Pseudo-learning effects in reinforcement learning model-based analysis: A problem of
misspecification of initial preference

Kentaro Katahira^{1,2*}, Yu Bai^{2,3}, Takashi Nakao⁴

Institutional affiliation:

1 Department of Psychology, Graduate School of Informatics, Nagoya University, Nagoya,
Aichi, Japan

2 Department of Psychology, Graduate School of Environment, Nagoya University

3 Faculty of literature and law, Communication University of China

4 Department of Psychology, Graduate School of Education, Hiroshima University

Abstract

In this study, we investigate a methodological problem of reinforcement-learning (RL) model-based analysis of choice behavior. We show that misspecification of the initial preference of subjects can significantly affect the parameter estimates, model selection, and conclusions of an analysis. This problem can be considered to be an extension of the methodological flaw in the free-choice paradigm (FCP), which has been controversial in studies of decision making. To illustrate the problem, we conducted simulations of a hypothetical reward-based choice experiment. The simulation shows that the RL model-based analysis reports an apparent preference change if hypothetical subjects prefer one option from the beginning, even when they do not change their preferences (i.e., via learning). We discuss possible solutions for this problem.

Keywords: reinforcement learning, model-based analysis, statistical artifact, decision-making, preference

Introduction

Reinforcement-learning (RL) model-based trial-by-trial analysis is an important tool for analyzing data from decision-making experiments that involve learning (Corrado and Doya, 2007; Daw, 2011; O’Doherty, Hampton, and Kim, 2007). One purpose of this type of analysis is to estimate latent variables (e.g., action values and reward prediction error) that underlie computational processes. Estimates of these variables are correlated with neural signals, allowing brain regions that represent variables to be identified. Another purpose is to characterize individual subjects by using model parameter estimates. These parameter estimates are related to physiological/personality characteristics or neuronal/mental deficits in individuals (Yechiam, Busemeyer, Stout, and Bechara, 2005; Kunisato et al., 2012; Huys, Pizzagalli, Bogdan, and Dayan, 2013).

Originally, RL model-based analysis was developed for the reward- or punishment-based learning paradigm. However, researchers have begun to apply RL model-based analysis to various other contexts, including the cognitive-control paradigm (Mars, Shea, Kolling, and Rushworth, 2012), the effect of self-choice on subsequent choices in perceptual decision-making (Akaishi, Umeda, Nagase, and Sakai, 2014), and free choice (Cockburn, Collins, and Frank, 2014).

In another context, the validity of the traditional analysis of the free-choice paradigm (FCP) has become controversial (Chen and Risen, 2010; Izuma and Murayama, 2013; Alós-ferrer and Shi, 2015). Researchers in this field have repeatedly demonstrated that choice induces preference change by showing that preference rating increases for chosen items and decreases for items that were not chosen (Brehm, 1956). This tendency has been explained by cognitive dissonance theory, which postulates that choice modulates decision-makers’ preferences to preserve consistency. A change in preference is quantified by positive spreading, that is, rating the difference between two

items is increased after choice, whereas a chosen item is more preferred, and the item that is not chosen is less preferred. However, Chen and Risen (2010) noted that this positive spread is possible due to certain statistical biases, even if a preference remains stable. This flaw is caused by the misspecification of true preference. When the initial preference of two options was deemed to be neutral despite their difference, subjects tended to choose the more preferred option the second time, which caused an apparent “choice-induced preference change”. Izuma and Murayama (2010) explained this problem by conducting computer simulations. Alós-ferrer and Shi (2010) corrected the proof of Chen and Risen (2010) and clarified the limitation of the problem.

The objective of this study is to prove that this methodological flaw can arise in RL model-based analyses in value-based decision-making studies. To achieve this objective, we conducted simulations that considered a hypothetical and reward-based choice experiment. Standard RL models compute the values of possible actions and do not directly represent preference. However, we can consider that the models represent preference for an option if a higher value is assigned to a given option compared to other options. The simulation shows that the RL model-based analysis reports an apparent preference change even when the hypothetical subject does not have any preference change (i.e., via learning) and even if the actual subject prefers one option. We refer to this effect as “the pseudo-learning effect” and show that it affects the model parameter estimates, model selection, and conclusions of an analysis. We also discuss certain possible solutions for this problem.

