Article:
Weiss-Cohen, L, Konstantinidis, E, Speekenbrink, M et al. (2018) Task complexity moderates the influence of descriptions in decisions from experience. Cognition, 170. pp. 209-227. ISSN: 0010-0277

https://doi.org/10.1016/j.cognition.2017.10.005

Task complexity moderates the influence of descriptions in decisions from experience

Leonardo Weiss-Cohen
University College London

Emmanouil Konstantinidis
University of Leeds and University of New South Wales

Maarten Speekenbrink and Nigel Harvey
University College London

Abstract

Decisions-makers often have access to a combination of descriptive and experiential information, but limited research so far has explored decisions made using both. Three experiments explore the relationship between task complexity and the influence of descriptions. We show that in simple experience-based decision-making tasks, providing congruent descriptions has little influence on task performance in comparison to experience alone without descriptions, since learning via experience is relatively easy. In more complex tasks, which are slower and more demanding to learn experientially, descriptions have stronger influence and help participants identify their preferred choices. However, when the task gets too complex to be concisely described, the influence of descriptions is reduced hence showing a non-monotonic pattern of influence of descriptions according to task complexity. We also propose a cognitive model that incorporates descriptive information into the traditional reinforcement learning framework, with the impact of descriptions moderated by task complexity. This model fits the observed behavior better than previous models and replicates the observed non-monotonic relationship between impact of descriptions and task complexity. This research has implications for the development of effective warning labels that rely on simple descriptive information to trigger safer behavior in complex environments.

Keywords: decisions from experience, decisions from description, description-experience gap, reinforcement learning, complexity, Iowa Gambling Task

Decisions in everyday life are often made using a combination of descriptive and experiential information. For example, consumers use descriptive reviews and personal experiences of similar items bought in the past; doctors rely on written published literature and their own clinical experience; and drivers pass road signs warning them of traffic queues on a familiar stretch of road. The ongoing proliferation of warning signs and labels can be considered as descriptive information that is added to an individual's own experience, reminders of high-loss small-frequency risks that are rarely experienced. For example, passengers frequently run at stations in order to catch their trains, and the overwhelming majority never directly experiences any accidents. But warnings signs are common, reminding individuals that running can be dangerous and cause harm. Despite the ubiquitous presence of both sources of information concurrently, the vast majority of decision-making research has exposed participants either to "decisions from description" or "decisions from experience" separately, very rarely combining the two in the same task (Fantino & Navarro, 2012).

Decisions from description vs. experience

Decisions from description are those in which a complete, idealized, and abstract set of information about the values and frequencies of potential outcomes from each choice is provided, typically in writing, to participants before choices are made (e.g., "50% chance to win 1,000; 50% chance to win nothing", from Kahneman & Tversky, 1979, p. 264). Decisions from experience, in contrast, do not provide any information before choices are made and, instead, require participants to form their own view of the potential outcomes from each choice via feedback provided after each selection is made (e.g., "You have won 100 dollars", from Bechara, Damasio, Damasio, & Lee, 1999, p. 5474). For the vast majority of the history of decision-making research, these two paradigms have been explored separately, each in their own individual domain.

One of the earliest attempts to empirically and systematically compare the two paradigms and study any differences in behavior was made by Barron and Erev (2003). They found that in decisions from experience, participants underweighted rare events, and

Leonardo Weiss-Cohen, Maarten Speekenbrink, and Nigel Harvey, Department of Experimental Psychology, Division of Psychology and Language Sciences, University College London, London, UK; Emmanouil Konstantinidis, Centre for Decision Research, University of Leeds, Leeds, UK, and School of Psychology, University of New South Wales, Sydney, Australia.

This research has been supported by a UCL Impact Award.

