Decision
2014, Vol. 1, No. 3, 184–191

2325-9965/14/$12.00 http://dx.doi.org/10.1037/dec0000016

To Simulate or Not? Comment on Steingroever, Wetzels, and Wagenmakers (2014)

Emmanouil Konstantinidis
and Maarten Speekenbrink
University College London

Julie C. Stout
Monash University

Woo-Young Ahn
Virginia Tech Carilion Research Institute,
Roanoke, Virginia

David R. Shanks
University College London

Steingroever, Wetzels, and Wagenmakers (2014) conducted a detailed investigation of 3 popular reinforcement-learning models for the Iowa gambling task using 2 model comparison techniques: a post hoc fit criterion and a simulation method. However, these 2 methods yield inconsistent results regarding which model should be preferred as a description of underlying psychological processes. Here, we describe the benefits of each method in an attempt to develop a more balanced view of how to utilize these model comparison techniques, and we outline the risks of focusing on a single method to make inferences about the overall utility of a model. Also, we make several suggestions about how applied research should evaluate candidate cognitive models, and we offer guidelines for future research aimed at identifying "good" models for decomposing and explaining participants' performance.

Keywords: experience-based decision-making, Iowa gambling task, mathematical modeling, reinforcement learning

One of the most popular and frequently used paradigms in experience-based decision-making is the Iowa gambling task (IGT; Bechara, Damasio, Damasio, & Anderson, 1994), which has become a test case for the development of learning models, as well as a standard screening tool for decision-making deficits in clinical populations. In this issue, Steingroever, Wetzels, and Wagenmakers (SWW, 2014) placed three reinforcement-learning (RL) models (Sutton & Barto, 1998) for the IGT under careful scrutiny: the Expectancy Valence-Learning (EVL; Busemeyer & Stout, 2002), the Prospect Valence-Learning (PVL; Ahn, Busemeyer, Wagenmakers, & Stout, 2008; Ahn, Krawitz, Kim, Busemeyer, & Brown, 2011), and the PVL-Delta (Ahn et al., 2008; Fridberg et al., 2010) models. Each of these models consists of different assumptions relating to the formation of utilities, the learning or updating of the expectancies of each deck, and the tradeoff between exploration and exploitation. The focus of SWW's examination is centered around two major points: first, which of these models should be preferred in order to isolate and identify the psychological processes that drive performance on the IGT, and second, which method should be used to select the best model.

Selection of the best-performing model generally involves sophisticated model comparison techniques that assess each model's ability to

Emmanouil Konstantinidis and Maarten Speekenbrink, Division of Psychology and Language Sciences, University College London; Julie C. Stout, School of Psychological Sciences, Monash University; Woo-Young Ahn, Virginia Tech Carilion Research Institute, Roanoke, Virginia; David R. Shanks, Division of Psychology and Language Sciences, University College London.

Correspondence concerning this article should be addressed to Emmanouil Konstantinidis, Department of Experimental Psychology, Division of Psychology and Language Sciences, University College London, 26 Bedford Way, London WC1H 0AP, England. E-mail: emmanouil.konstantinidis.09@ucl.ac.uk

reproduce the observed behavioral patterns and explain fundamental aspects of decision-making. Two methods are commonly used: a post hoc fit criterion (or one-step-ahead predictions) and a simulation method. However, as shown by SWW (see also Ahn et al., 2008; Yechiam & Busemeyer, 2008; Yechiam & Ert, 2007), these two methods yield inconsistent results as to which model should be preferred, a finding considered by SWW to indicate failure of at least one of the model comparison techniques (that is, one technique is better than the other—or more indicative—at disentangling the psychological processes that drive performance on the IGT). Following this, SWW argue that models' performance should be assessed using both methods (models have to pass a minimum threshold of adequacy under both methods), but the simulation method provides a better assessment of the underlying psychological processes because it is the only method that bases its predictions on newly generated payoff schedules and choices and thus reflects what participants would have done under novel conditions (i.e., new payoff sequences). They argue that the post hoc fit criterion is inferior in that it may only be able to reproduce the observed behavioral pattern by mimicking participants' previous choices, indicating that the estimated parameters may be biased by factors that do not invite a clear psychological interpretation.