Methods

The objective of the simulations that were performed in this study is to illustrate the situation

in which the pseudo-learning effect occurs. When the RL model is fitted to data for the case in which the preference is biased toward one option from the beginning of the experiment, learning appears to occur even though the real preference does not change (i.e., the learning does not occur). To clarify the mechanism of this phenomenon, we consider simple simulation settings.

Hypothetical data generation

We consider a simple probabilistic learning task, in which two options exist. One hundred hypothetical subjects perform 100 choice trials. We assumed that all subjects prefer option 1 without loss of generality. We also assume that the subjects select option 1 with a probability of 80% and select option 2 with a probability of 20% independent of their experiences (e.g., reward and choice histories). Thus, the hypothetical subjects do not change preference throughout the experiment (i.e., learning does not occur).

The reward schedule was established as follows. In Case 1, option 1 is associated with a reward with a probability of p_r during the entire experiment, and option 2 is associated with a reward with a probability of $1 - p_r$, i.e., reward probabilities are symmetric between two options. In Case 2, the reward probability p_r is shared between both options. In the simulations, we varied p_r from 0.05 to 1.0.

Models

We primarily consider a typical RL model—the Q-learning model (Watkins and Dayan, 1992)—which is the most commonly employed model for the model-based analysis of choice behavior [Footnote 1]. The model assigns the value $Q_i(t)$ to each action, where i is the index of

the action, and t is the index of the trial. The initial Q values are $Q_1(1) = q_1^0$, and $Q_2(1) = q_2^0$. In a common setting, the initial action values are set to zero (i.e., $q_1^0 = q_2^0 = 0$), and both options have a neutral value prior to the experiment. Based on the outcome of a decision, the action value for action i is updated as

$$Q_i(t + 1) = Q_i(t) + \alpha (R(t) - Q_i(t)),$$

where α is the learning rate that determines how much the model updates the action value depending on the reward prediction error $R(t) - Q_i(t)$. The learning rate is restricted to the range between 0 and 1. In a common RL-model based analysis, the action value of the option that is not chosen is not typically updated; however, an exception and related discussions are noted by Ito and Doya (2009) and Katahira (2015).

Based on the set of action values, the probability of choosing option i during trial t is given by the softmax function:

$$P(a(t) = i) = \frac{\exp(\beta \cdot Q_i(t))}{\sum_{j=1}^K \exp(\beta \cdot Q_j(t))},$$

where β is the inverse temperature parameter that determines the sensitivity of the choice probabilities to differences in the values, and K represents the number of possible actions. In this study, $K = 2$.

The primary model that is considered in this study is the Q-learning model that has initial Q values that are set to zero (i.e., $q_1^0 = 0, q_2^0 = 0$). In this model, the free parameters are the learning rate α and the inverse temperature β . In addition to this model, we consider two null models, which have no learning mechanisms. The first null model, which is the random-choice model, chooses options with equal probabilities for all options (i.e., the model selects both options with a probability of 0.5). This random-choice model does not include a free parameter. We show

that the analysis produces an incorrect conclusion due to the pseudo-learning effect if only this random-choice model is compared to the Q-learning model. The second null model, which is the biased choice model, includes the free parameter p_{bias} . This model chooses option 1 with a probability of p_{bias} and option 2 with a probability of $1 - p_{bias}$. The second null model (i.e., the biased choice model) can include the “true model” for the proposed hypothetical subject with $p_{bias} = 0.8$ (i.e., it can have a constant preference for option 1). This null model was included in the simulation to show that the pseudo-learning effect does not manifest if a researcher included the appropriate null model in the model comparison. The first null model corresponds to a special case of the Q-learning model, where $\beta = 0$, or equivalently, $\alpha = 0$ and $q_1^0 = q_2^0$. The second null model corresponds to a special case in which $\alpha = 0, q_1^0 \neq q_2^0$ and β has been appropriately scaled.

To estimate the model parameters, maximum likelihood estimation (MLE) was separately performed for each subject. MLE searches a single parameter set to maximize the log-likelihood of the model for all trials. This maximization was performed using the LBFGS algorithm that was implemented in Stan (Version 2.8.0, Stan Development Team, 2015), which was accessed from the R programming language (Version 3.2.0, R Core Team, 2015) via RStan (Version 2.8.0, Stan Development Team, 2015).

