Title: Development of a novel computational model for the Balloon Analogue Risk Task: The Exponential-Weight Mean-Variance Model

Authors:

Harhim Park¹, Jaeyeong Yang¹, Jasmin Vassileva^{2, 3}, and Woo-Young Ahn¹

¹Department of Psychology, Seoul National University, Seoul, Korea

²Department of Psychiatry, Virginia Commonwealth University, Virginia, United States of America

³Institute for Drug and Alcohol Studies, Virginia Commonwealth University, Virginia, United States of America

Corresponding author:

Woo-Young Ahn, Ph.D.

Department of Psychology

Seoul National University

Seoul, Korea 08826

Tel: +82-2-880-2538, Fax: +82-2-877-6428. E-mail: wahn55@snu.ac.kr

Keywords: Balloon Analogue Risk Task, risk-taking, Hierarchical Bayesian Analysis, computational modeling, substance use

Number of words in the abstract: 150

Number of words in the main text (including references): 9,203

Number of Figures/Tables in the main text: 7 / 1

Number of Figures/Tables in supplemental online material: 14 / 2

Number of References: 48

Conflict of Interest: The authors declare no competing financial interest.

Abstract

The Balloon Analogue Risk Task (BART) is a popular task used to measure risk-taking behavior. To identify cognitive processes associated with choice behavior on the BART, a few computational models have been proposed. However, the extant models either fail to capture choice patterns on the BART or show poor parameter recovery performance. Here, we propose a novel computational model, the exponential-weight mean-variance (EWMV) model, which addresses the limitations of existing models. By using multiple model comparison methods, including post hoc model fits criterion and parameter recovery, we showed that the EWMV model outperforms the existing models. In addition, we applied the EWMV model to BART data from healthy controls and substance-using populations (patients with past opiate and stimulant dependence). The results suggest that (1) the EWMV model addresses the limitations of existing models and (2) heroin-dependent individuals show reduced risk preference than other groups, which may have significant clinical implications.

1. Introduction

Computational modeling of cognitive tasks has been widely used to address the limitations of behavioral measures, with which it is often hard to identify underlying cognitive processes (Ahn et al., 2014; Daw, Gershman, Seymour, Dayan, & Dolan, 2011; Ratcliff, 1978). For example, research examining a series of models for the Iowa Gambling Task (IGT) has accounted for the various patterns in behavioral data, provided short- and long-term prediction and good parameter recovery (Ahn, Busemeyer, Wagenmakers, & Stout, 2008; Ahn et al., 2014; Busemeyer & Stout, 2002; Haines, Vassileva, & Ahn, 2018; Worthy & Maddox, 2014), and revealed decision-making deficits in several clinical populations that were not detected by traditional performance indices (see Ahn, Dai, Vassileva, Busemeyer, & Stout, 2016 for a review).

Like the IGT, the Balloon Analogue Risk Task (BART) (Lejuez et al., 2002) was originally designed for the clinical purpose of measuring risk-taking tendencies and the identification of individuals who are prone to take risks (Lejuez et al., 2002; Lejuez, Simmons, Aklin, Daughters, & Dvir, 2004), but its scope has expanded to other areas of psychology and cognitive science (van Ravenzwaaij, Dutilh, & Wagenmakers, 2011; Wallsten, Pleskac, & Lejuez, 2005). The BART has been shown to identify high risk-taking individuals (Aklin, Lejuez, Zvolensky, Kahler, & Gwadz, 2005; Lejuez et al., 2002; Lejuez et al., 2004) although several previous findings suggested low correlations between the BART performance and risky behavior (Frey, Pedroni, Mata, Rieskamp, & Hertwig, 2017; Hopko et al., 2006). Specifically, the behavioral performance of the BART is significantly correlated with self-report measures of risk-related constructs such as impulsivity and sensation-seeking (Lejuez et al., 2002), and past real-world risky behaviors such as drug use and risky sex (Lejuez et al., 2004). These results presumably reflect at least two features of the BART which make the task similar to real-world situations: First, each trial of the BART includes sequential risk-taking choices and terminates when the participant does not want to take the risk any more or encounters a certain condition that makes it impossible to proceed. Second, the riskiness of the risk-taking choices within a trial increases each time the participant makes a risky choice. Specifically, as the participant makes a risky choice, the loss amount gradually increases, whereas the reward amount remains constant or even decreases. Many actual risky behaviors involve these features. For example, patients with substance use disorders continue to take the drug even in the face of negative consequences until they are satisfied with it or cannot obtain the drug because of external factors such as lack of availability and insufficient money. Also, patients

1 often start with smaller amounts of the drug but gradually increase their dosage as their tolerance
2 increases. As their tolerance increases, they must (and are likely to) take more risks to get the same
3 amount of reward. The BART has the distinct advantage of effectively illustrating this kind of real-world
4 situation in a laboratory setting.

5 To quantitatively analyze the underlying cognitive processes on the BART, previous studies
6 have proposed two computational models: a four-parameter model (Wallsten et al., 2005) and a two-
7 parameter model (Pleskac, 2008; van Ravenzwaaij et al., 2011). The four-parameter model was
8 proposed as a winning model by comparing several computational models. The parameters of the four-
9 parameter model have been shown to correlate with the frequencies of past real-world risky behaviors
10 such as substance use, unprotected sex, and stealing (Wallsten et al., 2005). However, two parameters
11 related to the learning process of the four-parameter model showed poorer parameter recovery and
12 were systematically overestimated (Heathcote, Brown, & Wagenmakers, 2015; van Ravenzwaaij et al.,
13 2011). Since good parameter recovery performance is crucial for interpreting results based on the model
14 parameters (van Ravenzwaaij & Oberauer, 2009; Wagenmakers, Van Der Maas, & Grasman, 2007),
15 poor parameter recovery performance is a critical limitation of the four-parameter model.

16 The two-parameter model was proposed to overcome this limitation of the four-parameter model
17 (Pleskac, 2008; van Ravenzwaaij et al., 2011). To develop a model that shows good parameter recovery,
18 the authors simplified the original model by removing parameters that do not exhibit good recovery. The
19 two-parameter model is nested within the four-parameter model, and as a result of simplification, it
20 succeeded in recovering accurate parameter values (van Ravenzwaaij et al., 2011). However, the two-
21 parameter model also has a critical limitation; it is based on a strict assumption that participants do not
22 learn during the BART. The assumption is unrealistic unless the researcher tells the participant the
23 actual exploding probability of virtual balloons before starting the experiment. Consistently, Pleskac
24 (2008) showed that the two-parameter model provided a better fit than the four-parameter model when
25 the probability structure was explicitly informed, but a poorer fit than the four-parameter model when
26 the probability structure was uninformed.

27 Furthermore, because this issue is related to the task design, the two-parameter model may not
28 fit the original task design. The original task design has the advantage that it similarly illustrates real-
29 world situations, which means if we modify the task design to apply the two-parameter model, we lose
30 the advantage. Thus, there is a need to build a new model, which shows good parameter recovery and

fits the original task design.

Here, we propose a novel BART model, which shows good parameter recovery, fits the original task design, and provides an intuitive interpretation of the learning process. First, we introduce the existing model (the four-parameter model). We also consider a non-learning version of the four-parameter model to test the assumption that participants do not learn during the BART is unrealistic. Then, we reparameterize the four-parameter model to improve its parameter recovery. By modifying equations from the reparametrized version of the four-parameter model, we develop candidate models. Finally, we select the best model based on the leave-one-out information criterion (LOOIC) and the parameter recovery. To examine the implication of the new model, we compared the parameters with similar psychological constructs of competing models and applied the new model to BART data from patients with past opiate and stimulant dependence.

