost—a cost is associated with the changeover of behaviors. The proposed model has
three states instead of two, and we show that the third state is required to reproduce
bout-and-pause patterns through computer simulations. We further analyzed the
proposed model and found that it reproduced the relationship between experimental
manipulations and each of the three components that has been reported by previous
studies. The experimental results suggested that choice and cost are the mechanisms
that cause bout-and-pause patterns. Our results indicate that computational approaches
are useful to investigate and understand underlying mechanisms of animal behavior.

Keywords: bout-and-pause patterns, reinforcement learning, choice

Simulating bout-and-pause patterns with reinforcement learning

Introduction

Animals engage in various activities in their daily lives. For humans, it may be working, studying, practicing sports, or playing video games. For rats, it may be grooming, foraging, or escaping from a predator. Although specific activities are different between different species, common behavioral features are often observed.

Bout-and-pause patterns are one of the behavioral features commonly observed in many species. Activities engaged by an animal do not occur uniformly through time but often have short periods in which a burst of engaged responses is observed. For example, in an operant conditioning experiment, a rat presses a lever repeatedly in a short period and then it stops lever pressing. After a moment, the rat starts lever pressing again. The rat switches between the lever pressing state and the no lever pressing state repeatedly throughout the experiment. Such a temporal structure comprising of short-period response bursts and long pauses is observed in various species and activities; for example, email and letter communication by humans (Barabasi, 2005), foraging by cows (Tolkamp & Kyriazakis, 1999), and walking by Drosophila (Sorribes, Armendariz, Lopez-Pigozzi, Murga, & de Polavieja, 2011).

Shull, Gaynor, and Grimes (2001) showed that bout-and-pause patterns can be described with a broken-stick shape in the log-survivor plot of interresponse times (IRTs), which are characterized by a bi-exponential model. If IRTs follow a single exponential distribution, then the log-survivor plot shows a straight line. If IRTs follow a mixture exponential distribution called a bi-exponential model, the log-survior plot shows a broken-stick shape composed of two straight lines that have different slopes. Killeen, Hall, Reilly, and Kettle (2002) shows that lever pressing by rats is well described with a bi-exponential model, suggesting that this behavior has a bout-and-pause pattern. If IRTs follow a bi-exponential distribution, there are two different types of responses; within-bout responses, which have short IRTs, and between-bout responses, which have long IRTs. Each response type has its own exponential distribution in a bi-exponential model. Killeen et al. (2002) formulated

their bi-exponential model as follows.

$$p(IRT = \tau) = (1 - q)\omega e^{-\omega\tau} + qbe^{-b\tau} \qquad (1)$$

Their model has three free parameters: p, ω, and b, each of which correspond to

different components in the bout-and-pause pattern. First, p denotes the mixture ratio

of the two exponential distributions in the model and it corresponds to the length of a

bout. The bout length is the number of responses contained in one bout. Second, w

denotes the rate parameter for the exponential distribution of within-bout IRTs and it

corresponds to the within-bout response rate. Finally, b denotes the rate parameter for

the exponential distribution of between-bout IRTs and it corresponds to the bout

initiation rate.

The three model parameters, which define the overall response rate, are affected by

motivational and schedule-type variables (Brackney, Cheung, Neisewander, & Sanabria,

2011; Podlesnik, Jimenez-Gomez, Ward, & Shahan, 2006; Shull, 2004; Shull et al., 2001;

Shull, Grimes, & Bennett, 2004; Tanno, 2016). Motivational manipulations include the

reinforcement rate, the response-reinforcement contingency, and the deprivation level,

and they affect the bout initiation rate. An example of schedule-type manipulations is

adding a small variable ratio (VR) schedule in tandem to a VI schedule and they affect

the bout length and the within-bout response rate. The reinforcement rate (Shull et al.,

2001; Shull et al., 2004), the response-reinforcement contingency (Shull et al., 2001),

and the deprivation level (Podlesnik et al., 2006; Shull et al., 2001) affect the bout

initiation rate. As the reinforcement rate and the deprivation level increase, the bout

initiation rate increase. Extinction decreases the bout initiation rate. Brackney et al.

(2011) showed that the overall response rate decreased through time under extinction

and that the decrement is mainly caused by the decrease of the bout initiation rate.

