



UvA-DARE (Digital Academic Repository)

Safe models for risky decisions

Steingröver, H.M.

Publication date

2017

Document Version

Other version

License

Other

[Link to publication](#)

Citation for published version (APA):

Steingröver, H. M. (2017). *Safe models for risky decisions*. [Thesis, fully internal, Universiteit van Amsterdam].

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Validating the PVL-Delta Model for the Iowa Gambling Task

This chapter has been published as:
Helen Steingroever, Ruud Wetzels, and Eric-Jan Wagenmakers (2013).
Validating the PVL-Delta model for the Iowa gambling task.
Frontiers in Decision Neuroscience, 4, 898.¹

Abstract

Decision-making deficits in clinical populations are often assessed with the Iowa gambling task (IGT). Performance on this task is driven by latent psychological processes, the assessment of which requires an analysis using cognitive models. Two popular examples of such models are the Expectancy Valence (EV) and Prospect Valence Learning (PVL) models. These models have recently been subjected to sophisticated procedures of model checking, spawning a hybrid version of the EV and PVL models—the PVL-Delta model. In order to test the validity of the PVL-Delta model we present a parameter space partitioning (PSP) study and a test of selective influence. The PSP study allows one to assess the choice patterns that the PVL-Delta model generates across its entire parameter space. The PSP study revealed that the model accounts for empirical choice patterns featuring a preference for the good decks or the decks with infrequent losses; however, the model fails to account for empirical choice patterns featuring a preference for the bad decks. The test of selective influence investigates the effectiveness of experimental manipulations designed to target only a single model parameter. This test showed that the manipulations were successful for all but one parameters. To conclude, despite a few shortcomings, the PVL-Delta model seems to be a better IGT model than the popular EV and PVL models.

The Iowa gambling task (IGT; Bechara et al., 1994) is arguably the most popular neuropsychological paradigm to assess decision-making deficits in clinical populations. In order to isolate and identify the psychological processes that drive performance on the IGT, behavioral analyses of IGT data are insufficient. A promising alternative analysis approach is to use cognitive

¹The final publication is available at <http://dx.doi.org/10.3389/fpsyg.2013.00898>.

process models. The IGT imposes high demands on these models because it is a complex task producing various types of choice patterns that a good model should be able to generate (Steingroever, Wetzels, Horstmann, et al., 2013; Steingroever, Wetzels, & Wagenmakers, 2013a). In addition, the models should also account for individual differences and for participants' switch behavior on the task (e.g., Steingroever, Wetzels, & Wagenmakers, 2013a; Zhao & Costello, 2007). Despite the high demands, some plausible and elegant IGT models have been proposed. Two of the most frequently used representatives include the Expectancy Valence model (EV; see Steingroever, Wetzels, & Wagenmakers, 2013a; for references), and the Prospect Valence Learning model (PVL; see Steingroever, Wetzels, & Wagenmakers, 2013a; for references and a detailed description of the models). The parameters of these models correspond to distinct psychological processes such as motivation, learning/memory, and response consistency (Busemeyer, Stout, & Finn, 2003).

Since the development of the EV model in 2002, reinforcement-learning (RL) models for IGT data have been subjected to sophisticated procedures of model checking (e.g., Ahn et al., 2008; Busemeyer & Stout, 2002; Fridberg et al., 2010; Steingroever, Wetzels, & Wagenmakers, 2013a; Yechiam & Busemeyer, 2005; Yechiam & Ert, 2007; Yechiam & Busemeyer, 2008). These model comparison efforts spawned a hybrid version of the EV and PVL models—the PVL-Delta model (Ahn et al., 2008; Fridberg et al., 2010; Steingroever, Wetzels, & Wagenmakers, 2014; see next section for a detailed description of the PVL-Delta model and recent model comparison efforts). This model seems to be promising for IGT data because it can generate a variety of empirical choice patterns better than its competitors (Steingroever et al., 2014).

Whereas previous procedures of model checking focused mostly on relative comparisons of different RL models for IGT data, no efforts have been carried out to validate the PVL-Delta model (i.e., assess its adequacy in isolation). Here, we focus on two different ways of validating the PVL-Delta model: First, we conduct a parameter space partitioning (PSP) study that systematically assesses which choice patterns the PVL-Delta model generates across its entire parameter space. Thus, with this first validity check we aim to answer the question: Can the PVL-Delta model generate typical empirical choice patterns over a wide range of parameter settings? Second, we conduct a test of selective influence that investigates the effectiveness of experimental manipulations designed to target only one of the model parameters. Thus, with this second validity check we aim to answer the question: Do the parameters of the PVL-Delta model indeed correspond to the proposed psychological processes?

The outline of this article is as follows. In the first section, we explain the IGT, outline the PVL-Delta model, and review previous efforts to compare RL models for IGT data. In the second and third section, we present the PSP study and the test of selective influence. In the last section, we summarize our findings and discuss their ramifications. To anticipate our results, our PSP study shows that the PVL-Delta model can account for empirical choice patterns featuring a preference for the good decks or the decks with infrequent losses; however, the model fails to account for empirical choice patterns featuring a preference for the bad decks. Our test of selective influence shows that the manipulations were successful for all but one parameters.

4.1 The Iowa Gambling Task and the PVL-Delta Model

The Iowa Gambling Task

In this section we describe the IGT (see also Steingroever, Wetzels, & Wagenmakers, 2013a; Steingroever et al., 2014). The purpose of the IGT is to measure decision-making deficits of clinical populations in an experimental setting. In the traditional IGT, participants are initially given \$2000 facsimile money and are presented with four decks of cards. Participants are instructed to choose

Table 4.1: *Payoff scheme of the traditional IGT as developed by Bechara et al. (1994).*

	Deck A	Deck B	Deck C	Deck D
	Bad deck with fre- quent losses	Bad deck with infre- quent losses	Good deck with fre- quent losses	Good deck with infre- quent losses
Reward/trial	100	100	50	50
Number of losses/10 cards	5	1	5	1
Loss/10 cards	-1250	-1250	-250	-250
Net outcome/10 cards	-250	-250	250	250

cards in order to maximize their long-term net outcome (Bechara et al., 1994, 1997). Unbeknownst to the participants, the task typically contains 100 trials. After each choice, participants receive feedback on the rewards and the losses (if any) associated with that card, and the running tally.

The task aims to determine whether participants learn to prefer the good, safe decks over the bad, risky decks because this is the only choice pattern that maximizes the long-term net outcomes. The good, safe decks are typically labeled C and D, whereas the bad, risky decks are labeled A and B. Table 4.1 presents the traditional payoff scheme as developed by Bechara et al. (1994). This table illustrates that decks A and B yield high immediate, constant rewards, but even higher unpredictable, occasional losses: hence, the long-term net outcome is negative. Decks C and D, on the other hand, yield low immediate, constant rewards, but even lower unpredictable, occasional losses: hence, the long-term net outcome is positive. In addition to the different payoff magnitudes, the decks also differ in the frequency of losses: Two decks yield frequent losses (decks A and C) and two decks yield infrequent losses (decks B and D).

The PVL-Delta Model

In this section, we describe the PVL-Delta model in detail. The model formalizes participants' performance on the IGT through the interaction of four model parameters that represent distinct psychological processes (Ahn et al., 2008; Fridberg et al., 2010; Steingroever et al., 2014).

The first model assumption is that after choosing a card from deck $k \in \{1, 2, 3, 4\}$ on trial t , participants evaluate the net outcome associated with the just-chosen card by means of a non-linear utility function from prospect theory (Tversky & Kahneman, 1992)—the Prospect Utility function:

$$u_k(t) = \begin{cases} X(t)^A & \text{if } X(t) \geq 0 \\ -w \cdot |X(t)|^A & \text{if } X(t) < 0. \end{cases} \quad (4.1)$$

Here $X(t)$ represents the net outcome on trial t , that is, the sum of the experienced reward and loss (i.e., $X(t) = W(t) - |L(t)|$). The Prospect Utility function contains the first two model parameters—the shape parameter $A \in [0, 1]$, that determines the shape of the utility function, and the loss aversion parameter $w \in [0, 5]$. As A approaches zero, the shape of the utility function approaches a step function. The implication of such a step function is that given a positive net outcome $X(t)$, all utilities are similar because they approach one, and given a negative net outcome $X(t)$, all utilities are also similar because they approach $-w$. On the other hand, as A approaches one, the subjective utility $u_k(t)$ increases in direct proportion to the net outcome, $X(t)$. A value of w larger than one indicates a larger impact of negative net outcomes than positive net outcomes on the subjective utility, whereas a value of w approaching one indicates identical impact of negative

net outcomes and positive net outcomes. As w approaches zero, the model predicts that negative net outcomes will be neglected.

The PVL-Delta model further assumes that, after having formed the utility of the just chosen deck through Equation 4.1, decision makers update their expected utility of the just chosen deck, while keeping the expected utilities of the remaining decks unchanged. This updating process is described by the Delta learning rule:

$$Ev_k(t) = Ev_k(t - 1) + a \cdot (u_k(t) - Ev_k(t - 1)). \quad (4.2)$$

The Delta learning rule states that the expected utility of the chosen deck k is adjusted upward if the experienced utility $u_k(t)$ is higher than expected. If the experienced utility $u_k(t)$ is lower than expected, the expected utility of deck k is adjusted downward.² This updating process is influenced by the third model parameter—the updating parameter $a \in [0, 1]$. This parameter quantifies the memory for rewards and losses. A value of a close to zero indicates slow forgetting and weak recency effects, whereas a value of a close to one indicates rapid forgetting and strong recency effects.