2. Method

2.1 Participants

The initial sample included 593 individuals who had enrolled for a study of impulsivity in opiate and stimulant users in Sofia, Bulgaria. Then, only those who meet the following criteria were included: age between 18 and 50 years, more than 8 years of formal education, estimated IQ of 80 or above, no history of head injury or loss of consciousness for more than 30 minutes, no history of neurological illness or psychotic disorders, HIV-seronegative status, and not currently on opioid maintenance therapy. All participants had a negative breathalyzer test for alcohol and negative rapid urine toxicology screen for opiates, cannabis, amphetamines, methamphetamines, benzodiazepines, barbiturates, cocaine, MDMA, and methadone. We classified the included participants into three groups: healthy controls, heroin-dependent, and amphetamine-dependent groups. After that, group-specific criteria were applied to make each group include primarily mono-dependent ('pure') users. For the healthy control group, participants with any substance dependence or abuse symptom based on DSM-IV criteria were excluded (except for nicotine, caffeine, and past cannabis dependence). For the heroin and the amphetamine-dependent groups, mono-substance-dependent participants who met DSM-IV lifetime criteria for opiate or stimulant dependence with no dependence on any other substances were included. Finally, a total of 226 subjects (135 healthy controls, 47 heroin-dependent, and 44 amphetamine-

dependent individuals) were included in the analysis. For more details about the recruitment and screening procedures, see Ahn and Vassileva (2016). This study was approved by the Institutional Review Boards of the Virginia Commonwealth University and the Medical University in Sofia. All participants provided informed consent. See supplementary material for demographic and clinical characteristics of the participants (Table S1).

2.2 Task

In the BART, a virtual balloon is presented to the participant on each trial. Participants need to decide whether to pump the balloon to accumulate some predefined amount of reward (i.e., pump), or transfer and receive the reward that has been accumulated so far (i.e., transfer). Each trial ends when the participant chooses to transfer the accumulated reward or the balloon explodes. If the balloon explodes, the participant loses all the accumulated reward on that trial. We randomly determined explosion points for each trial; thus, each participant performed the task with a different set of explosion points. Participants are not informed about the probability of the balloon exploding. Typically, the degree of risk-taking on the BART is measured by the adjusted BART score, which is the average number of pumps for unexploded balloons (Lejuez et al., 2002). The adjusted BART score is preferable because it is not directly affected by the explosion probability. To examine group differences in the adjusted BART score between the three groups, we conducted the Bayesian t-test using the R package *BEST* (Kruschke, 2013; Meredith & Kruschke, 2018).

2.3 Models

2.3.1 The four-parameter model

The four-parameter model (Wallsten et al., 2005) is based on two assumptions. First, the participants update the belief about the probability of the balloon exploding after each trial. Second, the participants decide the optimal number of pumps before each trial.

From a computational modeling perspective, the first assumption means that p_k^{burst} , the participant's perceived probability that pumping the balloon on trial k will make the balloon explode, is constant during the trial k . The participant initially has a prior belief about the probability of the balloon exploding and updates the prior belief based on observation on each trial. The updating process is

described as follows:

$$p_k^{burst} = 1 - \frac{\alpha + \sum_{i=0}^{k-1} n_i^{success}}{\mu + \sum_{i=0}^{k-1} n_i^{pumps}} \text{ with } 0 < \alpha < \mu. \quad (1)$$

In Equation (1), the initial value of p_k^{burst} is $1 - \alpha/\mu$, which reflects the participant's initial belief that pumping will make the balloon explode. The magnitudes of α and μ indicate the degree of learning from observations; high values indicate that the prior belief is strong and the perceived probability is affected to just a small degree by the observed data. $\sum_{i=0}^{k-1} n_i^{success}$ is the sum of the number of successful pumps up to trial $k - 1$, and $\sum_{i=0}^{k-1} n_i^{pumps}$ is the sum of the total number of pumps up to trial $k - 1$.

The second assumption that the participant evaluates the optimal number of pumps before each trial is reflected in the equations for calculating the probability that the participant will pump the balloon. Adopting the prospect theory (Kahneman & Tversky, 2013), the expected utility after l pumps on trial k , U_{kl} , is given by:

$$U_{kl} = (1 - p_k^{burst})^l (lr)^\gamma. \quad (2)$$

In Equation (2), r is the amount of reward per successful pump, and γ is risk-taking propensity. We can calculate the optimal number of pumps by setting the first derivative of Equation (2) for l equals zero. Then, we can easily derive the optimal number of pumps on trial k , v_k , as follows:

$$v_k = \frac{-\gamma}{\ln(1 - p_k^{burst})} \text{ with } \gamma \geq 0. \quad (3)$$

Based on v_k , we can calculate the probability that the participant will pump the balloon on trial k for pump l , p_{kl}^{pump} :

$$p_{kl}^{pump} = \frac{1}{1 + e^{\tau(l - v_k)}} \text{ with } \tau \geq 0. \quad (4)$$

In this logistic equation, τ is the inverse temperature parameter. The inverse temperature of the choice rule determines how deterministic or random the choice is; The higher τ , the more deterministic. If $l < v_k$, p_{kl}^{pump} becomes greater than 0.5. Similarly, if $l > v_k$, p_{kl}^{pump} becomes less than 0.5. In sum, the four-parameter model has four parameters to be estimated: α, μ, γ , and τ .

We can calculate the likelihood of the data given the parameters by multiplying the probability that the participant will pump on trial k for pump l , p_{kl}^{pump} . The probability of the data given the parameters, $p(D|\alpha, \mu, \gamma, \tau)$, is given by:

$$p(D|\alpha, \mu, \gamma, \tau) = \prod_{k=1}^{k^{last}} \prod_{l=1}^{l_k^{last}} p_{kl}^{pump} \left(1 - p_{k, l_k^{last}+1}^{pump}\right)^{d_k}, \quad (5)$$

where k^{last} is the last number of trials, l_k^{last} is the last number of pumping opportunities on trial k . if the participant transfers the accumulated reward on trial k to the virtual bank account, $d_k = 1$ and if the balloon explodes on trial k , $d_k = 0$.

Although previous researchers have primarily used the four-parameter model, it has been known that α and μ of the four-parameter model are not well recovered. The two-parameter model was proposed to address this problem of the four-parameter model (Pleskac, 2008; van Ravenzwaaij et al., 2011). However, the two-parameter model does not apply to the original BART paradigm because participants are not informed of the probability structure of the task. Therefore, instead of the two-parameter model, we considered a non-learning version of the four-parameter model, which is based on an assumption that participants do not learn during the BART.

2.3.2 Non-learning version (Par3 model)

The non-learning version of the four-parameter model is based on the following assumption of the two-parameter model: participants do not learn during the BART. Since participants do not update their belief of the probability of the balloon exploding based on observations and they are not informed of the probability structure of the original BART paradigm, in this model, we assumed participant's belief of the probability of the balloon exploding as a parameter, θ . Then, similar to Equation (3), we can derive the optimal number of pumps, ν , as follows:

$$\nu = \frac{-\gamma}{\ln(1-\theta)} \text{ with } \gamma \geq 0. \quad (6)$$

Based on the optimal number of pumps, similar to Equation (4), we can calculate the probability that the participant will pump the balloon for pump l , p_l^{pump} :

$$p_l^{pump} = \frac{1}{1+e^{\tau(l-\nu)}} \text{ with } \tau \geq 0. \quad (7)$$

Notably, the optimal number of pumps, ν , and the probability that the participant will pump the balloon for pump l , p_l^{pump} , do not depend on the trial number because the participant's belief of the probability of the balloon exploding is a parameter with a fixed value. In sum, the non-learning version of the four-parameter model includes three free parameters to be estimated: θ , γ , and τ .