Correspondence regarding this manuscript should be addressed to Leonardo Weiss-Cohen, Department of Experimental Psychology, Division of Psychology and Language Sciences, University College London, Gower Street, London WC1E 6BT, UK. E-mail: leoweisscohen@gmail.com

were more risk seeking in the gain than in the loss domain. These are the reverse of well-established phenomena in the decision-making literature (Kahneman & Tversky, 1979), that is, overweighting of rare events and more risk seeking behavior in the loss domain, which had previously been explored exclusively in paradigms based in decisions from descriptions. This difference in behavior between decisions from descriptions and decisions from experience was later named the "description-experience gap" by Hertwig and Erev (2009). A substantial body of research has since been dedicated to studying the gap, by presenting different participants with the same choice scenarios, with either descriptions alone or experience alone (for a recent review, see Camilleri & Newell, 2013b). Camilleri and Newell suggested that the gap might be caused by differences in how information is presented, cognitively processed, stored, internally represented and compared. While the field has now extensively studied and compared the two paradigms side-by-side, limited research has been dedicated to tasks in which the two sources are combined and available simultaneously.

The limited previous research combining description and experience, "decisions from description-plus-experience", has shown no difference in behavior when adding descriptions to decisions from experience, if the two provide the same underlying information (Lejarraga & Gonzalez, 2011; Weiss-Cohen, Konstantinidis, Speekenbrink, & Harvey, 2016). In other words, behavior was similar in decisions from experience and decisions from description-plus-experience[1]. Lejarraga and Gonzalez proposed that this lack of observable differences in behavior was due to descriptions being ignored when experience was also available. However, descriptions do not appear to be fully ignored, because they influence behavior when they provide novel information (Barron, Leider, & Stack, 2008). We showed in our previous research (Weiss-Cohen et al., 2016), using cognitive modeling, that descriptions are not completely ignored, but instead they are discounted. A similar effect of discounting of descriptions, when experience was also present, was found in probability judgments by Shlomi (2014). Empirical research so far has shown that descriptions are discounted, to the point of apparent neglect, when combined with experience provided in the form of feedback.

The concept that feedback overwhelms descriptive information had been proposed

[1] The reason for not comparing them to decisions from description (without experience) is because the two paradigms are inherently different. Decisions from description are typically single-choice, single-outcome, while decisions from experience are multiple-choice, multiple-outcomes (Camilleri & Newell, 2013a). By adding descriptions to the latter, it is possible to keep their repeated-choice nature constant.

before (Jessup, Bishara, & Busemeyer, 2008; Yuviler-Gavish & Gopher, 2011). Lejarraga (2010) showed empirically that individuals prefer experiences over descriptions, by allowing participants to choose between the two types of information, which was then used to learn the probabilities associated with the options available and make their decisions. This preference for experiences can be explained by different cognitive processes being applied to descriptions and experiences, as suggested by Glöckner, Fiedler, Hochman, Ayal, and Hilbig (2012): when dealing with descriptions, individuals might engage in more complex computational processes, calculating the expected value of each option; conversely, personal experiences use simpler, more instinctive, and less demanding integration processes.

Research in other fields has shown similar findings arising from the additional cognitive burden consequent on processing descriptions and a resulting preference for experiences. For example, Gigerenzer and Hoffrage (1995) suggested that individuals are better at keeping track of sequentially acquired information, such as naturally presented frequencies experienced over time, and worse at processing percentages and probabilities presented descriptively. According to Hasher and Zacks (1984), individuals are able to learn from experience incidentally, automatically encoding frequencies with minimal effort and attention. On the other hand, Erev, Ert, Plonsky, Cohen, and Cohen (2017) have shown, using computational models, that decisions from descriptions can be explained by individuals mentally simulating outcomes from descriptions to arrive at expected values, a time consuming and costly process. Decision Field Theory, a model proposed by Busemeyer and Townsend (1993), is based around a similar concept that individuals mentally sample information over time until a decision threshold is reached. Overall, descriptions appear to be more cognitively demanding, whereas humans (and all other animals) are more naturally adapted to encode and process experiences.