In the next step, the model assumes that the expected utilities of each deck guide participants' choices on the next trial $t + 1$. This assumption is formalized by the softmax choice rule, also known as the ratio-of-strength choice rule. The PVL-Delta model uses this rule to compute the probability of choosing each deck on each trial (Luce, 1959; Equation 4.3). This rule contains a sensitivity parameter θ that indexes the extent to which trial-by-trial choices match the expected deck utilities. Values of θ close to zero indicate random choice behavior (i.e., strong exploration), whereas large values of θ indicate choice behavior that is strongly determined by the expected utilities (i.e., strong exploitation).

$$P[S_k(t + 1)] = \frac{e^{\theta \cdot Ev_k(t)}}{\sum_{j=1}^4 e^{\theta \cdot Ev_j(t)}} \quad (4.3)$$

The PVL-Delta model assumes a trial-independent sensitivity parameter θ , which depends on the final model parameter: the response consistency $c \in [0, 5]$ (Equation 4.4). Small values of c cause a random choice pattern, whereas large values of c cause a deterministic choice pattern.

$$\theta = 3^c - 1 \quad (4.4)$$

In sum, the PVL-Delta model has four parameters: (1) The shape parameter A , which determines the shape of the utility function, (2) the loss aversion parameter w , which quantifies the weight of net losses over net rewards, (3) the updating parameter a , which determines the memory for past expectancies, and (4) the response consistency parameter c , which determines the amount of exploitation versus exploration.

Previous Comparisons of RL Models

This section reviews previous model comparison studies. These studies compared the EV model, PVL model, and alternative RL models using a large variety of methods, for instance: the post hoc fit criterion (i.e., Ahn et al., 2008; Busemeyer & Stout, 2002; Fridberg et al., 2010; Yechiam & Busemeyer, 2005; Yechiam & Ert, 2007; Yechiam & Busemeyer, 2008);³ the simulation method (i.e., Ahn et al., 2008; Fridberg et al., 2010; Steingroever et al., 2014; Worthy, Hawthorne, & Otto,

²We initialized the expectancies of each deck k to zero, $Ev_k(0) = 0$.

³The post hoc fit criterion is also known as the one-step-ahead prediction method.

2013); tests of generalizability (i.e., Ahn et al., 2008; Yechiam & Busemeyer, 2005; Yechiam & Ert, 2007; Yechiam & Busemeyer, 2008); tests of parameter consistency (i.e., Yechiam & Busemeyer, 2008); and PSP (i.e., Steingroever, Wetzels, & Wagenmakers, 2013a).⁴

The above model comparison studies revealed many positive properties of RL models: First, RL models predict the choices on the next trial better than a Bernoulli baseline model (Ahn et al., 2008; Busemeyer & Stout, 2002; Fridberg et al., 2010; Yechiam & Busemeyer, 2005; Yechiam & Ert, 2007; Yechiam & Busemeyer, 2008).⁵ Second, parameters from the RL models estimated from one RL task can be used to predict performance on a different RL task (Ahn et al., 2008; Yechiam & Busemeyer, 2005; Yechiam & Ert, 2007; Yechiam & Busemeyer, 2008). Third, the loss aversion parameter and the updating parameter of the EV model are stable across different tasks (Yechiam & Busemeyer, 2008). Fourth, the estimated model parameters can be used to improve the prediction of group membership (i.e., chronic cannabis users vs. healthy controls; Fridberg et al., 2010).

These positive properties confirm that cognitive modeling analyses are indeed useful to learn more about the psychological processes that drive performance on the IGT. However, previous model comparison studies also revealed that, even though the EV and PVL models are frequently used, they fail to outperform their competitors consistently. It appears that the performance of the RL models depends on the data set and the method used to assess model performance (i.e., fit performance vs. simulation performance; see Steingroever et al., 2014, for a more detailed discussion on previous comparisons of RL models).

Instead of accepting the EV and PVL models as default models to describe IGT data, there is growing evidence that the PVL-Delta model may be a promising alternative IGT model: First, Ahn et al. (2008) showed that the PVL-Delta model results in the best simulation performance (i.e., prediction of the entire sequence of choices on the IGT under a new, unobserved payoff sequence) among the EV model, PVL model, and any combination of the components of the two models. Second, Fridberg et al. (2010) showed that, in two data sets, the PVL-Delta model outperforms the EV model in terms of post hoc fit and simulation performance. Third, Steingroever et al. (2014) showed that, among the EV, PVL, and PVL-Delta models, the PVL-Delta model is the only model that adequately generated the choice patterns shown by seven IGT data sets.

Even though the PVL-Delta model has recently come to the fore as a promising model for IGT data, it has not yet been sufficiently validated. Our goal here is to pursue two methods of validating the PVL-Delta model: a PSP study and a test of selective influence.

4.2 Parameter Space Partitioning

Method

We performed a PSP study to evaluate the flexibility of the PVL-Delta model (Pitt et al., 2006, Pitt et al., 2008; see also Steingroever, Wetzels, & Wagenmakers, 2013a, who performed a PSP study of the EV model, PVL model, and another hybrid model: the EV model with Prospect Utility function). The PSP method systematically assesses the choice patterns predicted by the PVL-Delta model across its entire parameter space. A model is overly flexible when it can generate not only all choice patterns that are observed empirically, but also choice patterns that are logically

⁴Note that the PSP study of Steingroever, Wetzels, and Wagenmakers (2013a) did not focus on the PVL-Delta model, but on the EV model, the PVL model, and another hybrid model: the EV model with the Prospect Utility function.

⁵The Bernoulli baseline model assumes that a participant's probability of choosing a given deck on a given trial equals the overall proportion of choices the participant actually made from that deck.

possible, but never observed. Instead, one should prefer a less flexible, parsimonious model that –ideally– only generates choice patterns that are also frequently observed in experiments (Pitt et al., 2006, 2008).

Note that PSP is a global method (i.e., the full range of parameter values is considered), whereas the other methods that were used to compare RL models are local (i.e., assessment at a particular point in the model’s parameter space; for instance, post hoc fit criterion, simulation method, tests of generalizability, and tests of parameter consistency). The advantage of global methods is that they enable one to assess the full range of choice patterns a model can generate, whereas the results of local methods always depend on the idiosyncrasies of any single data set (Pitt et al., 2006, 2008).

Pitt et al. (2006) describe a new search algorithm to implement PSP. In our implementation we did not use their sophisticated search algorithm, but followed the conceptual idea of PSP, and used a grid search that works as follows (see also Steingroever, Wetzels, & Wagenmakers, 2013a): For each parameter of the PVL-Delta model, we chose 60 values that were equally spaced over the corresponding parameter range. Each combination of these parameter values was used to generate data for 100 synthetic participants completing a 100-trial IGT. For all analyses in this paper, we scaled the traditional payoffs of the IGT as presented in Table 4.1 by dividing by 100 (cf. Ahn et al., 2011).

The generated data were used to analyze which choice patterns the PVL-Delta model can generate across its entire parameter space. Such analysis naturally requires a definition of choice patterns. Here we used two different definitions—the “broad definition of choice patterns” and the “restricted definition of choice patterns”. These definitions are the same as used by Steingroever, Wetzels, and Wagenmakers (2013a).

Broad definition of choice patterns

The “broad definition of choice patterns” is intended to provide a general idea of which choice patterns the PVL-Delta model can generate. Following Steingroever, Wetzels, and Wagenmakers (2013a), we defined five possible choice patterns: (1) Preference for the good decks over bad decks (i.e., $\{C, D\} \succ \{A, B\}$), (2) preference for the bad decks over good decks (i.e., $\{A, B\} \succ \{C, D\}$), (3) preference for the decks with infrequent losses over decks with frequent losses (i.e., $\{B, D\} \succ \{A, C\}$), (4) preference for the decks with frequent losses over decks with infrequent losses (i.e., $\{A, C\} \succ \{B, D\}$), and (5) remaining choice patterns. For each parameter combination, we computed the proportion of choices from each deck averaged across all 100 trials and all 100 repeated data generations. These average choice proportions were then sorted to determine the generated rank order of deck preferences for each parameter combination. Finally, we computed the proportion of the entire parameter space occupied by each of the defined choice patterns. Even though we defined five possible types of choice patterns, we assume based on the theory underlying the IGT (Bechara et al., 1994, 1997) and our IGT review (Steingroever, Wetzels, Horstmann, et al., 2013) that a good model for IGT data should only generate the first three types of choice patterns.

Restricted definition of choice patterns

Note that the broad definition of choice patterns only considers the rank order of the overall proportions of choices from each deck averaged over 100 repeated data generations with the same parameter combination. This means that it does not matter whether the PVL-Delta model generated, for example, a very strong or a very weak preference for the good decks over the bad decks. Both generated choice patterns are classified as the choice pattern “good decks over bad

Table 4.2: Mean proportions of choices from each deck and mean proportions of switches during the last 50 trials of healthy participants showing a pronounced deck preference (see Table 4 in Steingroever, Wetzels, & Wagenmakers, 2013a). Healthy participants are selected from the Steingroever, Wetzels, Horstmann, et al. (2013) data pool ($N = 394$).

Choice pattern	N	Deck A [sd]	Deck B [sd]	Deck C [sd]	Deck D [sd]	Switches during the last 50 trials [25%, 75% quantile] (min, max)
$(C + D) \geq .65$	54	.10 [.05]	.14 [.05]	.36 [.17]	.40 [.14]	.35 [.08, .52] (0.00, 0.96)
$(A + B) \geq .65$	18	.25 [.07]	.52 [.11]	.11 [.05]	.12 [.06]	.43 [.31, .58] (0.10, 0.86)
$(B + D) \geq .65$	139	.12 [.05]	.37 [.12]	.13 [.05]	.39 [.12]	.47 [.28, .66] (0.02, 1.00)

decks” (i.e., $\{C, D\} \succ \{A, B\}$). To go beyond this coarse classification, we also analyzed the model’s behavior when confronted with pronounced deck preferences. To get an indication of pronounced deck preferences shown by healthy participants on the IGT, we used Steingroever, Wetzels, and Wagenmakers (2013a)’s definition of pronounced deck preferences: Specifically, Steingroever, Wetzels, and Wagenmakers (2013a) searched their IGT data pool ($N = 394$; Steingroever, Wetzels, Horstmann, et al., 2013) for healthy participants that chose at least 65% cards from either the good decks (i.e., $(C + D) \geq .65$), the bad decks (i.e., $(A + B) \geq .65$), or the decks with infrequent losses (i.e., $(B + D) \geq .65$). By using the .65-criterion, Steingroever, Wetzels, and Wagenmakers (2013a) included healthy participants with pronounced deck preferences and excluded healthy participants with random choice behaviors. For each of these three groups, Steingroever, Wetzels, and Wagenmakers (2013a) computed the mean proportions of choices from each deck (as shown in Table 4.2). For instance, participants classified to the group “pronounced preference for the good decks” chose on average 36 cards from deck C and 40 cards from deck D. Note that 53.6% of all participants in the Steingroever, Wetzels, Horstmann, et al. (2013) data pool showed a pronounced deck preference by making at least 65% choices from the two most preferred decks. This empirical popularity of pronounced deck preferences underscores how important it is that a RL model for the IGT is able to produce such choice patterns.