2.3.3 Reparametrized version (Par4 model)

We suspected that the strong association between α and μ (the ratio reflects the participant's initial belief and the magnitudes indicate the degree of learning from observations) may be problematic and tested if reparametrizing α and μ would improve the parameter recovery performance of the model.

The parameters α and μ are associated with two processes. The ratio of α to μ , α/μ , refers to the participant's initial belief that pumping will make the balloon explode, and the magnitudes of both α and μ determine the degree of learning from observations. Thus, we wanted to reparametrize them so that each parameter is uniquely associated with just one process. Also, we wanted to remove the constraint that α is less than μ because the constraint might lead to inefficient sampling for Bayesian parameter estimation (see Section 2.4).

For the goal, we reparametrized α and μ into ϕ and η : $\phi = \alpha/\mu$ and $\eta = 1/\mu$. Substituting these parameters into Equation (1) yields:

$$p_k^{burst} = 1 - \frac{\phi + \eta \sum_{i=0}^{k-1} n_i^{success}}{1 + \eta \sum_{i=0}^{k-1} n_i^{pumps}} \text{ with } 0 < \phi < 1, \eta > 0. \quad (8)$$

After the reparameterization, the initial value of p_k^{burst} equals $1 - \phi$. Thus, ϕ indicates the participant's initial belief that pumping will not make the balloon explode. Also, η is an updating coefficient of the participant's belief by the observed data. If $\eta = 0$, p_k^{burst} is not affected by the observed data. If η is very large, p_k^{burst} rapidly comes close to the observed probability of burst. Equations (3) and (4) remain the same. Note that we included the reparametrized four-parameter version in the hBayesDM package as a function named *bart_par4* (Ahn, Haines, & Zhang, 2017).

2.3.4 The Exponential-Weight model (EW model)

The four-parameter model has two critical limitations, even after the reparameterization. First, although the model reflects that the participant updates p_k^{burst} from the prior belief with the observed data, Equation (8) hardly provides an intuitive interpretation of the learning process. We modified the updating equation in an attempt to show the learning process more clearly. Second, the assumption that the participant determines the optimal number of pumps before each trial may be unjustified. This assumption might be contradictory with the primary goal of the BART, measuring the risk-taking tendency of individuals, because it does not reflect some critical features of risk-taking such as loss

aversion and impulsive responses. Instead, the participant may decide whether to pump the balloon or not just before each pump.

To address the first issue, we defined a parameter, $\psi = 1 - \phi$, which is the initial value of p_k^{burst} . Substituting this parameter into Equation (8) yields:

$$p_k^{burst} = \omega_{k-1}\psi + (1 - \omega_{k-1})P_{k-1} \text{ with } 0 < \psi < 1, \eta > 0, \quad (9)$$

where $P_{k-1} = \frac{\sum_{i=0}^{k-1} (n_i^{pumps} - n_i^{success})}{\sum_{i=0}^{k-1} n_i^{pumps}}$, which is the observed probability that pumping has made the balloon explode up to trial $k - 1$, and $\omega_{k-1} = \frac{1}{1 + \eta \sum_{i=0}^{k-1} n_i^{pumps}}$, which is the weight indicating how much weight is given to the prior belief on trial k when estimating the probability of the balloon exploding. Each component of Equation (9) has a clear role, which is interpretable as a part of weight updating learning. Specifically, the current value (p_k^{burst}) is estimated as a weighted average of the initial value (ψ) and the observed value (P_{k-1}). As data accumulates, the participant updates the weight (ω_{k-1}) and the observed value (P_{k-1}). The weight (ω_{k-1}) and the observed value (P_{k-1}) are determined by the total number of the data ($\sum_{i=0}^{k-1} n_i^{pumps}$) and the number of the data that meet a certain condition (explosion, $\sum_{i=0}^{k-1} (n_i^{pumps} - n_i^{success})$). In this framework, η indicates how rapidly the participant depends on experience. If $\eta \rightarrow \infty$, learning entirely depends on the present outcome. If $\eta = 0$, no further learning occurs.

To improve the model performance within this framework, we modified the functional form of the weight, ω_{k-1} . If we define x as $\eta \sum_{i=0}^{k-1} n_i^{pumps}$, in Equation (9), $\omega_{k-1} = \frac{1}{1+x}$, which means the weight is hyperbolic. Other functional forms can be alternatives to the hyperbolic function if they meet two conditions: the participant's learning starts with the prior belief (if $x = 0$, $\omega_{k-1} = 1$) and primarily depends on the observed value after observing enough data (if $x \rightarrow \infty$, $\omega_{k-1} \rightarrow 0$). The exponential decay, e^{-x} , is a reasonable alternative because it meets the two conditions and is commonly used to describe natural phenomena such as the voltage of the resistor-capacitor circuit, the number of remain radioactive atoms, and the concentration of the first-order chemical reaction. Replacing ω_{k-1} with e^{-x} in Equation (9) yields:

$$p_k^{burst} = e^{-\xi \sum_{i=0}^{k-1} n_i^{pumps}} \psi + (1 - e^{-\xi \sum_{i=0}^{k-1} n_i^{pumps}}) P_{k-1} \text{ with } 0 < \psi < 1, \xi > 0. \quad (10)$$

To avoid confusion, we replaced η with ξ and named ξ an updating exponent.

For the second issue (the assumption that the participant determines the optimal number of

pumps before each trial), we tested a new model which assumed that the participant decides whether to pump the balloon or not before each pump instead of each trial. Using the prospect theory (Kahneman & Tversky, 2013), we calculated the subjective utilities for pumping and not-pumping a balloon on trial k for pump l as follows:

$$U_{kl}^{pump} = (1 - p_k^{burst})r^\rho - p_k^{burst}\lambda\{(l-1)r\}^\rho \quad \text{with } 0 < \rho < 2, \lambda > 0, \quad (11)$$

$$U_{kl}^{transfer} = 0, \quad (12)$$

where r is the amount of reward for each successful pump, ρ is risk preference, and λ is loss aversion. Then, we can calculate the probability that the participant will pump the balloon on trial k for pump l , p_{kl}^{pump} , by using these subjective utilities.

$$p_{kl}^{pump} = \frac{1}{1 + e^{\tau(U_{kl}^{transfer} - U_{kl}^{pump})}} \quad \text{with } \tau \geq 0, \quad (13)$$

where τ is inverse temperature. We noticed that this model (Equations (11), (12), and (13)) is similar to a model (Model1) reported in Wallsten et al. (2005). Although Model1 was not the best-fitting model, its model fit was close to that of the winning model (the four-parameter model). Considering that the EW model has a single parameter for risk preference instead of having separate risk preference parameters for gain and loss like Model1, we also included Model1 in the model comparison to examine its performance compared to other models. In summary, the EW model has five free parameters to be estimated: ψ (prior belief of burst), ξ (updating exponent), ρ (risk preference), τ (inverse temperature), and λ (loss aversion).

2.3.5 Model1

In Model1 (Wallsten et al., 2005), the equation calculating the probability of the balloon exploding is the same as the four-parameter model (Equation (1)). Because of the inefficient sampling issue, we utilized the reparametrized version (Equation (4)) instead of Equation (1).