However, some results imply that these variables affect not only the bout initiation rate

but also the bout length (Cheung, Neisewander, & Sanabria, 2012; Shull, 2004; Shull,

Gaynor, & Grimes, 2002; Shull et al., 2004). Schedule-type manipulations mainly

change the bout length (Shull et al., 2001; Shull et al., 2004) and the within-bout

response rate (Matsui, Yamada, Sakagami, & Tanno, 2018; Tanno, 2016). However,

there are opposite results that adding a small VR schedule in tandem to VI schedule has
no effect on the within-bout response rate (Shull et al., 2001; Shull et al., 2004). When
a VI schedule is followed by a small VR (tandem VI VR), an animal stays in a bout
longer and emits more responses in each bout. However, Brackney et al. (2011) showed
that if we add a small VR schedule in tandem to a VI schedule, the bout initiation rate
decreased slightly. These previous findings are summarized in Table 1, where we omit
the within-bout response rate since we will not study this bout component because of a
model complexity issue and insufficient experimental findings on it.

Although previous studies have investigated relationships between some
experimental manipulations and the bout components, underlying mechanisms that
generate bout-and-pause patterns have not been elucidated. Smith, McLean, Shull,
Hughes, and Pitts (2014) showed experimentally that choice and cost play important
roles in organizing responses into bout-and-pause patterns. When pigeons were trained
under a single schedule, the log-survivor plot showed a straight line instead of a
broken-stick shape (Bennett, Hughes, & Pitts, 2007; Bowers, Hill, & Palya, 2008).
Smith et al. (2014) trained pigeons under a concurrent VI VI schedule with and without
a changeover delay (COD). When pigeons were trained under the concurrent VI VI
schedule without a COD, the log-survivor plot still did not show a broken stick,
resulting in the same straight line as under the single schedule. However, under the
concurrent VI VI schedule with COD, the log-survivor plot showed a broken stick,
indicating that bout-and-pause patterns were clearly observed. From these experimental
observations, we extracted the following three facts. 1) When animals engage only in
one response in a given situation, bout-and-pause patterns are not observed. 2) If
animals can choose responses from two alternatives without a COD, bout-and-pause
patterns are still not observed. 3) Considering 1) and 2), we conclude that
bout-and-pause patterns emerge only when animals have two (or more) possible
alternative responses under a given situation and there is a COD between the start of
engagement and a reinforcement. These facts are interesting but they remain to be
inductive, and we still do not have constructive explanation that generate

bout-and-pause patterns. Existing studies on bout-and-pause patterns have investigated
to describe the phenomena rather than to provide causal models. Although rich
knowledge on changes in bout-and-pause patterns are available (Killeen et al., 2002;
Matsui et al., 2018; Tanno, 2016), they did not answer the question of "what
mechanisms make responses occur as bout-and-pause patterns?"

In this article, we propose a computational model based on reinforcement learning
that accounts for the causal generation mechanism of bout-and-pause patterns. We
assume that the bout-and-pause patterns is generated by two mechanisms: a choice
between an operant and other behaviors and a cost that is required to a transition from
engaged behavior to another. We suppose that motivational manipulations affect only
the choice mechanism and schedule-type manipulations affect the cost mechanism. To
incorporate those two mechanisms, we design a three-state Markov transition model,
which has an extra state in addition to the bout and pause states. We perform two
simulation studies to analyze the proposed model. In Simulation 1, we introduce our
model on the basis of the two different mechanisms; choice and cost. We show that the
proposed model can reproduce bout-and-pause patterns by confirming that the
log-survivor plot shows a broken-stick shape. We compare three models: a dual model,
a no cost model, and a no choice model. The dual model is composed of both the choice
and cost mechanisms. The no cost model has only the choice mechanism and the no
choice model has only the cost mechanism. Simulation experiments demonstrate that
the dual model can reproduce bout-and-pause patterns but the other two models failed
to reproduce it. It implies that both of choice and cost are indispensable for animal
responses to be organized into bout-and-pause patterns. In Simulation 2, we analyze
the dual model in depth and report its behavior under various experimental settings to
test if the dual model can reproduce relationships between the experimental
manipulation and the bout components discovered so far. Simulation results suggest
that the dual model can reproduce it not only qualitatively but also quantitatively. In
addition, the dual model. We speculate that real animals might have similar
mechanisms that generate bout-and-pause patterns as the dual model.

Simulation 1

Method: Simulation 1

Model.   We designed a three-state Markov process for modeling bout-and-pause patterns (Figure 1(a)). Two of the three states are "Operant" and "Others," in which the agent engages in an operant or other behaviors, respectively. In the third "Choice" state, the agent makes a decision between the operant and other behaviors. By having the Choice state in our model, we incorporate the fact that animals can choose their behavior from available options (e.g. grooming, exploration, and excretion) when they move freely during an experiment. The second fact that is observed in animal behavior is a cost required to make a transition from one behavior to another. Animals must decide whether to keep doing the same behavior or to make a transition, because fast switching is not optimal if a transition incurs a cost. Figure 1(b) and (c) shows two knockout models, the no choice model and the no cost model, respectively. In each model, one of the two mechanisms from the dual model is removed. In the no choice model, an agent can choose only one behavior in a given situation. In the no cost model, a cost is not required when a transition is made.