Table 4.2 thus provides an indication of pronounced deck preferences shown by healthy participants on the IGT. We used the mean proportion of choices from these three constructed groups for our second, restricted definition of choice patterns. Specifically, we define a pronounced preference for the good decks as at least 36 and 40 choices from decks C and D, respectively; we define a pronounced preference for the bad decks as at least 25 and 52 choices from decks A and B, respectively; and we define a pronounced preference for the decks with infrequent losses as at least 37 and 39 choices from decks B and D, respectively. Based on our simulations, we then determined the proportion of the parameter space of the PVL-Delta model that produced choice patterns that satisfy this second, restricted definition.

Switch behavior

Finally, a good RL model for the IGT should also capture the switches participants make on the IGT (Zhao & Costello, 2007). Steingroever, Wetzels, and Wagenmakers (2013a) therefore determined the mean proportion of switches during the last 50 trials for the three groups of healthy participants

showing pronounced decks preferences (revisited here in the last column of Table 4.2). The table contains for each of the three groups of healthy participants with pronounced choice patterns the mean proportion of switches during the last 50 trials and statistics quantifying the distribution of switch proportions (i.e., the interquartile range and the minimum and maximum switch proportions during the last 50 trials). This information is visualized by the boxplots shown in the left column of Figure 4.1. From Table 4.2 and Figure 4.1 it is evident that, in general, in all three groups participants switch frequently. However, the interquartile ranges and the minimum and maximum proportion of switches during the last 50 trials also indicate that there is large variability in the proportion of switches, such that the switch behavior of healthy participants varies between no switches at all to switches on every trial. This tendency to switch frequently, but also the large individual differences in the switch behavior of healthy participants are illustrated by Figures 4.2, 4.4, and 4.6 (see also Figures 3, 7, 10 in Steingroever, Wetzels, & Wagenmakers, 2013a) which show the trial-by-trial choices (i.e., choice profiles) of representative healthy participants with a pronounced preference for the good decks, bad decks, and decks with infrequent losses, respectively.⁶

We investigated whether the PVL-Delta model captures the empirical switch behavior by comparing the empirical and generated mean proportions of switches during the last 50 trials. Specifically, the generated mean proportions of switches were obtained by determining the mean proportions of switches during the last 50 trials for all parameter combinations that produced pronounced deck preferences. The code for the PSP study is available on www.helensteingroever.com.

Results

Broad definition of choice patterns

Table 4.3 presents the proportion of the parameter space of the PVL-Delta model occupied by each of the five different types of choice patterns. From this table, it is evident that the PVL-Delta model can generate all five different types of choice patterns. However, if we consider its partitioned parameter space more closely, we detect substantial differences between the popularity of the different choice patterns: The choice pattern “good decks over bad decks” is the most central to the model’s overall performance, as this choice pattern occupies the largest part of the model’s parameter space. The second and third largest part of its parameter space are occupied by the choice patterns “remaining” and “infrequent losses over frequent losses”. It is thus evident that choice patterns that are typically shown by healthy participants –the choice patterns “good decks over bad decks” and “infrequent losses over frequent losses” (Steingroever, Wetzels, Horstmann, et al., 2013)– occupy a major part of the model’s parameter space.

Table 4.3 also shows that the choice pattern “bad decks over good decks” is only generated over a minor part of the model’s parameter space. We have therefore grounds to conclude that this choice pattern is uncharacteristic of the PVL-Delta model, and is thus almost irrelevant to its overall performance (Pitt et al., 2006). This finding is important because the choice pattern “bad decks over good decks” is considered characteristic for participants with decision-making deficits (e.g., patients with lesions to the ventromedial prefrontal cortex; Bechara et al., 1994, 1997). These patients are thought to display decision-making deficits on the IGT because their inability to foresee the long-term consequences of their choice behavior leads them to only focus on the immediate rewards.

⁶See Steingroever, Wetzels, and Wagenmakers (2013a) for the deck selection profiles of all healthy participants that showed a pronounced deck preference (i.e., at least 65% choices from the two most preferred decks).

Table 4.3: *Proportions of choice patterns generated by the PVL-Delta model.*

Choice pattern		Proportion of all choice patterns
Good \succ bad decks	$\{C, D\} \succ \{A, B\}$.596
Bad \succ good decks	$\{A, B\} \succ \{C, D\}$.006
Infr. \succ frequent losses	$\{B, D\} \succ \{A, C\}$.118
Frequent \succ infr. losses	$\{A, C\} \succ \{B, D\}$.005
Remaining		.274

Restricted definition of choice patterns

Table 4.4 presents the proportion of all choice patterns generated by the PVL-Delta model that satisfy the restricted definition of choice patterns. The table also presents the mean and standard deviation of the parameter combinations that generated these pronounced deck preferences. The table shows that only minor parts of the parameter space of the PVL-Delta model are occupied by the three types of pronounced choice patterns, even though these patterns are frequently observed in experiments. For instance, 139 healthy participants from the Steingroever, Wetzels, Horstmann, et al. (2013) data pool (35.3%) show a pronounced preference for the decks with infrequent losses (i.e., $(B + D) \geq .65$). However, the PVL-Delta model only generates this choice pattern over 1.6% of its parameter space.

Switch behavior

In addition to the generated choice proportions, we also determined the generated proportion of switches during the last 50 trials for all parameter combination that satisfy the restricted definition of choice patterns (Columns 2 – 6 of Table 4.4). We averaged these generated switch proportions separately for each of the three types of pronounced deck preferences (last column of Table 4.4).

Table 4.4: *Proportion of choice patterns generated by the PVL-Delta model that satisfy the restricted definition of choice patterns. Note that this definition is only based on the mean proportion of choices of the two strongest preferred decks (first column). For the selected choice patterns, the corresponding mean and standard deviation of the model parameters, and the mean proportion of switches during the last 50 trials are presented.*

Choice pattern	Proportion of all choice patterns	A [sd]	w [sd]	a [sd]	c [sd]	Switches during the last 50 trials [25%, 75% quantile] (min, max)
$C \geq .36, D \geq .40$.0084	.66 [.21]	.62 [.38]	.30 [.21]	3.49 [.98]	.0571 [.0014, .0558] (.00, .5724)
$A \geq .25, B \geq .52$.0000028	.92 [.09]	.02 [.03]	.06 [.03]	3.07 [.40]	.0043 [.0003, .0055] (.00, .0210)
$B \geq .37, D \geq .39$.0162	.27 [.24]	.34 [.39]	.43 [.27]	2.80 [.88]	.0705 [.0034, .0918] (.00, .6450)

The table also contains statistics quantifying the distribution of the generated switch proportions, that is, the interquartile range and the minimum and maximum proportion of switches during the last 50 trials. This information is visualized by the right column of Figure 4.1.

When comparing the generated and observed mean proportion of switches during the last 50 trials given pronounced deck preferences, it is apparent that the PVL-Delta model underestimates the observed switch proportions, that is, the generated mean proportion of switches equals or falls below .07 for all generated pronounced choice patterns, whereas the observed mean proportion of switches equals or exceeds .35 for all observed pronounced choice patterns (Tables 4.2 and 4.4). In addition, for all three types of pronounced choice patterns, the interquartile range of the observed proportion of switches exceeds the interquartile range of the model-generated proportion of switches (Figure 4.1, Table 4.2 and 4.4). However, the largest generated switch proportion given a pronounced preference for the good decks and the decks with infrequent losses, respectively, lie within the corresponding interquartile ranges of the observed switch proportions. This suggests that for a few parameter combinations, the PVL-Delta model meets both empirical regularities—pronounced deck preferences and a tendency to switch frequently.

To illustrate the differences and commonalities between the data and the predictions, we plot in Figures 4.2 - 4.7 observed and generated choice profiles. Figures 4.2, 4.4, and 4.6 show the choice profiles of representative healthy participants with a pronounced preference for the good decks, bad decks, and decks with infrequent losses, respectively. Figures 4.3, 4.5, and 4.7 show the choice profiles that were generated with those parameter combinations that resulted in a pronounced preference for the good decks, bad decks, and decks with infrequent losses, respectively, and the maximum number of switches during the last 50 trials. From the figures it is evident that there are large discrepancies between the observed and generated choice profiles in the case of the pronounced preference for the bad decks: The PVL-Delta model generates a few switches in the beginning of the IGT and then exploitation of a single deck, even though healthy participants keep switching across the entire IGT. However, the observed and generated choice profiles look very similar in the case of the pronounced preference for the good decks and the decks with infrequent losses.

To conclude, many healthy participants from the Steingroever, Wetzels, Horstmann, et al. (2013) data pool (53.6%) showed pronounced deck preferences, that is, a pronounced preference for the good decks ($(C + D) \geq .65$), a pronounced preference for the bad decks ($(A + B) \geq .65$), or a pronounced preference for the decks with infrequent losses ($(B + D) \geq .65$) (Table 4.2). This empirical popularity of pronounced deck preferences is only partly reflected by the PVL-Delta model; the model produces choice patterns that satisfy the restricted definition of choice patterns only within minor parts of its parameter space (Table 4.4). In addition, healthy participants in general show many switches during the last 50 trials. However, the PVL-Delta model in general predicts that participants who show pronounced deck preferences switch rarely during the last 50 trials; all generated mean proportion of switches during the last 50 trials equal or fall below .07 whereas the observed mean proportions of switches lie around .40. But compared to the popular EV and PVL models (Steingroever, Wetzels, & Wagenmakers, 2013a), the PVL-Delta model performs better: The PVL-Delta model generates higher mean proportions of switches than its two competitors for almost all pronounced choice patterns; the only exception is that the EV model generates a higher mean proportion of switches for the choice pattern featuring a pronounced preference for the bad decks than the PVL-Delta model.