Model evaluation

To evaluate the models, we computed three criteria including the log Bayes factor, pseudo- r^2 , and the p -value of the likelihood ratio test; these parameters have been employed for the model-based analysis of choice data. First, we compare the primary model with the two null models using the Bayes factor (BF), which is a ratio of the marginal likelihood of two models (Kass and

Raftery, 1995). The log marginal likelihood of model i is approximated using the Bayesian information criterion (BIC), which is given by $BIC_i = -2l + k \log T$,

where l is the log likelihood for the parameter set obtained using MLE, k is the number of free parameters of model i , and T is the number of trials (i.e., $T = 100$). Given the BIC, the Bayes factors of model i and model j can be approximated as

$$BF_{ij} = \frac{p(\text{Data}|\text{Model } i)}{p(\text{Data}|\text{Model } j)} \approx \exp\left(-\frac{BIC_i - BIC_j}{2}\right).$$

When BF_{ij} is greater than 1, model i is favored; when BF_{ij} is less than 1, model j is favored. We applied the natural logarithm of the Bayes factor. A guideline suggests that the log Bayes factor that exceeds three implies “strong” evidence of model i against model j (Kass and Raftery, 1995).

The pseudo- r^2 is a standardized measure of how a model fits a given dataset (Camerer and Ho, 1999; Daw, 2011) and quantifies the relative degree of improvement in predicting choice over pure chance. The pseudo- r^2 is given by $(R - L) / R = 1 - L / R$, where R is assumed to be the log data likelihood of chance (i.e., the random choice model; $100 \times \log(0.5)$ for the proposed case), and L is assumed to be the log likelihood of the fit model. In the simulation settings in this study, the best model is the biased model, whose bias parameter is $p_{bias} = 0.8$. The expected pseudo- r^2 for this “true” model is $1 - (0.8 \cdot \log(0.8) + 0.2 \cdot \log(0.2)) / \log(0.5) = 0.278$.

We also performed a classical hypothesis test (i.e., the likelihood-ratio test). We compared the log likelihood of the null model, which is denoted by l_{null} , with the log likelihood of the alternative model, which is denoted by l_{alt} . The null hypothesis is that the improvement in the likelihood of more complex models relative to simpler models occurred by chance. Under the null hypothesis, the likelihood-ratio statistic $D = -2 (l_{null} - l_{alt})$ is known to obey the chi-square distribution, in which the degree of freedom is the difference in the number of model parameters between the two models.

Results

First, we illustrate how the pseudo-learning effect arises. Figure 1 shows typical choice traces in Case 1. Figure 1A describes the case in which the reward probability associated with option 1 (i.e., p_r) is 0.7, and option 2 is associated with the reward probability of 0.3. Figure 1B describes the case in which $p_r = 0.3$. The prediction of the Q-learning model for the propensity to choose option 1 (i.e., bold lines in the upper panels) approaches 0.8 (i.e., the broken line in the figure), which is the true preference bias of the data generation model. This effect is identified as the pseudo-learning effect because the preference (i.e., the tendency to choose option 1) appears to change in the model behavior, even though the real preference did not change over the entire session.

These results are explained as follows. First, the action values (i.e., Q-values) in the Q-learning model approach the expected reward value from this option after a sufficient number of trials if the learning rate is sufficiently small. Thus, if the Q-value is near the expected reward value, the average of the reward prediction error $R(t) - Q_i(t)$ is near zero, and the average update of the Q-values is near zero. In the case in which the reward probability is greater for the preferred option (i.e., option 1; the expected reward value is 0.7), Q_1 becomes larger than Q_2 (Figure 1A, bottom panel). In this situation, the model can attain the given preference if the inverse temperature β is appropriately tuned. The MLE, which maximizes the log data likelihood of the model, performs this tuning.