In this comment, we argue against the idea that the inconsistency of model comparison methods is a failure. Instead, as we outline, such inconsistency in the results should be expected given the difference in approach of these two methods, a point also made by SWW: The post hoc fit criterion uses participants' *previous history of choices and experienced payoffs* to make predictions about subsequent trials, whereas the simulation method relies *on the generation of new histories for choices and payoffs*. Second, while we agree that the simulation method is important, we argue that the post hoc fit criterion also has value in that its parameters are meaningful and provide useful psychological insight.

In the remainder of the commentary, we outline and present theoretical and practical arguments for why relying solely on simulation performance in an experience-based decision-making setting may lead to misleading conclusions, and we highlight valuable aspects of the post hoc criterion method. By articulating the benefits of each approach, we offer a more balanced view of how to utilize these model comparison methods to yield the richest interpretations of behavior.

Sequential Dependencies

The main objective of RL models of the IGT is to account for how future decisions are shaped by the experienced history of previous decisions and their associated payoffs. This is how learning occurs, and these models are evaluated on how well they capture learning effects throughout an individual's observed choice history. In other words, these models try to account for the sequential dependencies between each current choice and the previous choices and payoffs. In fact, all three models under consideration share the assumption that the dependence of a current choice on previous choices is fully mediated by the payoffs (and not the past choices) experienced as a consequence of these previous choices. Thus, to predict a future choice, one only needs to consider the history of experienced payoffs.

In the simulation method employed by SWW, random draws from a posterior distribution over parameters, conditional upon observed choices and payoffs, are used to simulate a set of new choice and payoff sequences. Model selection is then based on the match between the marginal probabilities of the artificial choices in this simulated set and the proportions of observed choices. Our view, however, is that by marginalizing over the simulated payoff sequences and focusing only on the choice probabilities, the method ignores the fundamental objective of RL models, which is to account for sequential dependencies between choices (i.e., how choices are shaped by experience). Furthermore, by considering the ability to reproduce (postdict) marginal probabilities in the *same* dataset as used for model estimation, the simulation method employed by SWW can favor models that are unlikely to generalize to new datasets. For example, consider a model with one parameter, which identifies a sequence of choices within the set of all possible choice sequences. The model predicts that a choice sequence displayed by a participant will be identical to the sequence identified by the parameter value. While, for a standard IGT with 100 trials and 4 choice alternatives, the parameter has 4^{100} possible values, estimation is easy once a participant's complete choice sequence is known. Because the model is not stochastic, it will reproduce, for each participant, their choice sequence exactly, regardless of the payoff sequence generated in the simulations. Averaging over participants, such simulations will perfectly reproduce the observed choice probabilities. Nevertheless, most people will agree that this is not an overly useful model: It predicts that people make exactly the same choices when performing the task another time and it can only *post*dict choices, not *pre*dict subsequent choices from the previous history. Because of this inability to make one-step-ahead predictions, the post hoc fit method, unlike a simulation method that estimates a model and evaluates simulation performance with the same data, will clearly disfavor this hypothetical model.

For now, the claim that the simulation method offers a more useful method of model discrimination than the post hoc fit method seems at least premature.

Individual Differences and Research Applied to Clinical Populations

The IGT has been used extensively as a neuropsychological test to assess decision-making in clinical populations (for a review, see, e.g., Bechara & Damasio, 2005; Dunn, Dalgleish, & Lawrence, 2006). Thus, it is of great importance to present a model that can offer informative conclusions regarding individual differences in the underlying psychological processes. This can enable researchers to identify key differences between clinical groups and healthy controls and make strong connections between neurophysiology and behavior, leading to a better characterization of the psychological symptoms of the disorder under test (Stout, Busemeyer, Lin, Grant, & Bonson, 2004; Yechiam, Busemeyer, Stout, & Bechara, 2005). A serious drawback of the simulation method that SWW employed is that it evaluates the average model prediction and ignores the fact that different models are required to fit different individuals.