Moreover, healthy participants show large individual differences in the proportion of switches during the last 50 trials, such that their switch behavior varies between no switches at all to switches on every trial. However, the PVL-Delta model tends to generate very few switches, given pronounced deck preferences, and fails to generate large proportion of switches (i.e., switch proportions higher than .65). But compared to the popular EV and PVL models (Steingroever,

Figure 4.1: Boxplots of observed and generated proportions of switches during the last 50 trials, given a pronounced deck preference. Each row presents the results for different pronounced choice patterns: First row: Pronounced preference for the good decks; second row: Pronounced preference for the bad decks; third row: Pronounced preference for the decks with infrequent losses. The first column presents the switches of 211 healthy participants selected from the Steingroever, Wetzels, Horstmann, et al. (2013) data pool (cf. Table 4.2). The second column presents the switches generated by the PVL-Delta model (cf. Table 4.4).

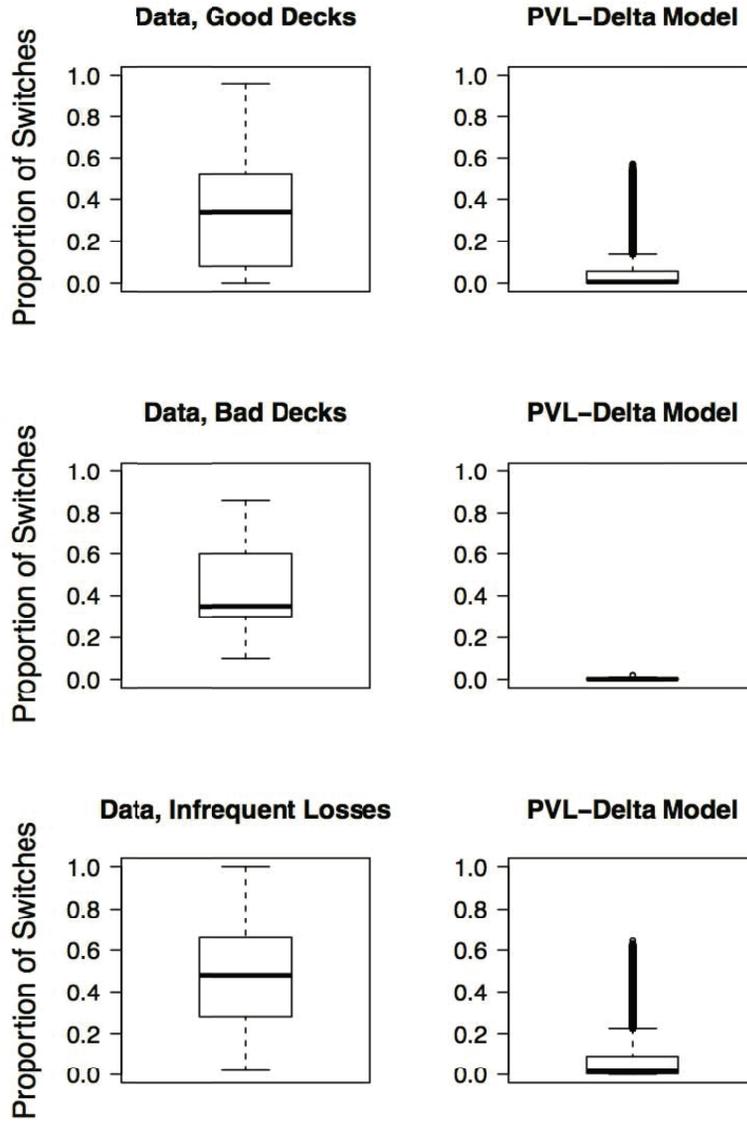


Figure 4.2: Deck selection profiles of four healthy participants showing a pronounced preference for the good decks. The filled dots indicate the occurrence of rewards and losses together; the empty dots indicate the occurrence of only rewards.

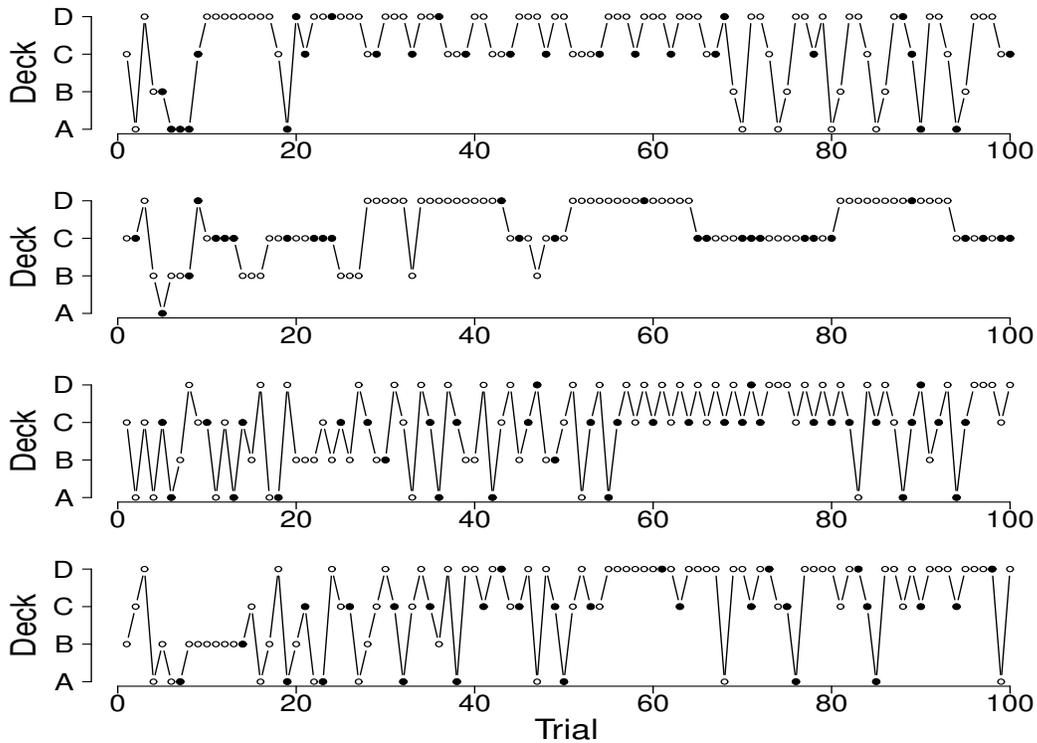


Figure 4.3: Deck selection profiles of four synthetic participants showing a pronounced preference for the good decks (generated by the PVL-Delta model; $A = 0.88, w = 0.68, a = 0.25, c = 1.27$).

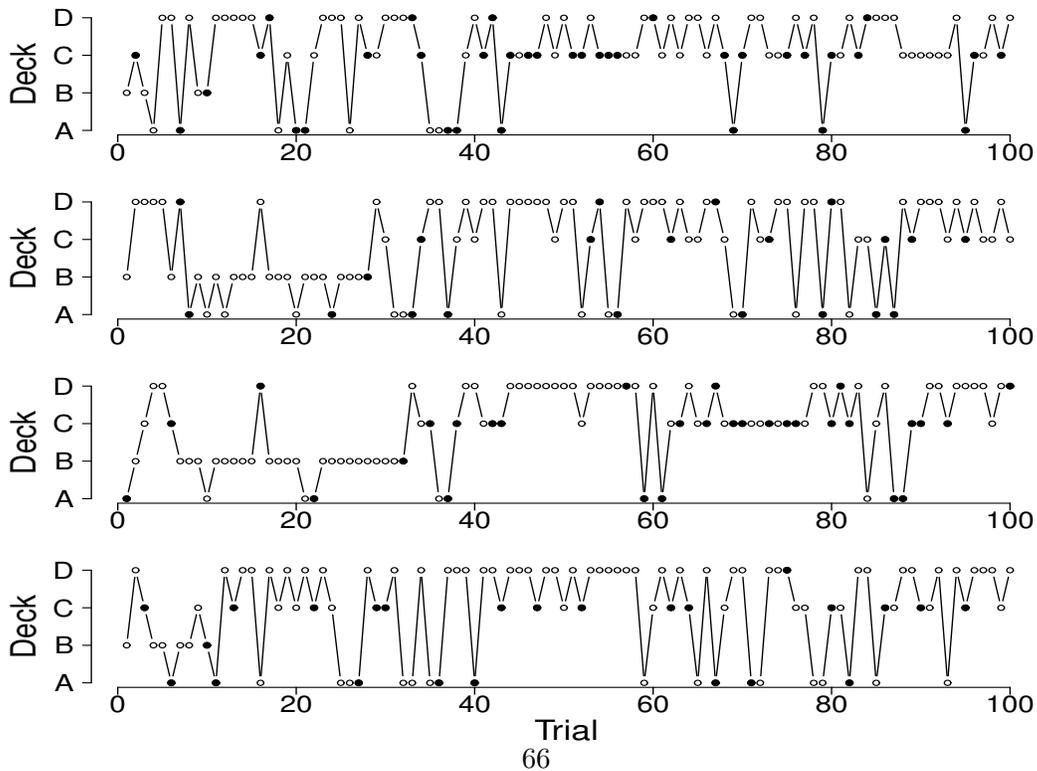


Figure 4.4: Deck selection profiles of four healthy participants showing a pronounced preference for the bad decks.