Conversely, when the reward probability is smaller for the preferred option compared with the less preferred option (e.g., the condition in Figure 1B), Q_1 becomes smaller than Q_2 after a sufficient number of trials; this situation appears after a greater number of more trials than shown

in Figure 1B. The model cannot easily represent the preference of option 1 in this situation. However, during the transient phase, in which the Q-values attain the expected rewards and the expected reward values are 0.3 for option 1 and 0.7 for option 2, as shown in Figure 1B, Q_1 may be large because the model experiences the outcome from option 1. Thus, Q_1 increases at a faster rate than Q_2 , and the smaller is the learning rate α , the longer is this transient phase. Thus, with a sufficiently small learning rate and a sufficiently large inverse temperature, the model can accurately represent the preference for option 1 during all trials. This situation is possible with a smaller learning rate, which enables the transient phase to cover all trials; the MLE selects this parameter set (Figure 1B).

Effect of reward probabilities in Case 1

Figure 2 shows the results of Case 1 with the reward probability systematically varied. The log BF quantifies how a model is favored over another model. Figure 2A shows the log BF of the Q-learning model over the random choice model (i.e., the null model). When the reward probability is small, the random choice model is favored because it provides a simpler explanation of the data. However, when the reward probability exceeds 0.2, which enables the Q-learning model to represent the preference for option 1 by the mechanisms shown in Figure 1, the log BF exceeds 10 and the posterior probability of the Q-learning model is 20 times larger than the posterior probability of the random choice model. These results may cause a researcher to conclude that the property of the choice behavior can be explained by the Q-learning model better instead of by chance.

The likelihood-ratio test produces a similar conclusion (Figure 2B) when examining the null hypothesis that the improvement of the likelihood by including model parameters (e.g., the

learning rate and the inverse temperature) only occurs by chance. When the reward probability trends away from 0.2, the p-value decreases, and when the reward probability exceeds 0.4, the majority of samples from the simulations become significant ($p < .05$; Figure 2A). Pseudo- r^2 , which is a commonly employed measure for the goodness of fit of RL models, shows similar tendencies (Figure 2D) because pseudo- r^2 quantifies the improvement in the log likelihood compared with the random-choice model. As the reward probability for option 1 increases, the mean pseudo- r^2 approaches the theoretical upper bound of 0.278 (i.e., the broken line; see the Methods section).

Conversely, when the Q-learning model is compared with the biased choice model, in which the bias parameter can represent an inherent preference that is independent of experience, the log BF favored the biased choice model because the Q-learning model no longer yields a better fit to the data, and the biased choice model has fewer parameters (Figure 2C).

Parameter estimates are also dependent on the reward probability. When the reward probability for option 1 is small, the Q-learning model represents the bias using the transient phase of learning. In this case, the learning rate tends to be small (Figure 2E). Conversely, when the reward probability exceeds 0.5, the learning rate tends to increase, which enables the Q-values to approach the expected value, as shown in the example in Figure 1A. Because an excessively large learning rate can produce unstable Q-values, which reduces the likelihood, MLE produces estimates of the learning rate at an appropriate point. The estimated inverse temperature parameter shows the opposite pattern (Figure 2F). When the reward probability for the preferred option is small, the difference between the two action values is small. Because the inverse temperature parameter is multiplied by the Q-values when computing the choice preference in the softmax function, a large inverse temperature can leverage a small difference in Q-value; a larger inverse

temperature can compensate for a small difference in Q-values and accurately represent the intrinsic preference.

Effect of reward probabilities in Case 2

Figure 3 shows the results when both options produce the reward with the same probability (Case 2). Although this situation is less common in reward-learning tasks, we examined this case to clarify the generality of the proposed situation. For this case, the apparent preference due to the pseudo-learning effect does not easily arise because the expected reward is common to both options. However, the pseudo-learning effect can arise in this case because the model can use the transient phase to represent the preference to option 1 by establishing a small learning rate, as shown in the example in Figure 1B. This case is not dependent on the reward probability; for all reward probabilities, the fit of the Q-learning model is better than that of the random-choice model (Figure 3A, B, and D). Based on this explanation, the learning rate is smaller for all reward probabilities (Figure 3E). Conversely, the inverse temperature was large when the reward probability was small, in which the discrepancy between the difference between two action values is small.

Discussion

The RL model-based analysis has been recognized as a useful framework for estimating internal psychological processes using experimental data. However, when the fitted model is misspecified, researchers can attain an incorrect conclusion. This pitfall of model-based analysis has been documented elsewhere (Nassar and Gold, 2013) with other types of model-

misspecification in which assumed models have a constant learning rate while the true learning rate dynamically changes. The results of the present study can be considered to be an example of this problem.