Here is how the agent travels in the proposed model. In the Choice state, the agent chooses either the operant or other behaviors. As a result of the choice, it moves from the Choice state to one of the Operant or Others states. It makes the choice based on preferences for each behavior, which is denoted by Q_{pref}. We will explain how to calculate Q_{pref} in the next paragraph. In the Operant state, the agent engages in the operant behavior and decides whether to stay in the Operant state or to move to the Choice state after every response. It decides to stay or switch based on Q_{cost}, which represents a transition cost to the Choice state. The Others state is the same as the Operant state except for that the agent engages in other behaviors.

The preference Q_{pref} is a function that compares the operant and other behaviors when the agent makes a choice between operant and other behaviors. The Q_{pref} function changes over time since it is updated based on the presence (or absence) of a

reinforcer per bout. The following equation describe updating rules for Q_{pref}.

$$Q_{\text{pref}}^{(i)}(t + 1) = Q_{pref}^{(i)}(t) + \alpha_{\text{rft}}(r(i) - Q_{\text{pref}}^{(i)}(t)) \tag{2}$$

$$Q_{\text{pref}}^{(i)}(t + 1) = Q_{\text{pref}}^{(i)}(t) + \alpha_{\text{ext}}(0 - Q_{\text{pref}}^{(i)}(t)) \tag{3}$$

If a reinforcer is presented in a bout, Q_{pref} is updated by Equation 2; otherwise, if a reinforcer is not presented in a bout, Q_{pref} is updated by Equation 3. Parameter α_{rft} denotes a learning rate. Parameter r denotes the value of a reinforcer and we assume $r > 0$ when a reinforcer is present and $r = 0$ when a reinforcer is absent. Parameter t denotes time in session and i denotes each option, that is, i is 1 if the operant behavior is chosen and if i is 2 if other behaviors are chosen.

The cost Q_{cost} is another function that defines a barrier in making a transition from the engaged behavior to the Choice state. We assumed that the cost is independent from the preference and only depends on the number of responses that is emitted to obtain a reinforcer from a bout initiation. When a reinforcer is presented , the cost function Q_{cost} is updated according to:

$$Q_{\text{cost}}^{(i)}(t + 1) = Q_{\text{cost}}^{(i)}(t) + \alpha_{\text{rft}}(\log x(i) - Q_{\text{cost}}^{(i)}(t)) \tag{4}$$

Parameter x denotes the number of responses that is emitted to obtain a reinforcer from a bout initiation. The other parameters are the same as Equations 2 and 3.

In the Choice state, the agent chooses either of the Operant or Others states according to the probability distribution calculated from the preferences for the two behaviors. The transition probability is defined as follows:

$$p_i = \frac{\exp\{\beta Q_{\text{pref}}^{(i)}(t)\}}{\sum_i \exp\{\beta Q_{\text{pref}}^{(i)}(t)\}}, \tag{5}$$

where the softmax inverse temperature parameter β represents the degree to which a choice is focused on the highest-value option.

If the agent is in either of the Operant or Others states, it makes a decision whether to continue engaged behavior or to go back to the Choice state. A decision is made according to the probability calculated from the cost and the preference for the

engaged behavior, which is defined as follows:

$$p_{\text{stay}} = \exp\left\{\frac{-1}{w_{\text{pref}}Q_{\text{pref}}(t) + w_{\text{cost}}Q_{\text{cost}}(t)}\right\}, \tag{6}$$

where w_{pref} and w_{cost} are weighting parameters for Q_{pref} and Q_{cost}, respectively. We assumed $w_{\text{cost}} > w_{\text{pref}}$ because schedule-type operations have stronger effects on the bout length than motivational manipulations.

Simulation.   In Simulation 1, we compared the three possible models; the dual model, the no choice model, and the no cost model. The dual model (Figure 1(a)), which includes both the choice and cost mechanisms as we described in the Model section. The second model was the no choice model (Figure 1(b)), which has only the cost mechanism and it can be thought of as a model made by removing the choice mechanism from the dual model. In the model, the agent engages only in the operant behavior. In other words, this model chooses only the operant behavior in the Choice state. The third model was the no cost model, which has only the choice mechanism without the cost mechanism. The no cost model chooses either operant or other behavior independent of the previous behavior; that is, according to this model, the agent does not continue to be in the currently engaged behavior and comes back to the Choice state after each response.