Like the EW model, Model1 is based on an assumption that the participant decides whether to pump the balloon or not before each pump based on the subjective utilities for pumping and not-pumping. The only difference is that Model1 has separate risk preference parameters for gain and loss. The subjective utilities for pumping and not-pumping a balloon on trial k for pump l are calculated as follows:

$$U_{kl}^{pump} = (1 - p_k^{burst})r^{\rho^+} - p_k^{burst}\lambda\{(l-1)r\}^{\rho^-} \quad \text{with } 0 < \rho^+, \rho^- < 2, \lambda > 0, \quad (14)$$

$$U_{kl}^{transfer} = 0, \quad (15)$$

where r is the amount of reward for each successful pump, ρ^+ is risk preference for gain, ρ^- is risk preference for loss, and λ is loss aversion. Then, we can calculate the probability that the participant will pump the balloon by using these subjective utilities (Equation (13)). In summary, Model1 includes six free parameters to be estimated: ϕ (prior belief of success), η (updating coefficient), ρ^+ (risk preference for gain), ρ^- (risk preference for loss), τ (inverse temperature), and λ (loss aversion).

2.3.6 The Exponential-Weight Mean-Variance model (EWMV model)

The existing models and EW model utilize the prospect theory (Kahneman & Tversky, 2013) to calculate subjective utilities for pumping and not-pumping. Given that previous studies have suggested that the performances of the prospect theory and the mean-variance analysis (Markowitz, 1952) are comparable (Boorman & Sallet, 2009; Hens & Mayer, 2014; Levy & Levy, 2004), we tested a model by applying the mean-variance analysis to the EW model. According to the mean-variance analysis, the subjective utility of an option can be formulated by a linear combination of the expected value and the variance of potential outcomes. Applying the mean-variance analysis, we calculated the subjective utilities for pumping and not-pumping a balloon on trial k for pump l as follows:

$$U_{kl}^{pump} = (1 - p_k^{burst})r - p_k^{burst}\lambda(l-1)r + \rho p_k^{burst}(1 - p_k^{burst})\{r + \lambda(l-1)r\}^2 \quad \text{with } \lambda > 0, \quad (16)$$

$$U_{kl}^{transfer} = 0, \quad (17)$$

where r is the amount of reward for each successful pump, ρ is risk preference, and λ is loss aversion. In Equation (16), the first two terms indicate the expected value of a pump, and the last term indicates the impact of the variance of potential outcomes on the subjective utility. Notably, in the mean-variance framework, the risk preference is defined as a coefficient of the variance term, which is a proxy for risk. If $\rho < 0$, the participant prefers to choose an option with a large variance of potential outcomes. If $\rho = 0$, the expected value of an option determines the participant's subjective utility. If $\rho > 0$, the participant prefers to choose an option with a small variance of potential outcomes. Like the EW model, we can calculate the probability of pumping the balloon by using these subjective utilities based on the mean-variance analysis (Equation (13)). In summary, the EWMV model includes five free parameters: ψ (prior belief of burst), ξ (updating exponent), ρ (risk preference), τ (inverse temperature), and λ (loss aversion).

2.4 Hierarchical Bayesian Analysis (HBA)

We used hierarchical Bayesian Analysis (HBA) for parameter estimation (Berger, 2013; Gelman, Carlin, Stern, & Rubin, 2004; Lee, 2011). HBA offers several benefits over conventional non-hierarchical approaches such as individual-level ordinary least squares and maximum likelihood estimation (MLE) methods. First, HBA estimates parameters as posterior distributions instead of point estimates. Posterior distributions provide us with more information about the parameters than point estimates because distributions show the uncertainty of the estimated values. Second, with HBA, we can systematically characterize similarities and differences across subjects within a Bayesian framework based on the amount of information from each individual. Previous studies suggest that HBA allows us to estimate model parameters more accurately than individual- or group-level MLE methods (Ahn, Krawitz, Kim, Busemeyer, & Brown, 2011).

We conducted HBA by using Stan (version 2.15.1), a probabilistic programming language for specifying statistical models (Carpenter et al., 2017). Stan uses Hamiltonian Monte Carlo (HMC) for sampling from high-dimensional parameter space. Specifically, we implemented the models in the hBayesDM (Ahn et al., 2017) environment, which uses Stan. For parameter estimation, we used flat or weakly informative priors (Ahn, Haines, & Zhang, 2017; Haines, Vassileva, & Ahn, 2018) to minimize the influence of the priors and the Matt trick (Papaspiliopoulos, Roberts, & Sköld, 2007) to facilitate the sampling process. We will make the Stan codes and precise prior settings publicly available through GitHub. We tested several other types of priors and confirmed that the priors hardly affected the results as long as the priors are approximately uninformative. We used a large enough sample size (4000 samples, including 2000 burn-in samples per chain) to assure that parameters converge to the target distributions, with four independent chains to check that the posterior distributions are not dependent on initial starting points. Note that Vehtari, Gelman, Simpson, Carpenter, and Bürkner (2019) recommended running at least four chains by default. The trace plots indicated that chains were well mixed and the $\hat{R}$ values (Gelman & Rubin, 1992) for all model parameters were lower than 1.1, which indicates that the estimated parameter values converged to their target posterior distributions. Thinning was not applied because thinning of chains is rarely useful for the precision of estimates (Link & Eaton, 2012). See supplementary material for detailed information of HBA (Table S2, Figure S3, and Figure

S4).

2.5 Model comparison

2.5.1 Leave-one-out information criterion (LOOIC)

LOOIC is an information criterion calculated from the Leave-one-out cross-validation. Leave-one-out cross-validation is a method to estimate out-of-sample prediction accuracy from a fitted Bayesian model based on the log-likelihood evaluated from the posterior distributions (Vehtari, Gelman, & Gabry, 2017). It is well-known that LOOIC has various benefits over simpler estimates such as Akaike Information Criterion (AIC, Akaike, 1998) and Bayesian Information Criterion (BIC, Schwarz, 1978). We used the R package *loo* (Vehtari et al., 2017) to estimate LOOIC for each model. Because LOOIC is calculated from the log-likelihood, the lower LOOIC is, the better its model fit is. For model selection, LOOIC weights are defined as Akaike weights (Wagenmakers & Farrell, 2004) calculated based on LOOIC values, and the detailed information is provided in the supplementary material (see Model comparison section).

2.5.2 Parameter recovery

We also used parameter recovery to evaluate how accurate a model estimates true parameter values from the simulation data generated from the true parameter values (e.g., Ahn et al., 2011; Ahn et al., 2014; Haines et al., 2018; Wagenmakers et al., 2007). For the comparison, we did parameter recovery analysis for the reparametrized version of the four-parameter model, the EW model, and the EWMV model. For each parameter in each competing model, we randomly sampled true parameter values from the truncated normal distribution with a mean and a standard deviation estimated from data of healthy controls, heroin-dependent, and amphetamine-dependent groups to investigate a broader range of parameter values. For each group, we calculated individual posterior means as individual parameter values for participants and used the mean and standard deviation of the individual parameter values as the mean and standard deviation for the truncated normal distribution. Then, we generated simulation data of 30 trials per subject by using the true parameter values. Lastly, we estimated parameter values from the simulation data. Correlations between the true parameter values and the predicted parameter values and regression coefficients were used to evaluate the model performance.