How the pseudo-learning effect can occur in realistic situations

The problem that was investigated in this study can easily occur in real situations. The choice task that is frequently combined with RL model-based analysis employs neutral images that represent the options that subjects should choose. However, such a task is frequently used without examining the neutrality of the images. If the subjects prefer one image to others due to certain personal reasons, their preferences can affect the results of the RL model-based analysis.

The results of the simulation that was performed in this study indicates that the pseudo-learning effect can arise when subjects prefer either option with commonly employed reward probabilities (e.g., 70% for the first option and 30% for the second option). This result indicates that even when the preference of subjects varied and the preferred options are balanced across subjects, the pseudo-learning effect cannot be canceled by averaging across individuals and can appear in the results of a group analysis.

This problem also arises when subjects shape their preference based on the outcome of their first choice, which describes the effect of “first impression” (Shteingart, Neiman, and Loewenstein, 2013). Shteingart et al. (2013) reported that the subject choice behaviors in certain gambling tasks are best described by the Q-learning model, in which the initial action value is replaced with the reward value of the first outcome from this option. Thus, when one option produces a reasonable outcome while another option produces a poor outcome (e.g., an absence

of reward), the effective initial action value of the former may be higher than the later. The problem that was investigated in this study may arise for this situation if the standard Q-learning model is employed.

Relation to a free-choice task

This pseudo-learning problem is closely related to the methodological flow in the FCP, which was described by Chen and Risen (2010), as we discussed in the Introduction of this paper. Studies that apply the FCP have demonstrated that choosing one option increases the preference for the chosen item, whereas subjects tend to lose their preference for the item that is not chosen. However, Chen and Risen (2010) noted that this effect (i.e., choice-induced preference change) can arise even when the subjects' true preferences remain the same under certain reasonable conditions. Thus, the apparent preference change may be a statistical artifact. Thus, the problem is similar to the problem of the misspecification of the initial value of the RL-model based analysis. The preference change in the FCP can be modeled as an RL framework. If the outcome of the chosen item is always assumed to be positive (i.e., the reward probability is 1.0 for the chosen option), the choice data in the FCP can be modeled by an RL model (Akaishi et al., 2013; Nakao et al., 2016). Thus, the problem described in this study can be considered to be an extension of the problem described by Chen and Risen (2010).

How the pseudo-learning effect can be problematic in research

The pseudo-learning effect can be problematic in real situations for at least two reasons. First, certain psychological research questions are translated into the differences among certain model parameters (e.g., how the mood of subjects influence the learning rate; Bakic, Jepma, De Raedt,

and Pourtois, 2014). However, the degree of initial preference can affect the estimation of the learning rate, as shown in the simulation results in this study. Second, other research questions assess whether subjects can learn via reinforcement and if RL models can better explain other models. When learning does not occur but a subject has a high preference for one option and RL models are only compared to unbiased choice models, as with pseudo- r^2 , the data can appear to be explained by the RL theory.

Possible solutions

We discuss several remedies for this problem. The problem arises because only misspecified models are employed. Thus, a principled approach is to include models that capture true intrinsic preferences. These models may include the initial values of each option as free parameters. Alternatively, a constant term may be added to an action value to provide a choice bias. By including a bias component in a null model (i.e., the biased choice model, which has been considered in this study) and comparing it with the RL models, researchers may quantify the effect of true learning processes (e.g., Katahira, Fujimura, Okanoya, and Okada, 2011).

Basic analyses that do not rely on model fits are also important. For example, if a researcher wants to confirm that learning occurred, the primary effect of binned blocks should be statistically significant, which can be confirmed by an analysis of variance (ANOVA). In addition, when the reward probabilities are switched during a session (e.g., Katahira et al., 2011; Katahira et al., 2014), the choices should be significantly biased to the optimal option in each interval with constant reward probabilities.