Experimental conditions used in the simulation were as follows. The schedule for the operant behavior was VI 120 s without an intertrial interval, and the schedule for the other behavior was FR 1. The maximum number of reinforcer in the Operant state was 1,000; that is, if the number of reinforcers reached 1,000, the simulation was terminated. The model parameters were α_{rft}, α_{ext}, β, ω_{pref}, and ω_{cost}. We set $\alpha_{\text{rft}} = .05$, $\alpha_{\text{ext}} = .01$, $\beta = 12.5$, $\omega_{\text{pref}} = 1.0$ and $\omega_{\text{cost}} = 3.5$. Response probabilities in the Operant and the Others states were fixed at $1/3$ in each time step. Before the start of the simulation, we initialized the agent and the experimental environment. Initial values of $Q_{\text{pref}}(i)$ and $Q_{\text{cost}}(i)$ were $0, 0$, respectively, and we created a VI table according to Fleshler and Hoffman (1962). We set a time step in the simulation to be 0.1 s.

We show pseudo code of the model and simulation in Algorithm 1.

Algorithm 1 Pseudocode of simulation

Input: MeanInterval

Initialize: $t = 0$, $NumRewards = 0$, $ResponseTimes = \{\}$

    while $NumRewards < 1000$ **do**

        $t = t + 0.1$

        if $i =$ Choice **then**

            Select a state $i \in \{$Operant, Others$\}$ with probability defined by Equation (5)

        end if

        if $i =$ Operant **and** $uniform(0,1) \leq 1/3$ **then**

            Append t to $ResponseTimes$

            if reward is presented **then**

                Update $Q_{\text{pref}}^{\text{Operant}}(t)$ according to Equation (2)

                Update $Q_{\text{cost}}^{\text{Operant}}(t)$ according to Equation (4)

                $NumRewards \mathrel{+}= 1$

            end if

            if reward is absent **then**

                Update $Q_{\text{pref}}^{\text{Operant}}(t)$ according to Equation (3)

            end if

            if $uniform(0,1) \leq$ Equation (6) **then**

                $i = i$

            else

                $i =$ Choice

            end if

         end if

        if $i =$ Others **and** $uniform(0,1) \leq 1/3$ **then**

            Update $Q_{\text{pref}}^{\text{Others}}(t)$ according to Equation (2)

            Update $Q_{\text{cost}}^{\text{Others}}(t)$ according to Equation (4)

            if $uniform(0,1) \leq$ Equation (6) **then**

                $i = i$

            else

                $i =$ Choice

            end if

Results: Simulaton 1

Figure 2(a) shows event records of IRTs generated by each model and Figure 3
shows the model schemes with transition probabilities between each state. The no
choice model generated a dense repetition of only the operant behavior with a high rate
without long pauses (top panel). The probability the agent stayed in the Operant state
was empirically 0.95, which is a good match with the theoretical value of p_{stay} defined
by Equation 6. In the middle panel, the response rate under the no cost model was low
and each response was separated by long pauses. The probability of the agent choose to
transit to the Operant state was empirically 0.06 and the agent returned to the Choice
state immediately after it responded. In the bottom panel, the agent with the dual
model generated a repetitive pattern of responses with a high rate in a short period
followed by a long pause. The agent in the Choice state made a transition to the
Operant state with a .12 probability and it stayed in the Operant state with a .71
probability.

Figure 2(b) show log-survivor plots in to see whether they show a straight line or
a broken stick. We used the IRTs from after the agent was presented 500 reinforcers to
the end of the simulation. The log-survivor plots of the no choice model and the no cost
model were described by one straight line whereas that of the dual model was described
with a broken-stick shape. The no choice model has a steeper slope than the no cost
model and is tangential to the curve of the dual model at the leftmost position. The no
cost model slope was slightly steeper than that of the dual model at the right side.

Discussion: Simulaton 1

Both the event records and log-survivor plots in Figure 2 imply that only the dual
model generates bout-and-pause patterns and the other two models failed to reproduce
bout-and-pause patterns. The event records in Figure 2(a) suggests only the dual model
exhibit bout-and-pause patterns. The log-survivor plot of only the dual model in
Figure 2(b) showed not a straight but a broken-stick shape, which is an evidence that
the underlying IRTs follow a bi-exponential distribution. Thus, only the dual model

reproduced bout-and-pause patterns.