3. Results

3.1 Model comparison

3.1.1 Leave-one-out information criterion (LOOIC)

Table 1 shows the LOOIC for five competing models. In all three groups, the EWMV model was the best-fitting model followed by the EW, Model1, Par4, and Par3 in the order of model fits. The Par4 model outperformed the Par3 model, which is consistent with the result of the previous study, showing that the two-parameter model provided a poorer fit than the four-parameter model when the probability structure was uninformed (Pleskac, 2008). LOOIC weights, the relative likelihoods of the models calculated based on LOOIC values, strongly favored the EWMV model as indicated by its LOOIC weights of 1. Considering that the Model1 includes one more parameter but shows poorer model fits than the EW model, we excluded the Model1 from subsequent analyses. Also, given that the LOOIC values for the Par3 model were much larger than those of the other models (i.e., the Par3 model provided much poorer model fits than the other models), we decided to exclude the Par3 model from further analyses.

Group	Model	LOOIC	Δ LOOIC	LOOIC Weights
HC	EWMV	20297.650	0	1.000
	EW	20454.267	156.617	0.000
	Model1	20488.310	190.660	0.000
	Par4	20666.613	368.963	0.000
	Par3	22482.830	2185.180	0.000
Her	EWMV	7055.653	0	1.000
	EW	7080.670	25.017	0.000
	Model1	7104.407	48.754	0.000
	Par4	7182.598	126.945	0.000
	Par3	7811.477	755.824	0.000
Amp	EWMV	6761.845	0	1.000
	EW	6793.605	31.760	0.000
	Model1	6805.140	43.295	0.000
	Par4	6877.514	115.669	0.000
	Par3	7264.684	502.839	0.000

Table 1. Leave-one-out information criterion for each model. EWMV: the exponential-weight mean-variance model, EW: the exponential-weight model, Model1: a model similar to the EW model from a previous study (Wallsten et al., 2005), Par4: the reparametrized version of the four-parameter model, Par3: the non-learning version of the four-parameter model, HC: healthy control group, Her: heroin-dependent group, Amp: amphetamine-dependent group. Lower LOOIC indicates a better model fit. The LOOIC weight is the relative likelihood of the model calculated

based on its LOOIC.

3.1.2 Parameter recovery

We evaluated the quality of parameter recovery as the correlation between the true and estimated values. Also, we considered the regression coefficients to characterize the degree of associations and biases. In the main text, we report the parameter recovery results from the healthy control group. The parameter recovery results from the heroin- and amphetamine-dependent groups are reported in the supplementary material (Figures S5-S10) and are not qualitatively different from the those from the healthy control group.

Figure 1 shows the results of parameter recovery from the healthy control group for the Par4 model. As shown in Figure 1, overall all parameters were relatively well recovered in the Par4 model including the prior belief of success (ϕ) and the updating coefficient (η), which were not well recovered and systematically overestimated in the previous studies (Heathcote et al., 2015; van Ravenzwaaij et al., 2011). This suggests that our reparameterization may have improved the parameter recovery performance by separating the roles of the two parameters. To directly compare the parameter recovery of the four-parameter model and the Par4 model, we attempted to recover the parameters of the four-parameter model, but the parameters of the original four-parameter model failed to converge even after many (e.g., 4000) burn-in samples. We suggest two possible reasons underlying the failure. First, given that the magnitudes of α and μ commonly indicate the degree of learning from observations, the high correlation between the two parameters might make the sampling process fail to work well even with HMC. Second, the constraint that μ is always larger than α may cause issues in the sampling process.

Figure 2 shows the results of parameter recovery from the healthy control group for the EW model. The EW model showed poorer parameter recovery performance than the Par4 model in two aspects. First, the risk preference (ρ , EW model) exhibited relatively weak recovery compared to the risk-taking propensity (γ , Par4 model). Second, for the loss aversion (λ , EW model), the estimated values were shrunk towards the mean value. This shrinkage effect indicates that the model parameter might not be accurately estimated from the information the data contain. Figures S7 and S8 show that the loss-aversion (λ , EW model) was also not recovered well from the heroin-dependent and amphetamine-dependent groups, which is consistent with the recovery results from the healthy control group.

Figure 3 shows the results of parameter recovery from the healthy control group for the EWMV model. Most parameters, except for the risk preference (ρ , EWMV model), showed good parameter recovery, including the loss aversion parameter (λ , EWMV model), which was not recovered well for the EW model. For the risk preference (ρ , EWMV model), the regression line reveals that the estimated values were slightly shrunk towards the mean value. Although the risk preference (ρ , EWMV model) was not recovered well, considering the EWMV model includes one more parameter than the Par4 model (i.e., the EWMV model is more complex than the Par4 model), we decided to scrutinize the parameter recovery more to accurately compare the parameter recovery performances of the EWMV and Par4 models.

The parameter recovery results for the prior belief and updating rate provide support for the EWMV model. For the Par4 model, some parameter values of the prior belief of success (ϕ) and the updating coefficient (η) deviated from the diagonal in all three groups (Figures 1, S5, and S6). Notably, the updating coefficient (η) showed poor parameter recovery when we used the mean and standard deviation estimated from the amphetamine-dependent group (Figure S6). In contrast, for the EWMV model, most parameter values of the prior belief of burst (ψ) and the updating exponent (ξ) were well recovered (Figures 3, S9, and S10). Considering the LOOIC and parameter recovery results, we selected the EWMV model as the winning model and compared it with the Par4 model in further analyses.

Figure 1. Parameter recovery results for the reparametrized version of the four-parameter model (Par4 model). The red lines denote $y = x$. The blue lines indicate the regression lines of each graph. Shaded regions indicate 95% confidence intervals. The correlation and regression coefficients of each scatter plot is as follows [correlation, slope, intercept]. ϕ (prior belief of success): [0.605, 0.896, 0.102], η (updating coefficient): [0.702, 0.991, 0.003], γ (risk-taking propensity): [0.918, 0.903, 0.079], τ (inverse temperature): [0.882, 0.878, 0.034]. The average of the correlation coefficients is 0.777.

Figure 2. Parameter recovery results for the exponential-weight model (EW model). The red lines denote $y = x$. The blue lines indicate the regression lines of each graph. Shaded regions indicate 95% confidence intervals. The correlation and regression coefficients of each scatter plot is as follows [correlation, slope, intercept]. ψ (prior belief of burst): [0.818, 0.757, 0.003], ξ (updating exponent): [0.764, 0.830, 0.005], ρ (risk preference): [0.666, 0.811, 0.138], τ (inverse temperature): [0.913, 0.973, 0.522], λ (loss aversion): [0.694, 0.260, 3.350]. The average of the correlation coefficients is 0.771.

Figure 3. Parameter recovery results for the exponential-weight mean-variance model (EWMV model). The red lines denote $y = x$. The blue lines indicate the regression lines of each graph. Shaded regions indicate 95% confidence intervals. The correlation and regression coefficients of each scatter plot is as follows [correlation, slope, intercept]. ψ (prior belief of burst): [0.847, 0.810, 0.003], ξ (updating exponent): [0.798, 0.667, 0.005], ρ (risk preference): [0.746, 0.610, 0.000], τ (inverse temperature): [0.812, 0.870, 1.495], λ (loss aversion): [0.933, 0.860, 0.314]. The average of the correlation coefficients is 0.827.

3.1.3 Posterior prediction

We evaluated the posterior prediction performance of each model as a correlation between the observed and simulated adjusted BART scores. For the goal, we estimated parameters for each model from the observed data and generated simulation data by using the estimated parameters. Then, we calculated the adjusted BART score from the simulation data and compared the simulated score with the observed score for each participant.