A fundamental goal of clinical research is to establish an explanatory framework for particular clinical groups, to connect pathological behavior to patterns of behavior on clinical and cognitive experimental tasks, and to map these deficits to neurophysiological mechanisms. Computational models serve as the intermediate step between the brain and observed behavior where performance on a task can be decomposed into its constituent cognitive processes and mapped to neural mechanisms (Busemeyer, Stout, & Finn, in press). In this regard, model comparison should also be informed by psychological measures from relevant clinical assessments, which will facilitate the selection of the best models. In other words, a good model will not only perform well on a quantitative statistical index (e.g., goodness of fit) but will also provide explanations and make connections to findings or observations from the existing literature and to validated characteristics of a particular clinical sample. For example, if the model parameters are correlated with clinically relevant characteristics derived from psychometric scales and personality questionnaires, then the model can serve as a good representation of the underlying psychological processes. Also, another criterion to assess model performance is individual parameter consistency (SWW also propose the use of this analysis); that is, comparison of correlations of model parameters estimated from the same individual in more than one task. This is an important step toward the identification of "stable" internal characteristics that drive performance on different tasks and can be used as an extra assessment of model performance (Yechiam & Busemeyer, 2008). This stability of the internal characteristics driving individual performance across somewhat disparate tasks helps to reinforce psychological explanations regarding underlying pathology in individuals.

Generalization, Biases, and Inertia

One of the most fundamental properties of a good model is its ability "to make predictions about what will be observed in the future or generalizations about what would be observed under altered circumstances" (Shiffrin, Lee, Kim, & Wagenmakers, 2008, p. 1249). Similar views have been expressed by other researchers who have discussed the importance of generalization in identifying a good candidate model and have introduced methods to assess a model's generalizability, such as the minimum description length (Pitt, Kim, & Myung, 2003; Pitt, Myung, & Zhang, 2002) and the generalization criterion (Busemeyer & Wang, 2000). If an (estimated) model is to reliably predict behavior in new settings, it should apply independently of task-specific effects, biases that may arise from the experimental setup, and idiosyncratic strategies or heuristics individuals may adopt to deal with uncertainty. In experience-based tasks there are two sources of information that drive participants' choices: frequency of past choices from different options and payoffs experienced from sampling each of these options (Yechiam & Ert, 2007). According to SWW, a bias-free model (one that generalizes efficiently) has to base its predictions only on past payoffs. We note that this is logical because the parameters of the RL models under investigation measure underlying psychological processes related to how participants respond to the payoffs they experience. However, we argue that the goal of predicting behavior in different contexts is not well served by simulation methods.

Adding to this argument, Erev and Haruvy (2005) distinguish between two types of predictions in descriptive learning models: first, predictions for a task that are based on the interaction of a participant with that same task (*within-game* predictions), and second, predictions for different tasks with which participants are unfamiliar (*new-game* predictions). Following this distinction, the post hoc fit criterion is a within-game prediction, whereas the simulation method can be seen as a new-game prediction. SWW argue that the simulation performance of a model should be of greater importance when our goal is to find the best model or assess the psychological processes underlying performance on the IGT. While we argue that simulation is indeed an important comparison technique, we present arguments why (a) simulation is a crude generalization test, which (b) may not reflect stable psychological processes.