The proposition that descriptions are more costly and effortful to process than experiences seems to support the evidence observed so far that descriptions are typically discounted in description-plus-experience paradigms, and that individuals prefer to rely on experiences rather than descriptions. However, we believe that the strength of this preference may not necessarily be static and that it is calibrated according to the situation. Some factors, such as plausibility and description complexity, have already been shown empirically to influence the strength of this preference. Less plausible descriptions, in comparison to the actual experienced feedback, received lower weights than more plausible ones (Weiss-Cohen et al., 2016). Lejarraga (2010) and Lejarraga and Gonzalez (2011) explored situations where descriptions were made less attractive to participants by increasing their perceived complexity, therefore making them even harder to process cognitively while keeping the underlying experiential task unchanged. By increasing the cognitive cost of processing descriptions, the authors showed an increase in preference for experiences. One limitation of these previous studies, however, was that the researchers did not change the complexity of the task itself, only the complexity of the descriptions used to label the same underlying processes by using simpler or more complex notation.

Task complexity

While complexity can be a subjective construct, and significantly dependent on individual differences, it is also related to certain underlying task characteristics that can be defined objectively (Campbell, 1988; Wood, 1986). Halford, Wilson, and Phillips (1998) have defined complexity as “the number of related dimensions or sources of variation” (p.803), in terms of cognitive and computational processing loads and its influence on learning difficulty. The complexity of patterns of data can thus be quantified in relation to the ease of learning the simplest set of rules, with the minimum number of dimensions (or the most compressed set of information), which is required to represent all of the data’s potential sources of variability (Mathy & Feldman, 2012). More complex rules are the ones that require more information, are not as compressible, and therefore harder to learn (Feldman, 2000). For categorization tasks, for example, complexity increases, and learning deteriorates, in proportion to the minimum number of dimensions or components needed to identify items (Briscoe & Feldman, 2011; Mathy & Bradmetz, 2004). Comparably, memory tasks can be made more difficult by increasing the number of items that individuals are asked to recall (Miller, 1956), although if some of those items can be compressed together into fewer chunks of information, then empirical performance improves, and complexity is deemed to be lower (Cowan, 2001).

In the decision-making domain, Thorngate (1980) and Johnson and Payne (1985) defined task complexity in relation to the number of different alternatives from which participants can select, and the number of possible outcomes available from each alternative. Increasing the number of alternatives and outcomes increases the entropy of the task, which can be associated with higher task complexity (Fasolo, Hertwig, Huber, & Ludwig, 2009). Entropy is an objective measure that has been used to quantify task complexity, based on information theory, with higher entropy associated with higher complexity (Swait & Adamowicz, 2001). Despite some research dedicated to studying decision making with multiple alternatives and multiple outcomes (Ert & Erev, 2007; Hills, Noguchi, & Gibbert, 2013; Noguchi & Hills, 2016), most research uses relatively simple tasks, both in descriptive and experiential paradigms (Hertwig & Erev, 2009; Rakow & Newell, 2010). Building upon the simple experimental paradigms commonly used to study general decision making, the extant description-plus-experience research has also utilized relatively simple tasks, based on two options, each of which has one or two potential outcomes (the single-outcome option providing a guaranteed result when selecting that option). This canonical preference for simple tasks, with their associated low costs of learning from experience, might be the driver behind the limited influence of descriptions on description-plus-experience tasks, and might explain why participants have shown preference for experiences over descriptions so far.

In this article, we extend the research in the field of decisions from description-plus-experience into the domain of more complex tasks, increasing both the number of alternatives available, and the number of unique possible outcomes from each alternative. We believe that task complexity will moderate the influence that descriptions have on decisions from experience. In simple tasks, participants quickly learn to identify the structure of the environment experientially, and no additional information is needed or desired. Descriptions, if available, are not very useful and do not help participants. These results has been empirically observed before, with simple tasks (Lejarraga & Gonzalez, 2011; Weiss-Cohen et al., 2016). Increasing the complexity of the task should make learning experientially slower, more costly, and more cognitively demanding (Ashby, Konstantinidis, & Yechiam, 2017; Fasolo et al., 2009; Frey, Mata, & Hertwig, 2015). This should make descriptions relatively more attractive than before, as relying on descriptions should provide an advantage to participants by giving them additional information that reduces learning time by lowering the need to explore the environment experientially. Higher task complexity should lead to situations in which engaging the extra processing effort associated with descriptions becomes cost-efficient. Therefore, an increase in task complexity should lead to an increase in the influence of descriptions on behavior.