Figure 4 shows the correlations between observed and simulated adjusted BART scores for the Par4, EW, and EWMV models. All models showed good predictive performance, and their predictive performances are comparable. The regression lines display that the simulated values were slightly shrunk towards the mean value.

Figure 4. Correlations between observed and simulated adjusted BART scores for the models. The red lines denote $y = x$. The blue lines indicate the regression lines of each graph. Shaded regions indicate 95% confidence intervals. The correlation coefficient of each scatter plot is as follows. Par4: 0.770, EW: 0.765, EWMV: 0.792.

3.2 Correlation analysis

To examine whether the EWMV model includes advantageous features of the four-parameter model, we investigated correlations between seemingly corresponding model parameters of the Par4 model and the EWMV model: (γ, ρ) , (ϕ, ψ) , (η, ξ) , and (τ, τ) ; the former one is the parameter of the Par4 model, and the latter one is the parameter of the EWMV model. The risk-taking propensity (γ) and the risk preference (ρ) are related to the risk-taking tendency. The prior belief of success (ϕ) and the prior belief of burst (ψ) correspond to the participant's prior belief about the balloon. The updating coefficient (η) and the updating exponent (ξ) mean updating rate of observation. The two inverse temperatures (τ) reflect how much the participant is deterministic. In the main text, we report the correlations between corresponding model parameters from the healthy control group. The correlation analysis results from the heroin- and amphetamine-dependent groups are reported in the supplementary material (Figures S11 and S12), which show similar patterns with the results from the healthy control group.

Figure 5 shows the correlations between the corresponding parameter pairs. All of the pairs had strong correlations. Although the correlation between the two inverse temperatures (τ) was relatively weak, it is acceptable as they are related to different quantities; one is related to the number of pumps, and the other is related to the subjective utility. The prior belief of success (ϕ) and the prior belief of burst (ψ) were negatively correlated because the sum of the two probabilities should be 1 in an ideal case. The updating coefficient (η) and the updating exponent (ξ) were positively correlated

because both of them represent how rapidly the participant updates the belief based on past experiences. Notably, the risk-taking propensity (γ) and the risk preference (ρ) showed a strong positive correlation, which implicates that, like the risk-taking propensity, the risk preference may reflect risk-taking tendency and be correlated with the frequencies of the past real-world risky behaviors.

Figure 5. Correlations between the corresponding parameter pairs of the models. The blue lines indicate the regression lines of each graph. Shaded regions indicate 95% confidence intervals.

3.3 Group difference

As a way of evaluating the utility of the EWMV model, we applied the EWMV model to healthy and substance-dependent populations (patients with past heroin or amphetamine dependence). We analyzed the group differences of three groups (healthy control, heroin, and amphetamine-dependent groups; see below for the details) for their behavioral performance and the parameter estimates of the EWMV model (we also tested the Par4 model).

3.3.1 Behavioral Performance

The heroin-dependent group displayed a marginally lower adjusted BART score (95% HDI: [-9.73, 0.629], mean = -4.59; 95.9% of the posterior samples were smaller than 0) than the amphetamine-dependent group. The result suggests that heroin users might show lower risk-taking than amphetamine users during the BART. See supplementary material for detailed information on the behavioral performance and the group difference in behavioral performance (Figures S1 and S2).

3.3.2 Model Parameters

We estimated parameters of the EWMV model and the Par4 model for each group separately

to compare the parameter values between the groups. Figure 6 shows the posterior distributions of the group parameters for each group with the EWMV model. The heroin-dependent group displayed credibly lower risk preference (ρ) than the healthy control group (95% HDI of the group difference: [-0.0064, -0.0005], mean: -0.0033) and the amphetamine-dependent group (95% HDI of the group difference: [-0.0073, -0.0013], mean: -0.0042). The amphetamine-dependent group displayed credibly higher risk preference (ρ) than the healthy control group (95% HDI of the group difference: [0.0001, 0.0018], mean: 0.0009). Additionally, the amphetamine-dependent group displayed credibly lower updating exponent (ξ) than the healthy control group (95% HDI of the group difference: [-0.0046, -0.0001], mean: -0.0023).

Figure 7 shows the posterior distributions of the group parameters for each group with the Par4 model. The heroin-dependent group displayed credibly lower risk-taking propensity (γ) than the amphetamine-dependent group (95% HDI of the group difference: [-0.280, -0.012], mean: -0.151). Also, the heroin-dependent group displayed higher inverse temperature (τ) than the amphetamine-dependent group (95% HDI of the group difference: [0.019, 0.070], mean: 0.044). See supplementary material for detail information about the group differences of the model parameters (Figure S13 and S14).

Figure 6. Posterior distributions of the group parameters with the exponential-weight mean-variance model (EWMV model). Tick marks on the bottom and top of each graph indicate 95% highest density intervals (HDIs). Points in the middle of each graph indicate mean values. Asterisks indicate that the 95% HDIs of the posterior distributions of group mean differences do not include zero (group differences were credible). See supplementary material for more information about the group differences of model parameters (Figure S9). ψ : prior belief of burst, ξ : updating exponent, ρ : risk preference, τ : inverse temperature, λ : loss aversion. HC: healthy control group, Her: heroin-dependent group, Amp: amphetamine-dependent group.

Figure 7. Posterior distributions of the group parameters with the reparametrized version of the four-parameter model (Par4 model). Tick marks on the bottom and top of each graph indicate 95% highest density intervals (HDIs). Points in the middle of each graph indicate mean values. Asterisks indicate that the 95% HDIs of the posterior distributions of group mean differences do not include zero (group differences were credible). See supplementary material for more information about the group differences of model parameters (Figure S10). ϕ : prior belief of success, η : updating coefficient, γ : risk-taking propensity, τ : inverse temperature. HC: healthy control group, Her: heroin-dependent group, Amp: amphetamine-dependent group.

The results of the behavioral performance and the model parameters are consistent. Among the three groups, the differences between the heroin-dependent and amphetamine-dependent groups were the most noticeable. The heroin-dependent group displayed a marginally lower adjusted BART score, lower risk preference (ρ), and lower risk-taking propensity (γ) compared to the amphetamine-dependent group. These results consistently show that heroin users show lower risk-taking than amphetamine users during the BART.

4. Discussion

The main focus of this study is on the development of a novel BART model that addresses the limitations of existing models. We proposed a non-learning version of the four-parameter model (Par3 model) and a reparametrized version of the four-parameter model (Par4 model). By modifying equations from the reparametrized version, we developed candidate models and selected the best model (EWMV model) based on the leave-one-out information criterion (LOOIC) and the parameter recovery. The model comparison results suggest that the EWMV model shows better prediction performance across all populations than the other models and good parameter recovery. To examine whether the EWMV

model includes advantageous features of the four-parameter model, we calculated the correlations between corresponding parameter pairs for the Par4 and EWMV models. All of the corresponding parameter pairs had strong correlations, which implies that the EWMV model may include advantageous features of the four-parameter model. As a way of evaluating the utility of the EWMV model, we analyzed differences among substance-using populations in behavioral performance and model parameters of the Par4 and EWMV models. The group differences in behavioral performance and model parameters of the Par4 and EWMV models were consistent. The results of the group differences show that the EWMV model reveals group differences among the groups more clearly than the behavioral performance and the Par4 model, and provides a measure of an additional core psychological construct of risk-taking behavior. Overall, these results suggest that the EWMV model has distinct merits as a computational model for the original BART paradigm.