We posit both of the choice and cost mechanisms are necessary to organize responses into bout-and-pause patterns. The no choice model failed because it lacks the choice mechanism. Without the choice mechanism, the agent almost always stayed in the Operant state and responded with a high rate without pauses. The reason behind the failure of the no cost model was the knockout of the cost mechanism. When the cost of switching behaviors is zero, the agent easily return to the Choice state, resulting in operant responses followed by long pauses. The choice and cost mechanisms contribute differently to generate bout-and-pause patterns; the choice mechanism generates pauses and the cost mechanism produces response bursts. Since the dual model has the both mechanisms, it reproduced bout-and-pause patterns.

Simulation 2

Having demonstrated in Simulation 1 that the dual model successfully reproduced bout-and-pause patterns, in Simulation 2 we analyzed this model under various environments. The previous studies (Brackney et al., 2011; Podlesnik et al., 2006; Shull, 2004; Shull et al., 2001, 2002; Shull et al., 2004; Tanno, 2016) have applied various experimental manipulations to animals to understand bout-and-pause patterns. We applied manipulations to the agent in the model by changing the following four environmental settings: the rate of reinforcement, the deprivation level, the schedule type, and the presence of extinction. The rate of reinforcement, the deprivation level, and the presence of extinction have similar effects on bout-and-pause patterns. These manipulations affect the bout initiation rate and the bout length (Shull, 2004; Shull et al., 2001, 2002; Shull et al., 2004). The schedule type affects only the bout length (Shull et al., 2001; Shull et al., 2004). Here we examined whether the dual model reproduces these relationships between the bout components and the experimental manipulations.

Method: Simulation 2

We used only the dual model in Simulation 2. We performed four experiments by manipulating only one of the four variables while keeping the other three variables the same as Simulation 1. The procedure of simulation is also the same as Simulation 1.

The four experimental manipulations are applied independently to each of the four variables: 1) the rate of reinforcement, 2) the deprivation level, 3) the schedule type, and 4) the presence of extinction. 1) We manipulated the rate of reinforcement by varying mean intervals of the VI schedule. Mean intervals used in this simulation were VI 30, 120 and 480 s. 2) We varied reward values obtained in the Operant state to control the deprivation level of the agent. Those values were .5, 1.0, and 1.5 to induce low deprivation, baseline, and high deprivation levels, respectively. The reward value that the agent received by taking other behaviors was the same as Simulation 1 throughout all the simulations. 3) We manipulated the schedule type by adding a small VR schedule in tandem to a VT schedule. The mean interval of the VT schedule was fixed to 120 s and VR values were 0, 4, and 8. 4) To attenuate the engagement to the operant response, we switched the schedule from VI 120 s to extinction after the agent obtained 1,000 reinforcers. The extinction phase finished when 3,600 s (36,000 time steps) elapsed.

When we analyzed the IRTs data from the extinction simulation, we used a dynamic bi-exponential model (Brackney et al., 2011), in which extinction causes the exponential decay of model parameters, p, ω, and b, according to the following equations.

$$1 - q_t = (1 - q_0)e^{-\gamma t} \tag{7}$$

$$\omega_t = \omega_0 e^{-\alpha t} \tag{8}$$

$$b_t = b_0 e^{-\beta t} \tag{9}$$

The parameters γ, δ, and ϵ denote the decay rates of p, ω, and b, respectively. Thus, Equation 1 can be rewritten as follows.

$$p(IRT = \tau) = (1 - q_t)\omega_t e^{-\omega_t \tau} + q_t b_t e^{-b_t \tau} \tag{10}$$

Since the decay of any of the three model parameters p, ω, and b can cause extinction,
we need to find which of these parameters actually decayed during the extinction
simulation. We excluded ω because it was fixed during the simulation. To identify
whether one or both of the p and/or b parameters decayed, we compared three models,
that is, the pb-decay, p-decay, and b-decay models. We calculated WAIC (widely
applicable information criterion) for each model.

Results: Simulation 2

Figure 4 shows the log-survivor plots of IRTs from each of the four simulations.
Figure 4(a) and (b) shows that manipulating the rate of reinforcement or the
deprivation level changed the slope of the right limb. As the rate of reinforcement or
the deprivation level increased, the slope of the right limb became steeper, indicating
that the bout initiation rate became larger. Figure 4(c) shows that adding the tandem
VR schedule to the VT schedule affected only the y-axis intercept of the right limb
without changing its slope. As the required response increased from the baseline to VR
4 or VR 8, the bout length became larger. The broken sticks in Figure 4(d) have
different slopes and y-axis intercepts, suggesting both the bout initiation rate and the
bout length were changed. Table 2 shows estimated parameters of bi-exponential model,
p, w, and b in each manipulation. Parameter p increased as reinforcement rate,
deprivation level, and number of required responses increased. Parameter w is not
changed in all manipulations. Parameter b is increased as rate of reinforcement and
deprivation level increased. However, the right limb was not stable and we performed a
fitting analysis described in the next paragraph.