Not a Direct Generalization Test

SWW suggest that the simulation method provides a good test of generalizability because it assesses models' predictions under new payoff sequences that participants have not encountered. To make these new predictions, the simulation method uses the best-fitting parameters from one-step-ahead predictions (post hoc fit criterion).[1] Two possible problems may arise from the application of the simulation method: first, the use of parameters that optimize one-step-ahead predictions to predict new unobserved choice sequences (new-game predictions) might be far from ideal because these parameters carry information about participants' history of experience with the task and serial dependencies. In other words, the use of parameters estimated with the post hoc fit method could result in a considerable underestimation of a model's simulation performance. Future research can examine the benefits of the simulation method (i.e., model predictions under novel conditions/payoff sequences) by using simulated participants to estimate the model parameters at an individual or aggregate level (a common practice in other experience-based tasks, see Erev & Barron, 2005; Gonzalez & Dutt, 2011). Then, these parameters can be compared with those from the one-step-ahead method (and the model predictions of each method) in order to provide better inferences regarding the level of generalizability and utility of each model.

The second point is related to model complexity, a model's capacity to fit different patterns of data (Pitt & Myung, 2002). Complexity is dependent on two different factors that affect model fit: the free parameters of the model and its functional form; that is, how the parameters and the mathematical equations are combined (Myung, 2000). Complexity is a very important concept in model comparison and is related to the generalizability of the model. SWW argue that the simulation method should be preferred over the post hoc fit criterion because it leads to a better assessment of generalizability, but it does not take into account either of the two factors described earlier. For example, the EVL model has 3 free parameters, whereas the PVL and the PVL-Delta have 4. Even though the differences in simulation performance between the PVL and PVL-Delta models cannot be ascribed to differences in the number of parameters or their functional form (there are qualitative differences in the patterns of choices that the two models predict), it could be the case that more (or less) complex models can be applied to

[1] SWW used Hierarchical Bayesian estimation for the model parameters, which also accounts for the uncertainty in the parameter estimates.

IGT data (see, e.g., the Value-Plus-Perseveration model; Worthy, Pang, & Byrne, 2013).

The question is how does one deal with complexity in model comparison? While there are a few techniques to capture complexity (see Pitt et al., 2002), we propose the use of the generalization criterion (GC) method (Busemeyer & Wang, 2000). The reason is that the GC is a better and more sophisticated modification of the cross-validation method and is sensitive to both factors of complexity (Pitt & Myung, 2002). Also, it has already been applied for model comparison in experience-based tasks (see Ahn et al., 2008; Yechiam & Busemeyer, 2008).[2] A critical difference between the GC and the simulation method is that the former uses the participants' actual choice and payoff histories (i.e., there is intermediate feedback from participants' choices), whereas the latter takes no input from what participants have actually selected and observed. Hence, when we assess model generalizability, the GC is a more appropriate method because it deals with the concept of model complexity and also takes into account how participant's future choices are dependent on their actual choice and payoff histories.

Updating of Expectancies, Choice Mimicry, and Inertia

One pronounced difference among the models that SWW examined relates to the updating of each deck's expectancy. The EVL and PVL-Delta models make use of the delta rule whereas the PVL uses a decay rule (Erev & Roth, 1998) to update the expectancies. In the case of the delta learning rule, only the expectancy of the selected alternative on each trial is updated while the expectancies of the unchosen alternatives remain as they were. Under the decay RI rule, the expectancy of each alternative changes as a function of the time and frequency of selections.

Yechiam and Ert (2007) evaluated the reliance of each learning rule on the past choices that each individual had made. The learning rule that relies least on past choices should yield a better measure of the payoff-related variables and thus be able to provide better predictions in new situations. Yechiam and Ert showed that when information about the payoffs is eliminated, (i.e., each alternative/option yields the exact same payoff—equal payoff series extraction method; EPSE) the decay RI rule was superior to the delta rule under the post hoc criterion method, indicating its ability to base its predictions on participants' choice history. Using the same technique, Konstantinidis, Busemeyer, Speekenbrink, and Shanks (2014) show that the model that has the worst performance under the EPSE method (PVL-Delta) is the best model under the simulation method.