Furthermore, we do not expect this relationship between task complexity and influence of descriptions to be monotonic. Task complexity is also closely linked to information processing, with more complex tasks being defined as those in which there is increased information loads, diversity and rate of change (Campbell, 1988). Research on information load has uncovered a non-monotonic inverted U-shape relationship between amount of information available and decision accuracy (Eppler & Mengis, 2004; Hwang & Lin, 1999). It has been shown empirically that an increase in information aids the decision-making process initially, but only up to a certain point, after which any additional information is actually detrimental and reduces the quality of the decisions (Jacoby, Speller, & Berning, 1974; Jacoby, Speller, & Kohn, 1974). This is commonly called "information overload", and describes the negative effects of receiving too much information. Potential causes for information overload include complexity of information provided, number of items of information and number of alternatives, among others (see Eppler & Mengis, 2004, p. 332). Similar U-shaped patterns peaking for medium complexity tasks have been found across other dimensions, such as choice satisfaction (Reutskaja & Hogarth, 2009), purchasing intentions (Shah & Wolford, 2007), the ability to accurately assess values (Keller & Staelin, 1987), the extent of information processing (Paul & Nazareth, 2010), and overall effort allocation (Swait & Adamowicz, 2001).

Consequently, we expect the influence of descriptions on our experiments to reduce in very complex tasks, after peaking in medium complexity tasks. If the task becomes too complex, then the descriptions required to summarize the task also become overly complex. The excess of information available, both in experience and description, should lead to information overload. Descriptions might also become too unwieldy to process cognitively, reducing their attractiveness as a source of information. We expect to find the maximum influence of descriptions in tasks of medium complexity, at the point where performance starts to suffer but descriptions are still not too complex. It is at this point that descriptions should be able to provide the most assistance.

Overview of current research

In the first two experiments, we begin by investigating the effects of introducing descriptions in a relatively complex task which has been widely used to study ambiguous decisions from experience, the Iowa Gambling Task (IGT: Bechara, Damasio, Damasio, & Anderson, 1994). The IGT has a particularly complex payoff structure which we believe can be better exploited with the benefit provided by the presence of descriptions. In this case, we expect descriptions to influence behavior, speed up learning and lead participants to perform better on the task, by choosing the advantageous decks earlier and more frequently. We also expect participants to gather less information experientially, by exploring less when descriptions are available, as descriptions already provide additional information to participants. In our third experiment, we extend our research into the manipulation of task complexity itself, comparing the influence of descriptions in tasks with different levels of complexity. We explored two separate dimensions of complexity, the number of choices and the number of outcomes from each choice. We tuned Experiment 3 in order to create substantially simpler and more complex tasks in comparison to Experiments 1 and 2, to empirically test a wide spectrum of complexity. In addition to the experimental behavioral analysis, we also fit cognitive models to the human data. We present a description-plus-experience cognitive model that integrates descriptive information into a reinforcement learning framework, with a novel approach for weighing descriptions using entropy as a proxy for the complexity of the task.

Experiment 1

Method

Design. The first experiment was based on the original Iowa Gambling Task (Bechara et al., 1994), with the addition of descriptive information for half of the participants in a new experimental manipulation, creating a described IGT (DIGT). The experiment followed a two-way between-subjects design controlling for the presence or absence of descriptions: in the experience-only (E) condition, participants relied on experience alone (in the form of feedback after each trial) to learn about the options available to them, without any descriptions; and in the description-plus-experience (DE) condition, participants were shown a full description of the distribution of outcomes available for each option, in addition to the experiential feedback after each trial.

Participants. We recruited 100 participants on-line using Amazon's Mechanical Turk service (47 females; age: M=31.0 years, SD=10.5 years), half in each experimental condition. Participation was restricted to individuals whose location was defined as in the United States. No participants were excluded from the analysis. Participants were paid a fixed amount of US$ 0.25 for participating and an additional bonus according to the outcomes of the choices they made during the experiment, with an average bonus of US$ 0.55 (SD=US$ 0.28)[2].