A fundamental solution attempts to minimize the initial preference that is independent of experience; this method includes making the option (e.g., images) as neutral as possible. For

example, a solution is achieved using a neutral stimulus after measuring preference via a subjective rating. However, researchers should be careful when interpreting this rating because a subjective rating does not always correspond to the effect of an image on choice behavior (Katahira et al., 2011). In addition, this rating can be noisy and cause the methodological flaw of the FCP that was reported by Chen and Risen (2011). Thus, careful interpretation is required.

Conclusion

Computational model-based analysis, which includes RL model-based analysis, is a valuable and powerful tool for analyzing behavioral data and identifying corresponding psychological and biological processes. However, if the model settings fail to capture the underlying characteristics of the behavior, the results may be less meaningful: The parameter estimates and the results of model selection may change due to the statistical artifact. The pseudo-learning effect due to misspecification of the initial preference, which we have discussed in the present study, may be a common scenario in which such problems occur. Researchers should be aware of the limitations of model-based analysis and carefully check the validity of the model assumptions.

Footnotes

1. The RL models considered in the present study are a special case of Q-learning where there is only one state and the delay in reward is not considered. Although they can also be regarded as special cases of SARSA, which is another type of RL, we term the present models as Q-learning according to convention.

References

- Akaishi, R., Umeda, K., Nagase, A., & Sakai, K. (2013). Autonomous Mechanism of Internal Choice Estimate Underlies Decision Inertia. *Neuron*, 81(1), 195-206.
doi:10.1016/j.neuron.2013.10.018
- Alós-ferrer, C., & Shi, F. (2015). Choice-induced preference change and the free-choice paradigm : A clarification. *Judgment and Decision Making*, 10(1), 34–49.
- Bakic, J., Jepma, M., De Raedt, R., & Pourtois, G. (2014). Effects of positive mood on probabilistic learning: Behavioral and electrophysiological correlates. *Biological Psychology*, 103, 223–232. doi:10.1016/j.biopsycho.2014.09.012
- Brehm, J. W. (1956). Postdecision changes in the desirability of alternatives. *The Journal of Abnormal and Social Psychology*, 52(3), 384–389.
- Camerer, C., & Ho, T.-H. (1999). Experienced-weighted attraction learning in normal form games. *Econometrica*, 827–874.
- Chen, M. K., & Risen, J. L. (2010). How choice affects and reflects preferences: revisiting the free-choice paradigm. *Journal of Personality and Social Psychology*, 99(4), 573–594.
doi:10.1037/a0020217
- Cockburn, J., Collins, A. G. E., & Frank, M. J. (2014). A Reinforcement Learning Mechanism Responsible for the Valuation of Free Choice. *Neuron*, 83(3), 551–557.
doi:10.1016/j.neuron.2014.06.035
- Corrado, G., & Doya, K. (2007). Understanding neural coding through the model-based analysis of decision making. *Journal of Neuroscience*, 27(31), 8178–80.
doi:10.1523/JNEUROSCI.1590-07.2007
- Daw, N. D. (2011). Trial-by-trial data analysis using computational models. *Decision Making, Affect, and Learning: Attention and Performance XXIII*, 23, 1–26.
- Huys, Q. J., Pizzagalli, D. a, Bogdan, R., & Dayan, P. (2013). Mapping anhedonia onto reinforcement learning: a behavioural meta-analysis. *Biology of Mood & Anxiety Disorders*, 3(1), 12. doi:10.1186/2045-5380-3-12
- Ito, M., & Doya, K. (2009). Validation of decision-making models and analysis of decision variables in the rat basal ganglia. *Journal of Neuroscience*, 29(31), 9861–74.
doi:10.1523/JNEUROSCI.6157-08.2009