An important finding of this study is that it suggests a way to improve parameter recovery. We showed that reparametrizing parameters associated with more than one role into parameters with unique roles might help the model recover accurate parameter values. Adequate parameter recovery is a fundamental assumption and necessary for analyzing parameters of a computational model, and it is noteworthy that we can improve parameter recovery by reparameterization alone. At the same time, it is notable that the information criteria such as AIC, BIC, and LOOIC for the reparametrized version and the original model are more or less the same. It suggests that the reparametrized version does not have additional explanatory power compared with the original model. The results demonstrate that parameter recovery and post hoc model fits measured with information criteria reflect different aspects of computational models, and we need to use both methods for comprehensive evaluation.

Besides the superior prediction performance and good parameter recovery performance, the EWMV model also has an advantage that it provides a more interpretable learning process: an agent estimates the present value as a weighted average of the initial and observed value and updates the weight and observed value as data accumulates. In addition, all parameters included in the EWMV model have distinct and interpretable roles. Also, the EWMV model might be applicable to a wide range of cognitive tasks other than the BART. The weight updating learning of the EWMV model is analogous to the Kalman filter, an algorithm to track unknown state variables with uncertainty (Welch & Bishop, 1995). Because the weight updating learning model might be applicable to all situations that include initial states and sequential observations, it might be an alternative to other well-established models to

1 quantify learning situations such as the Rescorla-Wagner model (Rescorla & Wagner, 1972).

2 Utilizing the mean-variance analysis (Markowitz, 1952) is another distinct feature of the EWMV
3 model. Previous studies, which compared the mean-variance analysis and the prospect theory
4 (Kahneman & Tversky, 2013), have suggested that their performances are comparable (Boorman &
5 Sallet, 2009; Hens & Mayer, 2014; Levy & Levy, 2004). However, only a few models (e.g., d'Acremont,
6 Lu, Li, Van der Linden, & Bechara, 2009) directly have utilized the mean-variance analysis to calculate
7 subjective utilities.

8 The group difference results show that the group differences in model parameters of the
9 EWMV model were consistent with the group differences in other indices, including the behavioral
10 performance and model parameters of the Par4 model. One seeming inconsistency is the results of the
11 inverse temperature parameters in the EWMV and Par4 models. For the EWMV model, all groups
12 displayed no credible group differences in the inverse temperature, whereas, for the Par4 model, the
13 heroin-dependent group displayed higher inverse temperature than the amphetamine-dependent group.
14 The reason for this discrepancy may be the two inverse temperatures are related to different quantities.
15 One is related to the subjective utility, whereas the other is related to the number of pumps.
16 Consequently, we did not consider this discrepancy as an inconsistency. The consistent group
17 difference result may indicate that the EWMV model appropriately reflects the participants' risk-taking
18 tendencies in their behaviors. It is also consistent with the results of previous studies showing that
19 opiates (heroin) and stimulants (amphetamine) addictions are behaviorally and neurobiologically
20 distinct (Badiani, Belin, Epstein, Calu, & Shaham, 2011), related to different dopamine modulation
21 mechanisms (Kreek et al., 2012), and characterized by different personality and neurocognitive profiles
22 (Ahn & Vassileva, 2016).

23 Another key implication from the group difference results is that the model parameters of the
24 EWMV model reveal the group differences more clearly than the adjusted BART score and the model
25 parameters of the Par4 model. It is of note that the group difference results still remain valid after
26 considering the bias in the parameter recovery result of the risk preference for the EWMV model
27 because the similar biased patterns appear in all three groups. Namely, the EWMV model aligns better
28 with the clinical function of the BART, whose original purpose is to identify individuals who are prone to
29 take risks. This implies that the EWMV model may be potentially useful for classifying individuals into
30 several clinical groups and establishing quantitative diagnostic criteria for risk-taking behavior.

Providing a measure of loss aversion, which is a core psychological construct of risk-taking behavior, is also advantageous to the EWMV model. Previous studies analyzing risk-taking behavior have consistently shown that loss aversion plays a crucial role in risk-taking behavior, and many computational models of experimental paradigms to investigate risk-taking tendency include parameters of loss aversion (Ahn et al., 2008; Ahn et al., 2011; Sokol-Hessner et al., 2009; Worthy, Pang, & Byrne, 2013). This feature makes the EWMV model comparable with the other computational models that include loss aversion.

In conclusion, we proposed a novel model for the BART, called the exponential-weight mean-variance (EWMV) model, using the weight updating learning and the mean-variance analysis, which addresses the limitations of existing models. The EWMV model outperformed other models in model fits and parameter recovery performance. Also, its distinct merits come with a more interpretable learning process, more salient group differences in model parameters between substance-dependent populations, and the existence of loss aversion parameter. Not limited to the BART, we hope that the weight updating learning model and the mean-variance analysis might apply to other cognitive tasks.

Declarations of interest: none

Acknowledgements

We thank the Seoul Science High School (SSHS) students for assistance in data analysis. Also, we would like to thank all volunteers for their participation in this study. We express our gratitude to Georgi Vasilev, Kiril Bozgunov, Elena Psederska, Dimitar Nedelchev, Rada Naslednikova, Ivaylo Raynov, Emiliya Peneva, and Victoria Dobrojalieva for assistance with recruitment and testing of study participants. The research reported in this publication was supported in part by the National Institute on Drug Abuse and Fogarty International Center under award number R01DA021421 to J.V, and the Basic Science Research Program through the National Research Foundation (NRF) of Korea funded by the Ministry of Science, ICT, & Future Planning (NRF-2018R1C1B3007313 and NRF-2018R1A4A1025891) to W.-Y.A., the Institute for Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2019-0-01367, BabyMind), and the Creative-Pioneering Researchers Program through Seoul National University to W.-Y.A.

References

- Ahn, W. Y., Busemeyer, J. R., Wagenmakers, E. J., & Stout, J. C. (2008). Comparison of decision learning models using the generalization criterion method. *Cognitive science*, 32(8), 1376-1402.
- Ahn, W. Y., Dai, J., Vassileva, J., Busemeyer, J. R., & Stout, J. C. (2016). Computational modeling for addiction medicine: from cognitive models to clinical applications. In *Progress in brain research* (Vol. 224, pp. 53-65): Elsevier.
- Ahn, W. Y., Haines, N., & Zhang, L. (2017). Revealing Neurocomputational Mechanisms of Reinforcement Learning and Decision-Making With the hBayesDM Package. *Comput Psychiatr*, 1, 24-57. doi:10.1162/CPSY_a_00002
- Ahn, W. Y., Krawitz, A., Kim, W., Busemeyer, J. R., & Brown, J. W. (2011). A model-based fMRI analysis with hierarchical Bayesian parameter estimation. *Journal of Neuroscience, Psychology, and Economics*, 4(2), 95-110. doi:10.1037/a0020684
- Ahn, W. Y., Vasilev, G., Lee, S.-H., Busemeyer, J. R., Kruschke, J. K., Bechara, A., & Vassileva, J. (2014). Decision-making in stimulant and opiate addicts in protracted abstinence: evidence from computational modeling with pure users. *Frontiers in Psychology*, 5, 849.
- Ahn, W. Y., & Vassileva, J. (2016). Machine-learning identifies substance-specific behavioral markers for opiate and stimulant dependence. *Drug alcohol dependence*, 161, 247-257.
- Akaike, H. (1998). Information theory and an extension of the maximum likelihood principle. In *Selected papers of hirotugu akaike* (pp. 199-213): Springer.
- Aklin, W. M., Lejuez, C., Zvolensky, M. J., Kahler, C. W., & Gwadz, M. (2005). Evaluation of behavioral measures of risk taking propensity with inner city adolescents. *Behaviour research therapy*, 43(2), 215-228.
- Badiani, A., Belin, D., Epstein, D., Calu, D., & Shaham, Y. (2011). Opiate versus psychostimulant addiction: the differences do matter. *Nature reviews neuroscience*, 12(11), 685.
- Berger, J. O. (2013). *Statistical decision theory and Bayesian analysis*. Springer Science & Business Media.
- Boorman, E. D., & Sallet, J. (2009). Mean–variance or prospect theory? The nature of value representations in the human brain. *Journal of Neuroscience*, 29(25), 7945-7947.
- Busemeyer, J. R., & Stout, J. C. (2002). A contribution of cognitive decision models to clinical assessment: decomposing performance on the Bechara gambling task. *Psychological assessment*, 14(3), 253.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., . . . Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of statistical software*, 76(1).
- d'Acremont, M., Lu, Z.-L., Li, X., Van der Linden, M., & Bechara, A. (2009). Neural correlates of risk prediction error during reinforcement learning in humans. *Neuroimage*, 47(4), 1929-1939.
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, 69(6), 1204-1215.
- Frey, R., Pedroni, A., Mata, R., Rieskamp, J., & Hertwig, R. (2017). Risk preference shares the psychometric structure of major psychological traits. *Science advances*, 3(10), e1701381.