In Figure 4(d), the total number of IRTs during the extinction phase was
insufficient to reliably estimate the right limb. Then, we analyzed a dynamic
bi-exponential model fitted to the IRTs during extinction. Table 3 shows the WAIC
values for the three models. The smallest WAIC was attained by the pb-decay model.
The model fit analysis suggested that the bout initiation rate and the bout length
decayed during extinction.

Figure 5 shows the boxplots of Q_{pref} and Q_{cost} in the three simulations except for extinction, which are to be used for assessing how the changes in the bout components are mediated. We excluded the extinction simulation because we already know that Q_{pref} causes the change of the bout components since Q_{cost} is fixed during the extinction phase. The top panel show that Q_{pref} and Q_{cost} increased as the rate of reinforcement increased. The middle panel indicates that increasing the deprivation level moved Q_{pref} and Q_{cost} upward. From the low panel, we can see that adding tandem VR schedule increased Q_{cost} without affecting Q_{pref}.

Discussion: Simulation 2

In Simulation 2, we tested whether the dual model has the same characteristics as animals reported by the previous studies. We analyzed the model with four experimental manipulations: the rate of reinforcement, the deprivation level, the schedule type, and the presence of extinction. The rate of reinforcement, the deprivation level and the presence of extinction affected the bout initiation rate and the bout length and adding the tandem VR schedule to the VT schedule affected only the bout length.

Table 4 summarizes the relationship between the experimental manipulations and the bout components observed in the dual model, which suggests that the behaviors of the dual model are consistent not only qualitatively but also quantitatively with the existing knowledge on animal behaviors. Therefore, the dual model satisfies at least the necessary conditions to be a model to be analyzed for the generation mechanism of bout-and-pause patterns.

How Q_{pref} and Q_{cost} depend on experimental manipulations can be understood according to the categorization of motivational and schedule-type manipulations proposed by Shull et al. (2001). The rate of reinforcement, the deprivation level, and extinction, which are categorized as motivational and, in our simulations, affected Q_{pref} and Q_{cost}. In contrast, schedule type only affected Q_{cost}. The point is that our distinction between Q_{pref} and Q_{cost} corresponds with distinction was proposed by Shull et al. (2001).

The dual model is limited to reproduce only some of the previous findings. Here are two examples of limitations. First, our model is not designed for analyzing the addition of a tandem VR schedule to a VT schedule, by which Tanno (2016) and Matsui et al. (2018) found the change of the within-bout response rate, which was fixed in our model. Second, the value and the delay between a response and a reinforcer is fixed in our model. Brackney et al. (2011) and Podlesnik et al. (2006) considered a delayed reinforcement from a bout initiation causes the inverse correlation between the bout initiation rate and the bout length. This result can be reproduce if required responses by tandem VR is very high (more than 32). Third, sometimes the bout length does not decrease during extinction Brackney et al. (2011). Our dual model cannot reproduce this result even if we change model parameters.

General Discussion

In this paper, we have developed a computational model with reinforcement learning. The model is meant to explain how bout-and-pause patterns can be generated and we examined its validity by comparing computer simulations and experimental findings. We hypothesized that a three-state model involving two independent mechanisms, the choice between Operant and Others and the cost in the changeover of behaviors, is necessary to organize responses into bout-and-pause patterns. We demonstrated in Simulation 1 that the dual model can reproduce bout-and-pause patterns under a VI schedule. Simulation 2 inspected the relationships between various experimental manipulations and the bout components in our model. These results support our hypothesis that assumes that an agent transitioning between three states driven by choice and cost organizes its responses into bout-and-pause patterns. This is our answer to "why bout-and-pause patterns emerge?"

Our generative model reproduced the descriptive results reported by (Shull, 2004; Shull et al., 2001, 2002; Shull et al., 2004). Although our dual model does not include Equation 1 explicitly, IRTs generated by the dual model followed the bi-exponential model. This is an example of the first advantage of reinforcement learning discussed in

the next paragraph.

Our model is developed based on reinforcement learning, which has two
advantages. First, reinforcement learning models consist of simple rules. Instead of
modeling complicated behaviors of animals by complicated rules, we implement simple
rules designed by decomposing them into basic rules that are likely to be implemented
in real animals. This is a divide-and-conquer approach to develop a model. Second,
reinforcement learning is a computational approach, where the model is described with
mathematical expressions that are implemented by programming code. We can run
simulations to analyze what happens inside the model. It is more quantitative than
models described by natural languages.