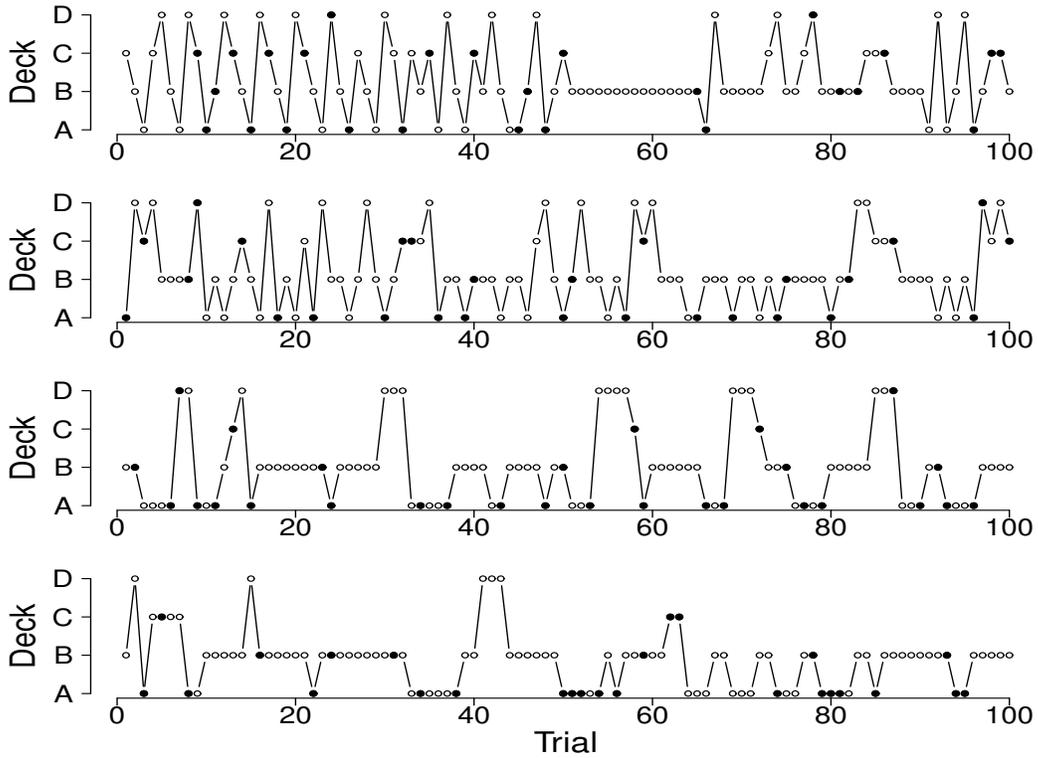


Figure 4.5: Deck selection profiles of four synthetic participants showing a pronounced preference for the bad decks (generated by the PVL-Delta model; $A = 1.00, w = 0.08, a = 0.05, c = 2.71$).

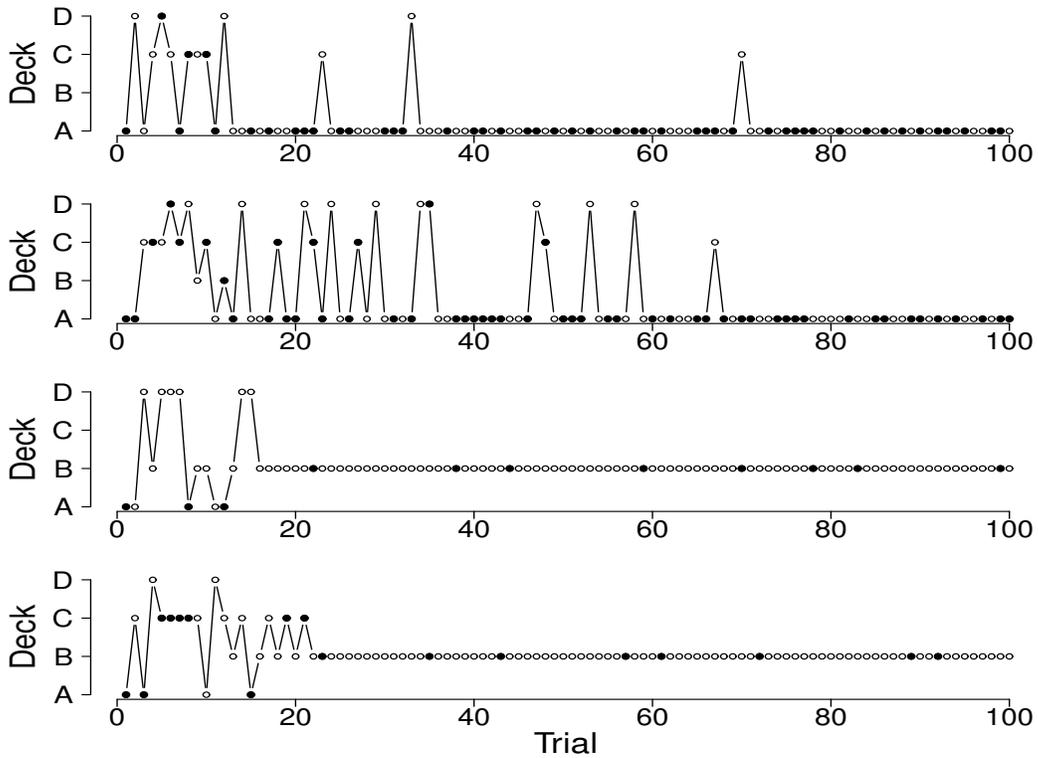


Figure 4.6: Deck selection profiles of four healthy participants showing a pronounced preference for the decks with infrequent losses.

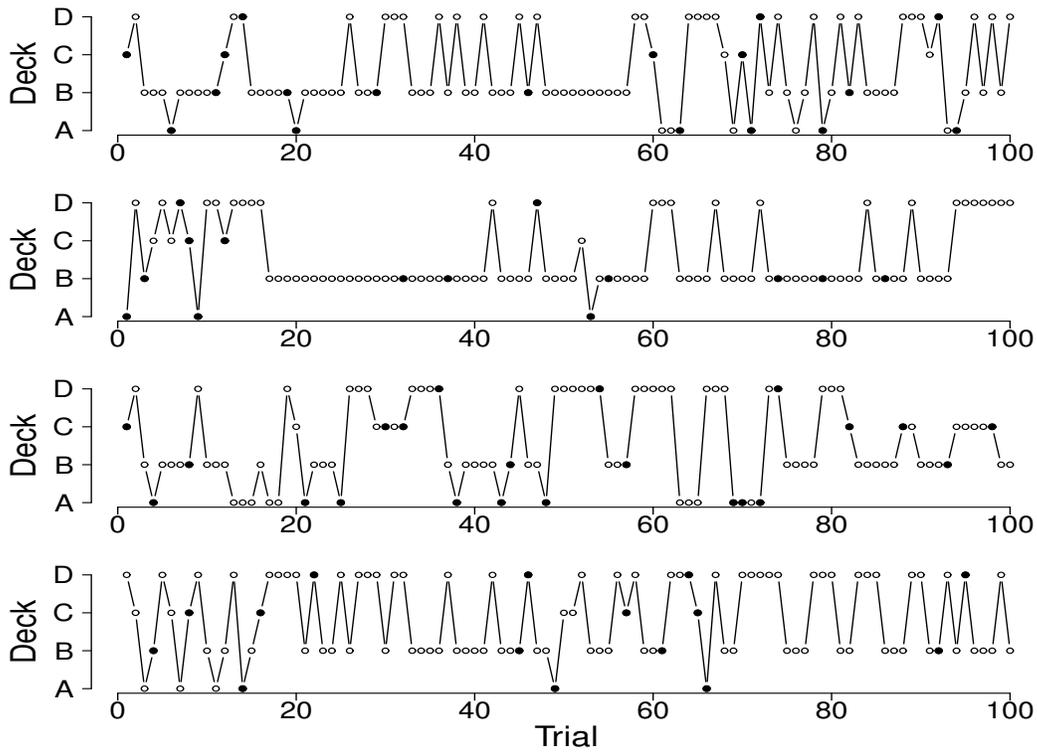
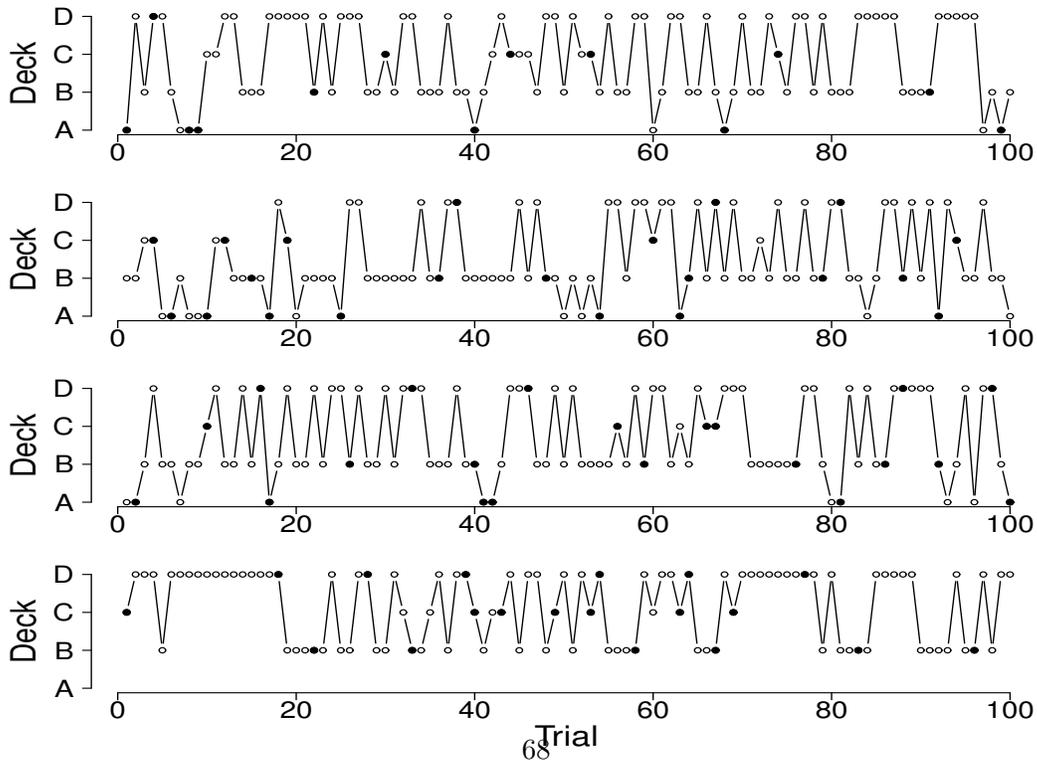


Figure 4.7: Deck selection profiles of four synthetic participants showing a pronounced preference for the decks with infrequent losses (generated by the PVL-Delta model; $A = 0.00, w = 0.00, a = 0.42, c = 1.19$).



Wetzels, & Wagenmakers, 2013a), the PVL-Delta model again performs better because the EV and PVL model’s failure to generate large proportions of switches, given a pronounced choice pattern, is even stronger: Given a pronounced choice pattern, the EV and PVL models fail to generate switch proportions higher than .35 and .46, respectively. Despite these discrepancies between the empirical and the generated switch behavior, we showed that –given a pronounced preference for the good decks or the decks with infrequent losses and those parameter combinations that yielded the maximum number of switches during the last 50 trials– the PVL-Delta model can produce choice patterns that strongly resemble the empirical choice patterns of healthy participants.