One reasonable question following the previous demonstration is whether the decay RI rule of the PVL model can capture stable psychological characteristics of each individual or just favors the selection of the decks that have been most frequently sampled. SWW suggest that choice mimicry is an undesirable characteristic of a model because it does not convey meaningful information about the psychological elements that drive performance on the task. When information about individuals' actual choices is removed (through the simulation method), then the decay RI rule does not predict the observed choice patterns as accurately as the delta rule.

However, choice inertia is a well-documented tendency in these kinds of tasks (Erev & Haruvy, 2005) according to which people show a propensity to repeat their last selections irrespective of the received payoffs and consequently of the expected value of each option/deck. Inertia is a manifestation of risk-taking behavior (Dutt & Gonzalez, 2012), and it has been implemented in various cognitive models of experience-based decision making (e.g., Biele, Erev, & Ert, 2009; Gonzalez & Dutt, 2011; Nevo & Erev, 2012). The application of

[2] The application of the criterion is straightforward: there are two experimental conditions or datasets. The first serves as the calibration set and the second as the test set. A model is fitted on the calibration set, and the estimated parameters are used to predict the data in the test set. Finally, the predictions of each of the candidate models in the test set are compared to assess the empirical validity of these models. Ahn et al. (2008) used the GC to evaluate different RL models in the context of experience-based decision-making. Specifically, the same group of participants completed two tasks, the IGT and the Soochow gambling task (SGT; Chiu et al., 2008), and the models were compared on their ability to predict one task's choices based on the estimated parameters of the other task. Using the GC, Ahn et al. found that the same pattern of results emerged as in the one-step-ahead prediction method; that is, the PVL model with the decay RI rule was the best-performing model (but see Yechiam & Busemeyer, 2008).

inertia can enhance the utility of cognitive models by making them more ecologically valid (Dutt & Gonzalez, 2012) and can improve predictions of serial dependencies in people's decisions (Erev & Haruvy, 2005).

Concluding Remarks

Selection among competing cognitive models is a core challenge for understanding cognitive processing. SWW attempted to tackle this issue in one of the most frequently used decision-making paradigms by using two methods: a post hoc fit criterion and a simulation method. While we agree that model selection should not be based only on the merits of model fit (post hoc fit criterion), the suggestion that the simulation method is more informative about whether a model captures the underlying psychological processes is problematic, for the reasons we have outlined. From a formal-technical perspective, the simulation method cannot provide a strong test of generalization (a point also hinted at by SWW), nor does it assess model complexity. According to Pitt et al. (2002) "only by taking complexity into account can a selection method accurately measure a model's generalizability" (p. 474). In this regard, we suggest that future evaluations of model performance and generalizability should employ measures that deal directly with model complexity, such as the GC. Second, the best model under the simulation method may not provide an adequate description of the stable psychological characteristics of each individual. Rather, the estimated parameters from one-step-ahead predictions may carry over artifacts generated from the experimental design. For this reason, even parameter consistency techniques might fail to provide an accurate answer if the parameters have been estimated on incorrect a priori assumptions about the mechanisms of the cognitive process under investigation.

These problems lead us to reconsider whether the existing models are able to capture the most fundamental aspects of experience-based decision-making. One suggestion would be to extend the RL models in order to account for effects such as inertia and perseveration. Even though the decay RI rule makes no explicit assumptions about inertia, the steady decrease in value expectations of the unselected options favors the selection of the option that has been selected more often. One problem with the decay RI rule is that inertia is confounded with the expectancy of each option. Because both of these dimensions are represented by a single numerical value, it is very difficult to ascertain which of these tendencies is responsible for a model's predictions (Worthy et al., 2013). The VPP model by Worthy et al. (2013) seems a step in the right direction (because it explicitly incorporates inertia), although it has to be tested further before making any strong conclusions about its overall utility. Specifically, the VPP model has 8 parameters, which may be too much given the small numbers of trials in the IGT (usually 100). Moreover, the behavioral pattern from experimental tasks may not be sufficient to extract all the relevant information regarding the underlying psychological processes. A careful examination and comparison of the model parameters with established psychometric measures and personality scales would reveal the extent to which the model captures observed behavior and the validity of its assumptions.