Task. The task closely followed the original IGT (Bechara et al., 1994). The instructions closely matched the original wording apart from changes needed for the computerized on-line delivery of this version of the task (Chiu & Lin, 2007). Participants were presented with four decks of cards (decks A, B, C, and D), side by side on the screen, with the backs of the cards displayed and their faces hidden (Figure 1). The faces of the cards provided the feedback after each selection, with the number of points earned or lost associated with each individual card. The naming of the decks given here was used for analysis only and not shown to participants. The order of the decks from left to right was randomized for each participant, as well as the pattern on the back of each deck. Choices were made using the mouse. To avoid rapid sequential clicking of the same choice repeatedly, participants were required to move the mouse cursor to a button at the bottom of the screen between selections. Participants' choices were financially consequent and accumulated towards their final pay. Participants started the task with 2,000 points and points earned or lost after each selection were added to or deducted from their total. Points were converted to money at a rate of US$ 0.20/1000 points. Accumulated amounts in points and U.S. dollars were shown on-screen and updated after each choice was made.

The schedule of outcomes from each deck was the same as in the original IGT: The

[2]Participants in the DE condition received a significantly 66% higher bonus than participants in the E condition, according to an asymptotic Wilcoxon rank sum test (DE=US$0.68, E=US$0.41, W=1919, z=4.47, p<.001).