4.3 Test of Selective Influence

In this section we investigate whether the parameters of the PVL-Delta model indeed correspond to distinct psychological processes. We will therefore carry out a test of selective influence for the PVL-Delta model. This means that we fit the model to data collected from the standard IGT, but also from conditions that were designed to affect selectively one of the model parameters. These data were collected by Wetzels, Vandekerckhove, et al. (2010), and their experiment was originally designed as a test of selective influence for the EV model. However, the experimental manipulations that were intended to affect the parameters of the EV model should also be reflected by the parameters of the PVL-Delta model because of the high similarity between the two models.

Method

We fit the PVL-Delta model separately to four data sets reported by Wetzels, Vandekerckhove, et al. (2010). Specifically, Wetzels, Vandekerckhove, et al. (2010) conducted an experiment with a standard condition and three additional conditions that were designed to affect selectively one of the model parameters:⁷ In the “standard condition”, 19 participants completed a 150-trial IGT under the standard administration. In the “rewards condition”, 20 participants completed a 150-trial IGT under the instruction to pay more attention to rewards and to consider losses as less important. We expected this manipulation to decrease the loss aversion parameter w .

In the “updating condition”, 19 participants completed a 150-trial IGT under the standard administration. However, each choice was followed by a on-screen presentation of five numbers that the participants had to remember because, after the next choice, participants were asked about the relative position of one of the numbers. We expected this manipulation to increase the updating parameter a .

In the “consistency condition”, 16 participants completed a 150-trial IGT under the standard administration. However, they were told that after every 10 trials the payoff schemes for the decks could have changed. We expected this manipulation to decrease the consistency parameter c .

To fit the PVL-Delta model, we used a Bayesian hierarchical approach detailed in the next section. This estimation procedure has been consistently shown to outperform alternatives such as maximum likelihood estimation and Bayesian individual estimation (Ahn et al., 2011; Wetzels, Vandekerckhove, et al., 2010).

To assess whether the chains of all parameters had converged successfully from their starting values to their stationary distributions, we visually inspected the Hamiltonian Monte Carlo (HMC) chains and used the \hat{R} statistic (Gelman & Rubin, 1992). The \hat{R} statistic is a formal diagnostic

⁷Note that we use the data sets that Wetzels, Vandekerckhove, et al. (2010) obtained after having eliminated two sources of contamination. Specifically, Wetzels, Vandekerckhove, et al. (2010) removed participants for whom one or more of the maximum likelihood point estimates were located on the boundary of the parameter space, and participants for whom the Bernoulli baseline model outperformed the EV model.

measure of convergence that compares the between-chain variability to the within-chain variability. Values close to 1.0 indicate convergence to the stationary distribution, whereas values greater than 1.1 indicate inadequate convergence.

To assess model performance in absolute terms, we used two different methods: the post hoc absolute fit method and the simulation method (see also Steingroever et al., 2014). These two methods allow us to assess the model’s ability to fit and generate the choice patterns present in each of the four conditions. Our implementation of both methods relies on visually contrasting—separately for each deck as a function of 15 bins each containing 10 trials—the observed mean choice proportions from the experiment against the mean choice probabilities from the model.

Both methods start by sampling parameter values from the joint posterior distributions over the individual-level parameters (hereafter individual-level joint posteriors). In the case of the post hoc absolute fit method, the model is provided with the sampled parameter values, but also with the actual choices and payoffs of each participant. The post hoc absolute fit method computes the probability of choosing each deck on the next trial based on the information on the observed choices and payoffs up to and including the current trial. The simulation method, on the other hand, is only provided with the sampled parameter values, and relies on generating choices for another sequence of payoffs that could have been observed.⁸ In particular, on each trial, the simulation method generates a choice based on the predicted choice probabilities. For both methods and for each participant, we repeated the process of obtaining the predicted choice probabilities 100 times to account for uncertainty in the individual-level joint posteriors (for detailed recipes see Steingroever et al., 2014).⁹

To investigate the effect of the experimental manipulations, we visually compared the posterior distributions of the group-level parameters of all four conditions.

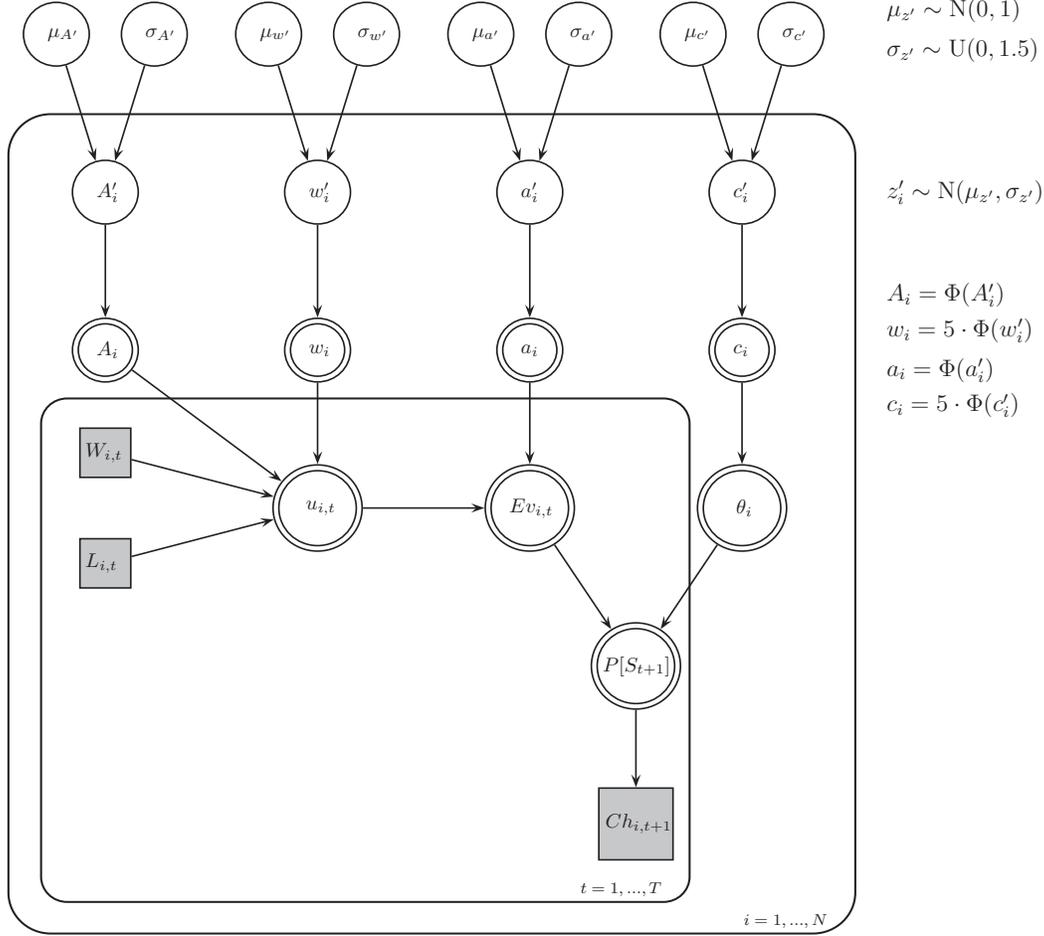
Bayesian hierarchical estimation procedure

To fit the PVL-Delta model to the data of the four experimental conditions, we used a Bayesian hierarchical estimation procedure (see Wetzels, Vandekerckhove, et al. (2010) for the same model specification in the case of the EV model). The Bayesian graphical PVL-Delta model for a hierarchical analysis is shown in Figure 4.8. This figure shows that the graphical model consists of two plates: The inner plate expresses the replications of the choices on $t = 1, \dots, T$ trials of the IGT, and the outer plate expresses the replications for $i = 1, \dots, N$ participants. For the sake of clarity, we omitted the notation that indexes the deck number k . The quantities $W_{i,t}$ (rewards of participant i on trial t), $L_{i,t}$ (losses of participant i on trial t), and $Ch_{i,t+1}$ (choice of participant i on trial $t + 1$) can directly be obtained from the data. The quantities $u_{i,t}$, $Ev_{i,t+1}$, and θ_i are deterministic because they can be calculated from Equations 4.1, 4.2, and 4.4. All individual-level parameters z_i , that is, $\{A_i, w_i, a_i, c_i\}$, are also deterministic because instead of modeling the individual-level parameters directly, we modeled their respective probit transformations z'_i , that is, $\{A'_i, w'_i, a'_i, c'_i\}$. This means that the parameters z'_i lie on the probit scale covering the entire real line. The probit transformation is the inverse of the cumulative standard normal distribution function. The parameters z'_i are assumed to be drawn from group-level normal distributions with mean $\mu_{z'}$ and standard deviation $\sigma_{z'}$. Only after the analysis was complete, we transformed the parameters $\mu_{z'}$ and z'_i back to the original scale.

⁸Note that we used the same payoff schedule as in the corresponding experiment.

⁹For completeness, we also produced predicted choices based on the joint posterior of the group-level parameters (hereafter group-level joint posterior); that is, we generated data with 1,000 parameter values that were randomly drawn from the group-level joint posterior. There are slight differences between the two types of posterior predictives, but the general conclusions are the same (see Appendix C for further details).

Figure 4.8: Bayesian graphical PVL-Delta model for a hierarchical analysis. $\Phi()$ is the cumulative standard normal distribution function.



The model specification requires a definition of priors for the group-level means and standard deviations. We assigned a normal prior to the group-level means, $\mu_{z'} \sim N(0, 1)$, and a uniform prior to the group-level standard deviations, $\sigma_{z'} \sim U(0, 1.5)$.

We implemented the PVL-Delta model in Stan (Stan Development Team, 2014b, 2014c; Hoffman & Gelman, 2014). The code to fit the PVL-Delta model in Stan is available on <http://www.helensteingroever.com>. To confirm that we correctly implemented the PVL-Delta model, we ran several parameter-recovery studies. The results of two such studies are presented in Appendix C.