Our model has at least two shortcomings, which are the range of parameters and
the redundancy of the model. The parameters α_{rft}, α_{ext}, β, ω_{pref}, and ω_{cost} in our model
have not been tuned to fit to behavioral data from real animals. The evidence that
supports our parameter selection is that our model quantitatively reproduced
bout-and-pause patterns. Second, although our model has five parameters, real animals
may have fewer parameters. To verify our model for these two points, it would be useful
to compare empirical data from real animals and our computational model would be
useful to further verify our model.

Standing on our model proposed in this paper, we can extend our research to
many directions to explain more various aspects of bout-and-pause patterns. Here we
discuss two of them. First, our results were retrospective to data from previous
behavioral experiments and the proposed model was not tested by its predictionary
ability to unseen data. The second approach is verification based on neuroscientific
knowledge. Even if the model can correctly predict unseen data from behavioral
experiments, it does not guarantee that animals employes the same the model. To
explore real mechanisms that animals implement, it would be effective to compare the
internal variables of the model with neural activities measured from real animals during
behavioral experiments, to perform knock out experiments by inducing lesions at
specific areas of the brain that should be active during the experiments, or to activate

or deactivate specific neurons during the experiment.

Reinforcement learning can be employed to model and explain animal behaviors other than bout-and-pause patterns, since it is a general framework where an agent learns optimal behaviors in a given environment through trial-and-error (Sutton & Barto, 2018). Such a reinforcement learning framework agrees well with the three-term contingency in behavior analysis. There are three essential elements in reinforcement learning; a state, an action, and a reward. The state is what agent observe and is information about an environment. The action is a behavior what the agent takes in a given state. The reward is what the agent obtains as the result of the action. These three elements are similar to a discriminative stimulus, a response, and an outcome. This similarity allows behavior analysts to employ reinforcement learning in their research. For example, Sakai and Fukai (2008) employed actor-critic reinforcement learning to modeling the matching law.

References

Barabasi, A.-L. (2005). The origin of bursts and heavy tails in human dynamics.
    Nature, *435*(7039), 207.

Bennett, J. A., Hughes, C. E., & Pitts, R. C. (2007). Effects of methamphetamine on
    response rate: a microstructural analysis. *Behavioural Processes*, *75*(2), 199–205.

Bowers, M. T., Hill, J., & Palya, W. L. (2008). Interresponse time structures in
    variable-ratio and variable-interval schedules. *Journal of the Experimental
    Analysis of Behavior*, *90*(3), 345–362.

Brackney, R. J., Cheung, T. H., Neisewander, J. L., & Sanabria, F. (2011). The isolation
    of motivational, motoric, and schedule effects on operant performance: a modeling
    approach. *Journal of the Experimental Analysis of Behavior*, *96*(1), 17–38.

Cheung, T. H., Neisewander, J. L., & Sanabria, F. (2012). Extinction under a
    behavioral microscope: isolating the sources of decline in operant response rate.
    Behavioural Processes, *90*(1), 111–123.

Fleshler, M. & Hoffman, H. S. (1962). A progression for generating variable-interval
    schedules 1. *Journal of the Experimental Analysis of Behavior*, *5*(4), 529–530.

Killeen, P. R., Hall, S. S., Reilly, M. P., & Kettle, L. C. (2002). Molecular analyses of
    the principal components of response strength. *Journal of the Experimental
    Analysis of Behavior*, *78*(2), 127–160.

Matsui, H., Yamada, K., Sakagami, T., & Tanno, T. (2018). Modeling bout–pause
    response patterns in variable-ratio and variable-interval schedules using
    hierarchical bayesian methodology. *Behavioural Processes*, *157*, 346–353.

Podlesnik, C. A., Jimenez-Gomez, C., Ward, R. D., & Shahan, T. A. (2006). Resistance
    to change of responding maintained by unsignaled delays to reinforcement: a
    response-bout analysis. *Journal of the Experimental Analysis of Behavior*, *85*(3),
    329–347.

Sakai, Y. & Fukai, T. (2008). The actor-critic learning is behind the matching law:
    Matching versus optimal behaviors. *Neural Computation*, *20*(1), 227–251.

Shull, R. L. (2004). Bouts of responding on variable-interval schedules: effects of
deprivation level. *Journal of the Experimental Analysis of Behavior*, *81*(2),
155–167.