- Izuma, K., & Murayama, K. (2013). Choice-induced preference change in the free-choice paradigm: a critical methodological review. *Frontiers in Psychology*, 4, 41. doi:10.3389/fpsyg.2013.00041
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430), 773–795.
- Katahira, K. (2015). The relation between reinforcement learning parameters and the influence of reinforcement history on choice behavior. *Journal of Mathematical Psychology*, 66, 59–69.
- Katahira, K., Fujimura, T., Matsuda, Y.-T., Okanoya, K., & Okada, M. (2014). Individual differences in heart rate variability are associated with the avoidance of negative emotional events. *Biological Psychology*, 103, 322–331. doi:10.1016/j.biopsycho.2014.10.007
- Katahira, K., Fujimura, T., Okanoya, K., & Okada, M. (2011). Decision-Making Based on Emotional Images. *Frontiers in Psychology*, 2, 311. doi:10.3389/fpsyg.2011.00311
- Kunisato, Y., Okamoto, Y., Ueda, K., Onoda, K., Okada, G., Yoshimura, S., ... Yamawaki, S. (2012). Effects of depression on reward-based decision making and variability of action in probabilistic learning. *Journal of Behavior Therapy and Experimental Psychiatry*, 43(4), 1088–94. doi:10.1016/j.jbtep.2012.05.007
- Nassar, M. R., & Gold, J. I. (2013). A Healthy Fear of the Unknown: Perspectives on the Interpretation of Parameter Fits from Computational Models in Neuroscience. *PLoS Computational Biology*, 9(4), e1003015. doi:10.1371/journal.pcbi.1003015
- Nakao, T., Kanayama, N., Katahira, K., Odani, M., Ito, Y., Hirata, Y., ... Northoff, G. (2016). Post-response $\beta\gamma$ power predicts the degree of choice-based learning in internally guided decision-making. *Scientific Reports*, 6, 32477. doi:10.1038/srep32477
- O'Doherty, J. P., Hampton, A., & Kim, H. (2007). Model-based fMRI and its application to reward learning and decision making. *Annals of the New York Academy of Sciences*, 1104, 35–53. doi:10.1196/annals.1390.022
- Shteingart, H., Neiman, T., & Loewenstein, Y. (2013). The role of first impression in operant learning. *Journal of Experimental Psychology. General*, 142(2), 476–88. doi:10.1037/a0029550
- Watkins, C. J. C. H., & Dayan, P. (1992). Q-learning. *Machine Learning*, 8(3), 279–292.
- Yechiam, E., Busemeyer, J. R., Stout, J. C., & Bechara, A. (2005). Using cognitive models to map relations between neuropsychological disorders and human decision-making deficits.

Psychological Science, 16(12), 973–8. doi:10.1111/j.1467-9280.2005.01646.x

Figures

Figure 1. Examples of model fits to hypothetical biased-choice data (Case 1), in which the reward probability of option 1 is 0.7 and the reward probability of option 2 is 0.3 (A); and the reward probability of option 1 is 0.3 and the reward probability of option 2 is 0.7 (B). For both cases, the upper panels show the simulated sequences of choice and reward and model prediction. The long vertical lines represent the choice (i.e., the upper line indicates choosing option 1 and the lower indicates choosing option 2) and the outcome (i.e., red represents a rewarded trial, and black represents a non-rewarded trial). Rewarded trials are also denoted by small vertical lines. The choice is generated from the biased choice model, in which the probability of choosing option 1 ($= 0.8$) is depicted by a broken horizontal line. The solid black lines indicate the probability of choosing option 1, which was computed from the fitted model. The bottom panels show the estimated time course of the action values Q_1 and Q_2 . The estimated parameters include the learning rate ($\alpha = 0.080$) and the inverse temperature ($\beta = 3.260$) for (A); and $\alpha = 0.017$ and $\beta = 17.34$ for (B).

Figure 2. The effect of the reward probabilities on the model evaluation and parameter estimation when the reward probabilities are symmetric between the two options (Case 1). Each dot represents a sample from a single hypothetical subject. The bold line represents the mean for each reward probability. (A): Log Bayes factor (BF) of the Q-learning model for a random unbiased model. A positive log BF value indicates that the Q-learning model is favored, and a negative value indicates that the random model is favored. (B): P-value of the likelihood ratio test, where the null hypothesis is that the improvement of the Q-learning model compared with the random-choice model only occurs by chance. (C): Log Bayes factor of the Q-learning model of the biased choice model. (D): Pseudo- r^2 . The broken line indicates the theoretical upper bound of the average pseudo- r^2 attained by the true model (i.e., the biased choice model). (E): Estimated values of the learning rate α in the Q-learning model. (F): Estimated values of the inverse temperature β in the Q-learning model.

Figure 3. Effect of reward probabilities on the model evaluation and estimation when the reward probabilities are shared by two options (Case 2). The conventions are the same conventions shown in Figure 2.