For each parameter, we ran three HMC chains simultaneously. The fitting procedure consisted of two steps: First, we initialized all chains with randomly generated starting values. We collected 1,000 samples of each chain after having discarded the first 9,000 samples of each chain as burn-in. However, this procedure did not result in successful convergence of the HMC chains of all parameters: for instance, for some parameters, two chains may appear to have converged to their stationary distributions and looked like “hairy caterpillars” that are randomly intermixed, whereas the third chain behaved differently and producing an inferior goodness of fit (GOF). Therefore, in

a second step, we again ran three HCM chains for each parameter, but this time, we initialized all chains with parameter values close to the mean of the HCM chain that produced the best GOF in the first step. However, even this procedure resulted in convergence problems for a few participants (e.g., bimodal posterior distributions). We therefore excluded participants with such convergence issues and repeated the first and second step. This explains why the sample sizes presented in Table 4.5 are slightly smaller than those reported by Wetzels, Vandekerckhove, et al. (2010).

Table 4.5 also presents, for each data set separately, the number of burn-in samples and posterior samples that we collected for each chain. These specifications differ across data sets to ensure that all chains reached convergence. We based our inferences on these posterior samples.

Table 4.5: *Sample size of the four data sets and number of burn-in samples and posterior samples that we collected for each chain.*

Experimental condition	Sample size	Burn-in samples	Posterior samples
Standard	17	37,000	3,000
Rewards	19	30,000	3,000
Update	16	23,000	1,500
Consistency	15	18,000	2,000

Results

In this section, we discuss the results of the test of selective influence. We first focus on the behavioral level by describing the choice patterns observed in the four experimental conditions. Second, we focus on the level of the cognitive modeling analyses; we describe tests confirming that the posterior distributions converged successfully from their starting values to their stationary distributions. In addition, we show that the PVL-Delta model results in a satisfactory fit performance and simulation performance for the four conditions. Finally, we visually compare the posterior distributions of the group-level parameters of all four conditions to draw inferences about the effect of the experimental manipulations.

Behavioral data

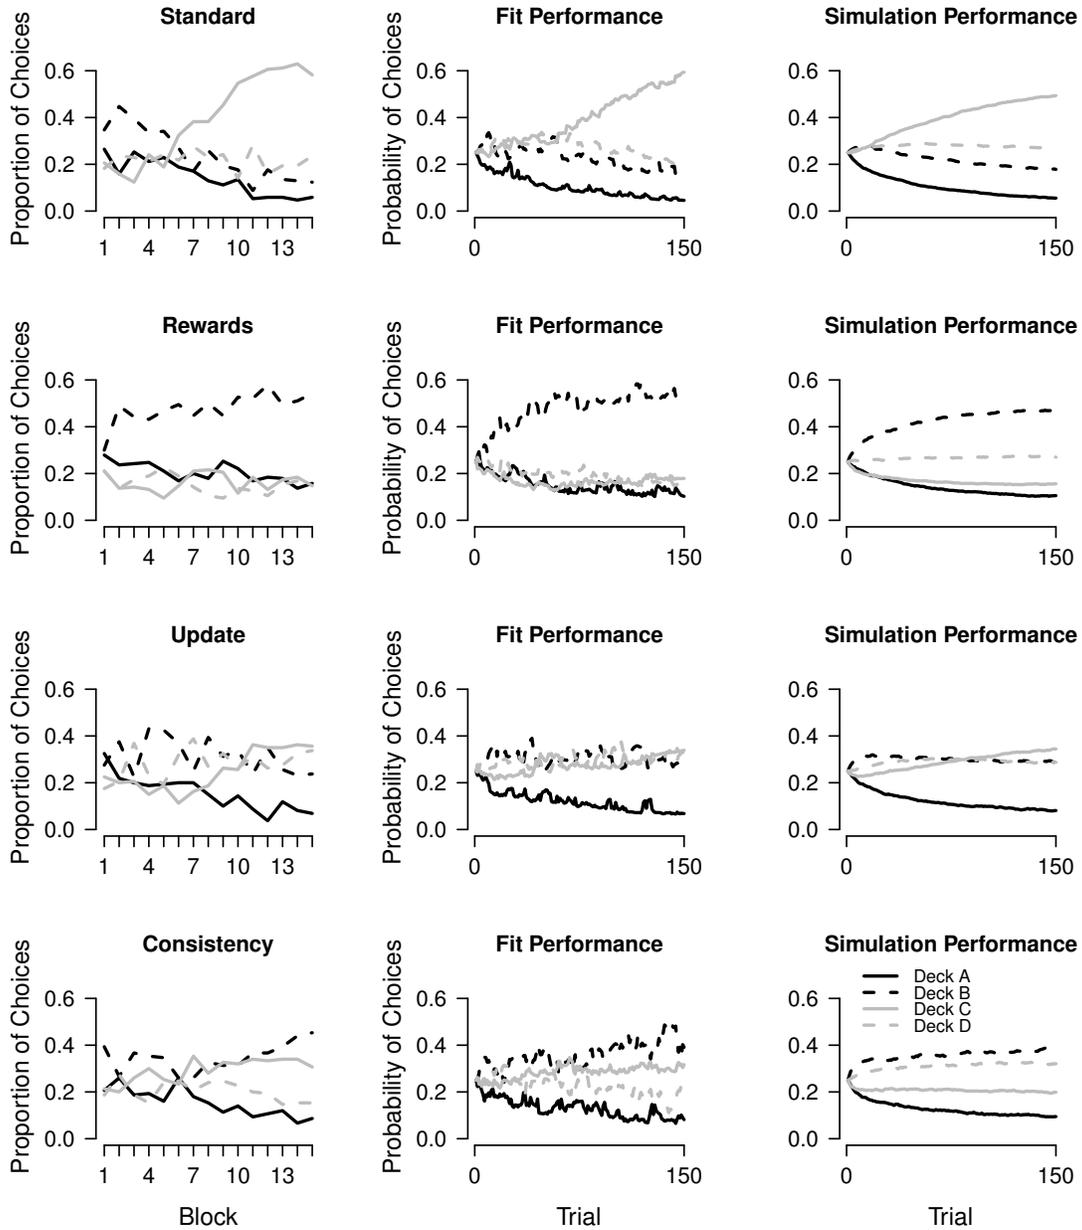
The mean proportion of choices from each deck within 15 blocks each containing 10 trials as observed in the four experimental conditions reported by Wetzels, Vandekerckhove, et al. (2010) are presented in the first column of Figure 4.9. In the standard condition, participants learned to prefer good deck C over all remaining decks; however, participants failed to learn that deck D is also a good deck.

In the rewards condition (i.e., participants were instructed to pay more attention to rewards and to consider losses as less important), participants learned to prefer bad deck B over all remaining decks. Note that even though bad decks A and B both yield high immediate rewards on every trial, participants did not learn to select deck A more often than good decks C and D. This may suggest that the experimental manipulation was only partly successful.

In the updating condition (i.e., each choice was followed by a on-screen presentation of five numbers that participants had to remember because, after the next choice, they were asked about the relative position of one of the numbers), participants show a very weak learning curve; they only learned to avoid deck A.

In the consistency condition (i.e., participants were told that after every 10 trials the payoff schemes for the decks could have changed), participants –in contrast to the intention of the

Figure 4.9: Observed choice behavior and assessment of absolute model performance. The first column shows the mean proportion of choices from each deck within 15 blocks as observed in the four experimental conditions reported by Wetzels, Vandekerckhove, et al. (2010). Each block contains 10 trials. The second and third column show the fit performance and simulation performance, respectively, for each of the four conditions. Fit performance and simulation performance are based on random draws from the individual-level joint posteriors.



experimental manipulation— did not evenly explore all decks across the entire 100 trials. Instead participants learned to prefer decks B and C over the remaining decks. It seems that participants prefer bad deck B because it yields high immediate rewards on the majority of the trials; however, participants prefer good deck C because it never yields a net loss and is therefore a safe option.

Convergence checks

Visual inspection of the HMC chains and consideration of the \hat{R} statistics for all parameters (all parameters had \hat{R} values below 1.045) suggest that all chains have converged successfully. To illustrate how we visually assessed convergence, we show the chains of one individual-level parameter in Appendix C. From the figure it is evident that the chains have converged successfully from their starting values to their stationary distribution, looking like “hairy caterpillars” that are randomly intermixed.