To summarize, SWW's attempt at a rigorous examination of RL models in the IGT can be seen as the starting point of a more careful investigation of model performance. However, the results and conclusions from their analyses are premature for establishing and identifying one single model based on which applied research should assess decision-making. In this comment, we highlighted the drawbacks of focusing on a single method, and we offered suggestions about how future research should tackle the most challenging issue of model selection in experience-based decision-making. Specifically, for assessments of generalizability, we suggested the use of the GC, individual parameter consistency, and parameter estimation using simulated participants; for psychological interpretation of model predictions, we proposed a direct comparison between the estimated parameters and measurements from psychological scales and questionnaires; finally, a better assessment of each model's basic properties is needed using model and parameter recovery techniques (e.g., Wagenmakers, Ratcliff, Gomez, & Iverson, 2004; Wetzels, Vandekerckhove, Tuerlinckx, & Wagenmakers, 2010) and tests of specific influence (e.g., Steingroever, Wetzels, & Wagenmakers, 2013).

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,* 1376–1402. doi:10.1080/03640210802352992

Ahn, W.-Y., Krawitz, A., Kim, W., Busemeyer, J. R., & Brown, J. (2011). A model-based fMRI analysis with hierarchical Bayesian parameter estimation. *Journal of Neuroscience, Psychology, and Economics, 4,* 95–110. doi:10.1037/a0020684

Bechara, A., & Damasio, A. R. (2005). The somatic marker hypothesis: A neural theory of economic decision. *Games and Economic Behavior, 52,* 336–372. doi:10.1016/j.geb.2004.06.010

Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition, 50,* 7–15. doi:10.1016/0010-0277(94)90018-3

Biele, G., Erev, I., & Ert, E. (2009). Learning, risk attitude and hot stoves in restless bandit problems. *Journal of Mathematical Psychology, 53,* 155–167. doi:10.1016/j.jmp.2008.05.006

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,* 253–262. doi:10.1037/1040-3590.14.3.253

Busemeyer, J. R., Stout, J. C., & Finn, P. (in press). Using computational models to help explain decision making processes of substance abusers. In D. Barch (Ed.), *Cognitive and affective neuroscience of psychopathology*. Oxford, England: Oxford University Press.

Busemeyer, J. R., & Wang, Y. M. (2000). Model comparisons and model selections based on generalization criterion methodology. *Journal of Mathematical Psychology, 44,* 171–189. doi:10.1006/jmps.1999.1282

Chiu, Y.-C., Lin, C.-H., Huang, J.-T., Lin, S., Lee, P.-L., & Hsieh, J.-C. (2008). Immediate gain is long-term loss: Are there foresighted decision makers in the Iowa Gambling Task? *Behavioral and Brain Functions, 4,* 13. doi:10.1186/1744-9081-4-13

Dunn, B. D., Dalgleish, T., & Lawrence, A. D. (2006). The somatic marker hypothesis: A critical evaluation. *Neuroscience and Biobehavioral Reviews, 30,* 239–271. doi:10.1016/j.neubiorev.2005.07.001

Dutt, V., & Gonzalez, C. (2012). The role of inertia in modeling decisions from experience with instance-based learning. *Frontiers in Psychology, 3,* 177. doi:10.3389/fpsyg.2012.00177

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review, 112,* 912–931. doi:10.1037/0033-295X.112.4.912

Erev, I., & Haruvy, E. (2005). Generality, repetition, and the role of descriptive learning models. *Journal of Mathematical Psychology, 49,* 357–371. doi:10.1016/j.jmp.2005.06.009

Erev, I., & Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *The American Economic Review, 88,* 848–881.