Shull, R. L., Gaynor, S. T., & Grimes, J. A. (2001). Response rate viewed as
engagement bouts: Effects of relative reinforcement and schedule type. *Journal of
the Experimental Analysis of Behavior*, *75*(3), 247–274.

Shull, R. L., Gaynor, S. T., & Grimes, J. A. (2002). Response rate viewed as
engagement bouts: Resistance to extinction. *Journal of the Experimental Analysis
of Behavior*, *77*(3), 211–231.

Shull, R. L., Grimes, J. A., & Bennett, J. A. (2004). Bouts of responding: The relation
between bout rate and the rate of variable-interval reinforcement. *Journal of the
Experimental Analysis of Behavior*, *81*(1), 65–83.

Smith, T. T., McLean, A. P., Shull, R. L., Hughes, C. E., & Pitts, R. C. (2014).
Concurrent performance as bouts of behavior. *Journal of the Experimental
Analysis of Behavior*, *102*(1), 102–125.

Sorribes, A., Armendariz, B. G., Lopez-Pigozzi, D., Murga, C., & de Polavieja, G. G.
(2011). The origin of behavioral bursts in decision-making circuitry. *PLoS
Computational Biology*, *7*(6), e1002075.

Sutton, R. S. & Barto, A. G. (2018). *Reinforcement learning: an introduction*. MIT
press.

Tanno, T. (2016). Response-bout analysis of interresponse times in variable-ratio and
variable-interval schedules. *Behavioural Processes*, *132*, 12–21.

Tolkamp, B. J. & Kyriazakis, I. (1999). To split behaviour into bouts, log-transform the
intervals. *Animal Behaviour*, *57*(4), 807–817.

Table 1

*Previous findings from animal experiments. The three cells marked with "?" do not
have agreement within the previous reports.*

	Reinforcement rate	Deprivation level	Schedule type (tandem VR)	Extinction
Bout initiation rate	↗	↗	↘ or →?	↘
Bout length	↗ or →?	↗ or →?	↗	↘ or →?

Table 2

Estimated parameters of bi-exponential model in each simulation.

Manipulation	Condition	ω	b	q
Rate of reinforcement	2.0 per min	3.08	0.23	0.17
	0.5 per min	3.06	0.09	0.27
	0.125 per min	3.19	0.03	0.31
Deprivation level	High depriva- tion	3.04	0.24	0.17
	Baseline	3.15	0.08	0.26
	Low depriva- tion	3.17	0.03	0.31
Tandem VT 120 VR x	0	3.07	0.09	0.26
	4	3.23	0.08	0.19
	8	3.11	0.08	0.16

Table 3

Parameter selection for the dynamic bi-exponential model with WAIC. The lower WAIC, the better the model.

Model	WAIC
pb-decay	1.936
b-decay	1.940
p-decay	1.980

Table 4

The behavior of the dual model. The cells marked with "" indicates the consistency with the animal findings shown in Table 1. The cells marked with "†" were the cells with "?" in Table 1.*

	Reinforcement rate	Deprivation level	Schedule type (tandem VR)	Extinction
Bout initiation rate	↗*	↗*	→†	↘*
Bout length	↗*	↗*	↗*	↘†

Table 5

The dependency of Q_{pref} and Q_{cost} to experimental manipulations in the dual model.

	Reinforcement rate	Deprivation level	Schedule type (tandem VR)	Extinction
Category	Motivational	Motivational	Schedule-type	Motivational
Q_{pref}	↗	↗	→	↘*
Q_{cost}	↗	↗	↗	→

Figure 1. Model schemes of the dual model, the no choice model and the no cost model.

(a) The model scheme of the dual model. The upper node, the bottom left node, and

the bottom right node correspond to the Choice state, the Operant state, and the

Others state respectively. Each arrow denotes the transition from one state to others.

(b) The model scheme of the no choice model. In this model, the Others state is

omitted. (c) The model scheme of the no cost model. In this model, self-transitions in

the Operant and Others state is omitted.

Figure 2. (a) Response event records in the (top) single, (middle) no cost, and

(bottom) dual models in the 50 s period just after 500 reinforcers were presented (event

records were stable after 500 reinforcer presentation). Each vertical line denotes one

response. (b) Log-survivor plots of the three models drawn by using all the IRTs after

500 reinforcers.

Figure 3. The schemes of the models with transition probabilities between states that

calculated from the simulation data after the agent obtained 500 reinforcer.

Figure 4. The log-survivor plots of (a) the rate of reinforcement, (b) the deprivation
level, (c) the schedule type, and (d) the presence of extinction drawn by all data after
the agent obtained 500 reinforcers.