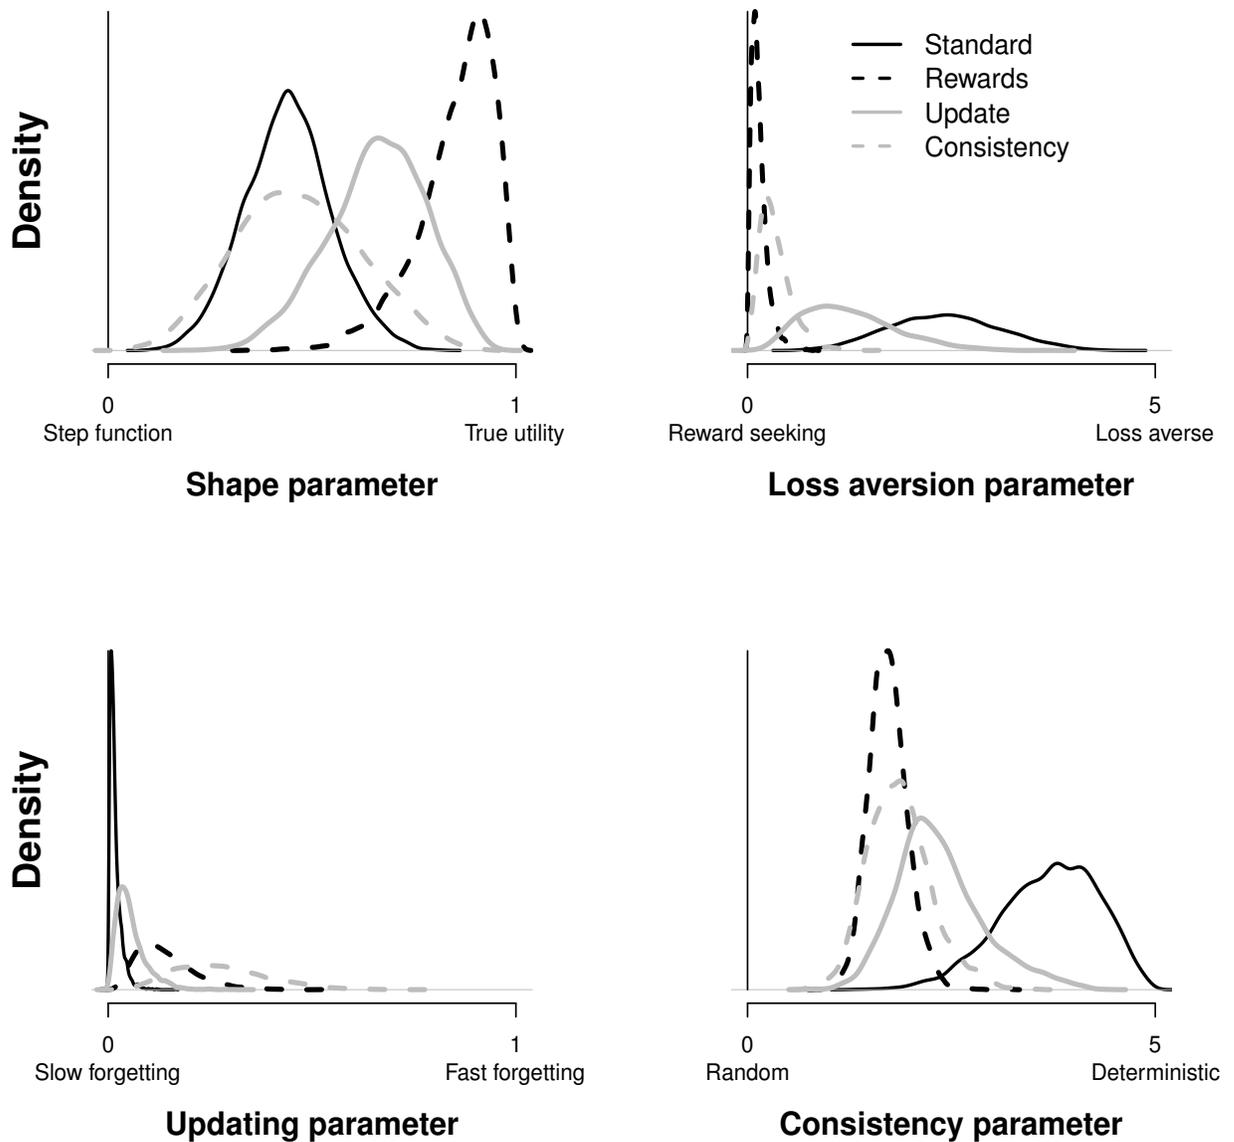
Absolute model performance

To assess the absolute model performance of the PVL-Delta model with respect to the four experimental conditions, the second and third column of Figure 4.9 show the fit performance and simulation performance, respectively. Fit performance and simulation performance are based on random draws from the individual-level joint posterior. From the second column of the figure it is evident that the PVL-Delta model provides a good fit to the data of all four conditions (i.e., the model makes accurate one-step-ahead predictions when provided with access to the observed sequence of choices and payoffs). In addition, the third column of Figure 4.9 illustrates that the PVL-Delta model adequately generates the choice pattern shown by the standard and update conditions. In the case of the rewards and consistency conditions, the simulation performance of the PVL-Delta model is acceptable; the model correctly predicts the most preferred deck, but fails to account for the rank order of the remaining three decks: In the reward condition, the model predicts that deck D is preferred over decks A and C even though the participants chose these three decks about equally often. In the consistency condition, the model predicts that deck D is preferred over deck C even though the participants showed the reverse pattern. To sum up, the PVL-Delta model captures the global patterns in the data providing an acceptable fit and simulation performance with respect to the four data sets at hand; this allows us to meaningfully compare the group-level parameters of the four conditions.

Test of selective influence

Figure 4.10 presents the posterior distributions of the group-level parameters of all four conditions. It is evident that the experimental manipulation is successfully reflected in the loss aversion parameter and the consistency parameter: First, compared to participants that received the standard instruction, participants who were instructed to focus on rewards (i.e., the rewards condition) had lower values for the loss aversion parameter indicating that they were indeed more reward-seeking. Second, fitting the PVL-Delta model to data of participants that were told that after every 10 trials the payoff schemes for the decks could have changed (i.e., the consistency condition) resulted in a smaller consistency parameter (i.e., a more random choice behavior) than fitting the PVL-Delta model to data of participants that received the standard instructions. However, in the update condition is no clear effect on the updating parameter. Yet, it is evident that the consistency parameter in the update condition is noticeably lower than in the standard condition (i.e., a more random choice behavior); this is consistent with the choice pattern

Figure 4.10: Posterior distributions for the group-level parameters of the PVL-Delta model in the four experimental conditions.



shown by the update condition; participants only learned to avoid deck A, but show a completely indistinguishable preference for the remaining three decks.

4.4 Discussion

In this article, we conducted two tests to validate the PVL-Delta model: a parameter space partitioning study and a test of selective influence. Applying PSP to the PVL-Delta model, we have obtained a deeper understanding of the model’s behavior. We used two different definitions of choice patterns; the broad definition allowed us to get an indication of how central each of the choice patterns are to the model’s overall performance, and the restricted definition allowed us to assess the model’s data-fitting potential when confronted with data featuring pronounced deck preferences.

Using the broad definition of choice patterns, the PSP study revealed that the PVL-Delta model can generate all typical empirical choice patterns. However, the PVL-Delta model generates the choice pattern featuring a preference for the bad decks only over a minor part of its parameter space suggesting that this choice pattern is virtually irrelevant to the model’s overall performance.

Using the restricted definition of choice patterns, the PSP study revealed that the PVL-Delta model can still generate all pronounced empirical choice patterns over a minor part of its parameter space. But for these pronounced choice patterns, the PVL-Delta model generally underestimates the empirical switch proportions during the last 50 trials. In particular, given pronounced preferences for the bad decks, the PVL-Delta model fails to account for the empirical switch behavior. This failure seems to be caused by the Prospect Utility function of the PVL-Delta model: In a previous PSP study, (Steingroever, Wetzels, & Wagenmakers, 2013a) showed that this failure is also present in the PVL and EV-PU model (i.e., models with the Prospect Utility function), but not in the EV model (i.e., a model without the Prospect Utility function). However, in the case of the other two pronounced choice patterns –the choice patterns favoring decks with high expected value or low loss frequency– we showed that the PVL-Delta model provides a good account for the empirical switch behavior for some parameter combinations.

The results of the PSP study for the PVL-Delta model and the earlier PSP studies for the EV and PVL models (Steingroever, Wetzels, & Wagenmakers, 2013a) suggest that the PVL-Delta model outperforms its two competitors. The EV model fails to generate a pronounced preference for the decks with infrequent losses; the PVL model is able to generate pronounced decks preferences, but underestimates the switch proportions even more strongly than the PVL-Delta model. This superiority of the PVL-Delta model is in line with the posterior predictive checks reported by Steingroever et al. (2014).

An important advantage of PSP is that it is a global analysis technique augmenting local methods that have previously been used to compare RL models (Pitt et al., 2006, 2008). Whereas local methods, such as the post hoc fit criterion or the generalization criterion, evaluate a model’s performance at a single point of a model’s parameter space, global methods such as PSP help us to determine the full range of choice patterns that a model can generate by varying its parameter values (see also Vanpaemel, 2009). This means that we can obtain a global perspective on the data-fitting potential of the PVL-Delta model. Thus, if researches wish to apply the PVL-Delta model to IGT data, they can decide based on the behavioral results whether it is appropriate to apply the PVL-Delta model or not.

The PSP results of this paper should be interpreted with care. PSP gives an indication of how central choice patterns are to the overall performance of the model. However, it is premature to conclude that the PVL-Delta model cannot generate the choice pattern “bad decks over goods

decks” at all, solely because the model generates this choice pattern over a small part of the parameter space. Instead, we can only conclude that this choice pattern is not central to the model’s overall performance.

It should also be noted that the inferences drawn from the PSP study strongly depend on our definitions of choice patterns. The restricted definition of choice patterns was based on IGT performance of healthy participants (Steingroever, Wetzels, & Wagenmakers, 2013a). We could thus detect inconsistencies between the empirical popularity of each pronounced choice pattern in the Steingroever, Wetzels, Horstmann, et al. (2013) data pool and the frequency predicted by the PVL-Delta model. It is troubling that the PVL-Delta model fails to generate a pronounced preference for the bad decks with many switches. But it should be acknowledged that this choice pattern is not central in healthy participants’ IGT performance: in the Steingroever, Wetzels, Horstmann, et al. (2013) data pool, only 5% ($N = 18$) of the healthy participants showed this choice pattern (Table 4.2). Still, this choice pattern is assumed to be characteristic for patients with decision-making deficits (Bechara et al., 1994, 1997), but a better empirical foundation (e.g., a literature review on the IGT performance of clinical groups) is required to accurately judge the gravity of the PVL-Delta model’s failure to generate a pronounced preference for the bad decks with many switches.

The test of selective influence revealed that the experimental manipulations had a noticeable effect on the loss aversion parameter and consistency parameter, but not on the updating parameter. However, it is premature to conclude that the updating parameter does not correspond to memory processes. It may be that the experimental manipulation did not work out properly. In addition, one should bear in mind that every data set is characterized by its own idiosyncrasies. IGT data generally are highly idiosyncratic—possibly because the IGT is a very complex task (Steingroever, Wetzels, Horstmann, et al., 2013). In order to be able to draw more accurate conclusions on whether the parameters represent distinct psychological processes, independent repetitions of the test of selective influence and even different experimental manipulations are necessary.

The results of this article confirm that the PVL-Delta model is an attractive alternative to the popular EV and PVL models. However, the PVL-Delta model is also characterized by a few shortcomings because it underrepresents the choice pattern featuring a preference for the bad decks. Nevertheless, we recommend that researchers use the PVL-Delta model to disentangle psychological processes underlying IGT performance, provided that they rigorously assess absolute model performance before interpreting the model parameters (Steingroever et al., 2014).

Acknowledgement

This publication was supported by the Dutch national program COMMIT, by an Open Access grant from NWO, and by the Netherlands Organisation for Scientific Research (NWO; 404-10-086).