- 1 Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2004). Bayesian data analysis, 2nd edn. Texts in
2 Statistical Science. In: Boca Raton, London, New York, Washington DC: Chapman & Hall, CRC.
- 3 Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences.
4 *Statistical science*, 7(4), 457-472.
- 5 Haines, N., Vassileva, J., & Ahn, W. Y. (2018). The Outcome-Representation Learning Model: A Novel
6 Reinforcement Learning Model of the Iowa Gambling Task. *Cognitive science*, 42(8), 2534-
7 2561.
- 8 Heathcote, A., Brown, S. D., & Wagenmakers, E.-J. (2015). An introduction to good practices in
9 cognitive modeling. In *An introduction to model-based cognitive neuroscience* (pp. 25-48):
10 Springer.
- 11 Hens, T., & Mayer, J. (2014). Cumulative prospect theory and mean variance analysis: a rigorous
12 comparison. *Swiss Finance Institute Research Paper*(14-23).
- 13 Hopko, D. R., Lejuez, C., Daughters, S. B., Aklin, W. M., Osborne, A., Simmons, B. L., & Strong, D. R.
14 (2006). Construct validity of the balloon analogue risk task (BART): Relationship with MDMA
15 use by inner-city drug users in residential treatment. *Journal of Psychopathology and*
16 *Behavioral Assessment*, 28(2), 95-101.
- 17 Kahneman, D., & Tversky, A. (2013). Prospect theory: An analysis of decision under risk. In *Handbook*
18 *of the fundamentals of financial decision making: Part I* (pp. 99-127): World Scientific.
- 19 Kreek, M. J., Levran, O., Reed, B., Schlussman, S. D., Zhou, Y., & Butelman, E. R. (2012). Opiate
20 addiction and cocaine addiction: underlying molecular neurobiology and genetics. *The*
21 *Journal of clinical investigation*, 122(10), 3387-3393.
- 22 Kruschke, J. K. (2013). Bayesian estimation supersedes the t test. *Journal of Experimental Psychology:*
23 *General*, 142(2), 573.
- 24 Lee, M. D. (2011). How cognitive modeling can benefit from hierarchical Bayesian models. *Journal*
25 *of Mathematical Psychology*, 55(1), 1-7.
- 26 Lejuez, C. W., Read, J. P., Kahler, C. W., Richards, J. B., Ramsey, S. E., Stuart, G. L., . . . Brown, R. A.
27 (2002). Evaluation of a behavioral measure of risk taking: The Balloon Analogue Risk Task
28 (BART). *Journal of Experimental Psychology: Applied*, 8(2), 75-84. doi:10.1037/1076-
29 898x.8.2.75
- 30 Lejuez, C. W., Simmons, B. L., Aklin, W. M., Daughters, S. B., & Dvir, S. (2004). Risk-taking propensity
31 and risky sexual behavior of individuals in residential substance use treatment. *Addictive*
32 *behaviors*, 29(8), 1643-1647.
- 33 Levy, H., & Levy, M. (2004). Prospect theory and mean-variance analysis. *Review of Financial Studies*,
34 17(4), 1015-1041.
- 35 Link, W. A., & Eaton, M. J. (2012). On thinning of chains in MCMC. *Methods in ecology and evolution*,
36 3(1), 112-115.
- 37 Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance*, 7(1), 77-91. doi:10.2307/2975974
- 38 Meredith, M., & Kruschke, J. K. (2018). Bayesian Estimation Supersedes the t-Test.
- 39 Papaspiliopoulos, O., Roberts, G. O., & Sköld, M. (2007). A general framework for the parametrization

- of hierarchical models. *Statistical science*, 59-73.
- Pleskac, T. J. (2008). Decision making and learning while taking sequential risks. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(1), 167.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological review*, 85(2), 59.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical conditioning II: Current research theory*, 2, 64-99.
- Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, 6(2), 461-464.
- Sokol-Hessner, P., Hsu, M., Curley, N. G., Delgado, M. R., Camerer, C. F., & Phelps, E. A. (2009). Thinking like a trader selectively reduces individuals' loss aversion. *Proceedings of the National Academy of Sciences*, 106(13), 5035-5040.
- van Ravenzwaaij, D., Dutilh, G., & Wagenmakers, E.-J. (2011). Cognitive model decomposition of the BART: Assessment and application. *Journal of Mathematical Psychology*, 55(1), 94-105. doi:10.1016/j.jmp.2010.08.010
- van Ravenzwaaij, D., & Oberauer, K. (2009). How to use the diffusion model: Parameter recovery of three methods: EZ, fast-dm, and DMAT. *Journal of Mathematical Psychology*, 53(6), 463-473.
- Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics Computing*, 27(5), 1413-1432.
- Vehtari, A., Gelman, A., Simpson, D., Carpenter, B., & Bürkner, P.-C. (2019). Rank-normalization, folding, and localization: An improved $\hat{R}$ for assessing convergence of MCMC. *arXiv preprint arXiv:1903.08008*.
- Wagenmakers, E.-J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic bulletin & review*, 11(1), 192-196.
- Wagenmakers, E.-J., Van Der Maas, H. L., & Grasman, R. P. (2007). An EZ-diffusion model for response time and accuracy. *Psychonomic bulletin review*, 14(1), 3-22.
- Wallsten, T. S., Pleskac, T. J., & Lejuez, C. W. (2005). Modeling behavior in a clinically diagnostic sequential risk-taking task. *Psychol Rev*, 112(4), 862-880. doi:10.1037/0033-295X.112.4.862
- Welch, G., & Bishop, G. (1995). An introduction to the Kalman filter. In: Citeseer.
- Worthy, D. A., & Maddox, W. T. (2014). A comparison model of reinforcement-learning and win-stay-lose-shift decision-making processes: A tribute to WK Estes. *Journal of Mathematical Psychology*, 59, 41-49.
- Worthy, D. A., Pang, B., & Byrne, K. A. (2013). Decomposing the roles of perseveration and expected value representation in models of the Iowa gambling task. *Frontiers in Psychology*, 4, 640.