Fridberg, D. J., Queller, S., Ahn, W.-Y., Kim, W., Bishara, A. J., Busemeyer, J. R., . . . Stout, J. C. (2010). Cognitive mechanisms underlying risky decision-making in chronic cannabis users. *Journal of Mathematical Psychology, 54,* 28–38. doi:10.1016/j.jmp.2009.10.002

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review, 118,* 523–551. doi:10.1037/a0024558

Konstantinidis, E., Busemeyer, J. R., Speekenbrink, M., & Shanks, D. R. (2014). *Cognitive modeling of experience-based decision-making: The differential effects of past payoffs and past choices*. Manuscript in preparation.

Myung, I. J. (2000). The importance of complexity in model selection. *Journal of Mathematical Psychology, 44,* 190–204. doi:10.1006/jmps.1999.1283

Nevo, I., & Erev, I. (2012). On surprise, change, and the effect of recent outcomes. *Frontiers in Psychology, 3,* 24. doi:10.3389/fpsyg.2012.00024

Pitt, M. A., Kim, W., & Myung, I. J. (2003). Flexibility versus generalizability in model selection. *Psychonomic Bulletin & Review, 10,* 29–44. doi:10.3758/BF03196467

Pitt, M. A., & Myung, I. J. (2002). When a good fit can be bad. *Trends in Cognitive Sciences, 6,* 421–425. doi:10.1016/S1364-6613(02)01964-2

Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review, 109,* 472–491. doi:10.1037/0033-295X.109.3.472

Shiffrin, R. M., Lee, M. D., Kim, W., & Wagenmakers, E.-J. (2008). A survey of model evaluation approaches with a tutorial on hierarchical Bayesian methods. *Cognitive Science, 32,* 1248–1284. doi:10.1080/03640210802414826

Steingroever, H., Wetzels, R., & Wagenmakers, E. -J. (2013). Validating the PVL-Delta model. *Frontiers in Decision Neuroscience, 4,* 898. doi:10.3389/fpsyg.2013.00898

Steingroever, H., Wetzels, R., & Wagenmakers, E. -J. (2014). Absolute performance of reinforcement-learning models for the Iowa gambling task. *Decision, 1,* 161–183.

Stout, J. C., Busemeyer, J. R., Lin, A., Grant, S. J., & Bonson, K. R. (2004). Cognitive modeling analysis of decision-making processes in cocaine abusers. *Psychonomic Bulletin & Review, 11,* 742–747. doi:10.3758/BF03196629

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: The MIT Press.

Wagenmakers, E.-J., Ratcliff, R., Gomez, P., & Iverson, G. J. (2004). Assessing model mimicry using the parametric bootstrap. *Journal of Mathematical Psychology, 48,* 28–50. doi:10.1016/j.jmp.2003.11.004

Wetzels, R., Vandekerckhove, J., Tuerlinckx, F., & Wagenmakers, E.-J. (2010). Bayesian parameter estimation in the expectancy valence model of the Iowa gambling task. *Journal of Mathematical Psychology, 54,* 14–27. doi:10.1016/j.jmp.2008.12.001

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. doi:10.3389/fpsyg.2013.00640

Yechiam, E., & Busemeyer, J. R. (2008). Evaluating generalizability and parameter consistency in learning models. *Games and Economic Behavior, 63,* 370–394. doi:10.1016/j.geb.2007.08.011

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,* 973–978. doi:10.1111/j.1467-9280.2005.01646.x

Yechiam, E., & Ert, E. (2007). Evaluating the reliance on past choices in adaptive learning models. *Journal of Mathematical Psychology, 51,* 75–84. doi:10.1016/j.jmp.2006.11.002

Received February 18, 2014
Revision received May 29, 2014
Accepted June 9, 2014 ■