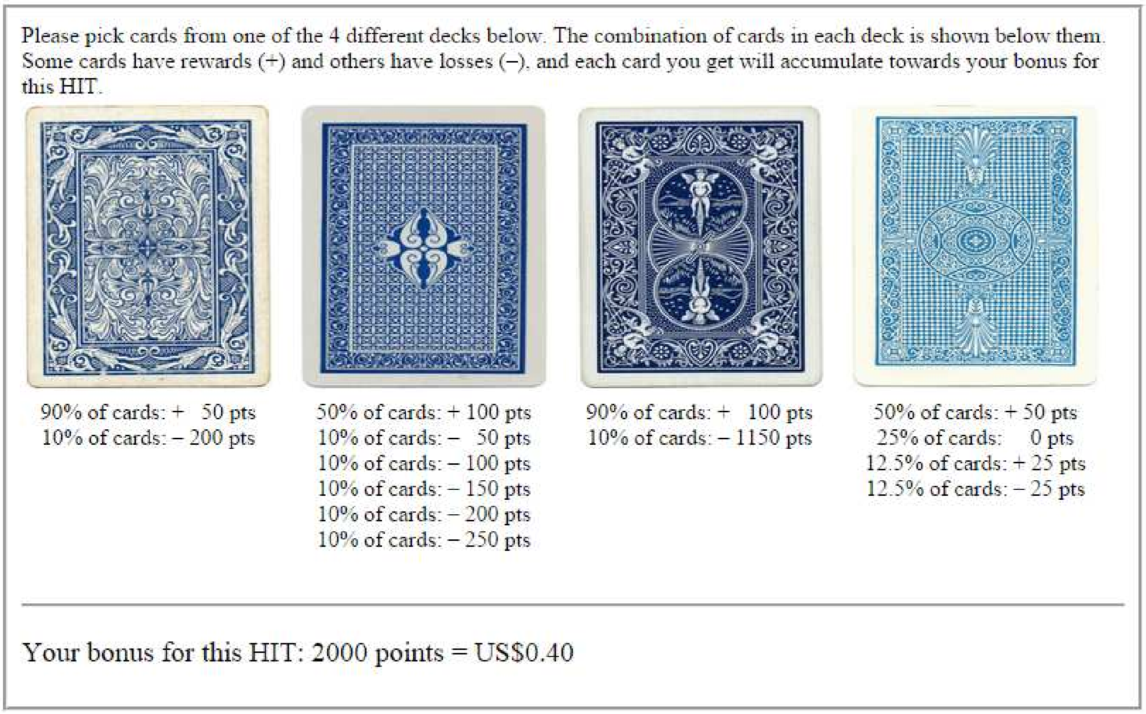

Figure 1. Screenshot of the first trial of Experiment 1 in the description-plus-experience (DE) condition, showing the DIGT (described IGT), with a random presentation of the decks in the order D, A, B, C from left to right. The patterns on the backs of each deck were also randomized. In the experience-only (E) condition, the space underneath the cards was left blank, with no descriptions shown, and the second sentence in the title regarding the combination of cards in each deck was replaced with "Each deck contains a different combination of cards". The "HIT" mentioned in the title refers to the terminology employed by Amazon's Mechanical Turk service to designate individual tasks on which participants can work, and stands for Human Intelligence Task.

order of the cards within each deck was not random but instead followed the fixed order given in the original task, with a repeating pre-defined sequential pattern of 40 cards for each deck (Bechara et al., 1994, Figure 1). In contrast to the original IGT, which showed rewards and losses separately for each card (e.g., "You have won 100 points, but you also have lost 150 points"), we opted to summarize outcomes as single net values (e.g., "–50 points"). This made the task simpler to describe, and circumvented the predictability of rewards associated with the original study (Steingroever, Wetzels, Horstmann, Neumann, & Wagenmakers, 2013). Decks A and B have a negative expected value of –25 points for each card, while decks C and D have a positive expected value of +25 points for each card. Hence decks A and B are referred to as the *disadvantageous* decks, and decks C and D are the *advantageous* decks. In order to maximize their bonus, participants have to select more often from the advantageous decks and avoid the disadvantageous ones.

In the experience-only (E) condition, participants did not receive any further information about the decks, and had to learn which decks were the advantageous ones via the feedback provided after each selection, in what was a close replication of the traditional IGT paradigm. In the description-plus-experience (DE) condition, a description of the cards contained in each deck (Table 1) was permanently displayed underneath the relevant deck, across all trials (Figure 1), in addition to the feedback provided, resulting in a described IGT (DIGT). After each selection, participants were shown only the outcome in points of the card from the deck they selected (i.e., partial feedback), and the card selected was replaced at the end of that deck, with no changes to the card order in the non-selected decks. Therefore participants learned only about the deck they selected on each trial, with no new feedback information for unselected decks. Participants were not told beforehand how many cards they would get to choose, and instead were instructed to choose cards from decks repeatedly until told to stop, which was after 100 choices. The task was self-paced and was completed on average in 8.30 minutes (*SD*=4.35).

Table 1
Actual card composition and wording of descriptions shown underneath each deck in Experiment 1. The expected value for each individual card in decks A and B was –25 points and in decks C and D was +25 points.

Experience-only condition (E), N=50			
Deck A	Deck B	Deck C	Deck D
(blank)	(blank)	(blank)	(blank)
Description-plus-Experience condition (DE), N=50			
Deck A	Deck B	Deck C	Deck D
50% of cards: +100 pts 10% of cards: –50 pts 10% of cards: –100 pts 10% of cards: –150 pts 10% of cards: –200 pts 10% of cards: –250 pts	90% of cards: +100 pts 10% of cards: –1150 pts	50% of cards: +50 pts 25% of cards: 0 pts 12.5% of cards: +25 pts 12.5% of cards: –25 pts	90% of cards: +50 pts 10% of cards: –200 pts

Results

Selections from advantageous decks. The main dependent variable was the frequency of cards selected from the advantageous decks, calculated as the total number of cards selected from decks C and D for each sequential block of 20 choices (Figure 2A). They were analyzed with a linear mixed-effects model using the lme4 package (Bates, Maechler, Bolker, & Walker, 2014) and post-hoc analyses using the lsmeans package (Lenth, 2016), with Tukey adjustments, in R (R Core Team, 2014). The fixed effects were the presence or absence of descriptions (DE or E), the blocks of 20 choices each (with polynomial contrasts), and their interaction. The model also contained a random intercept for each participant.

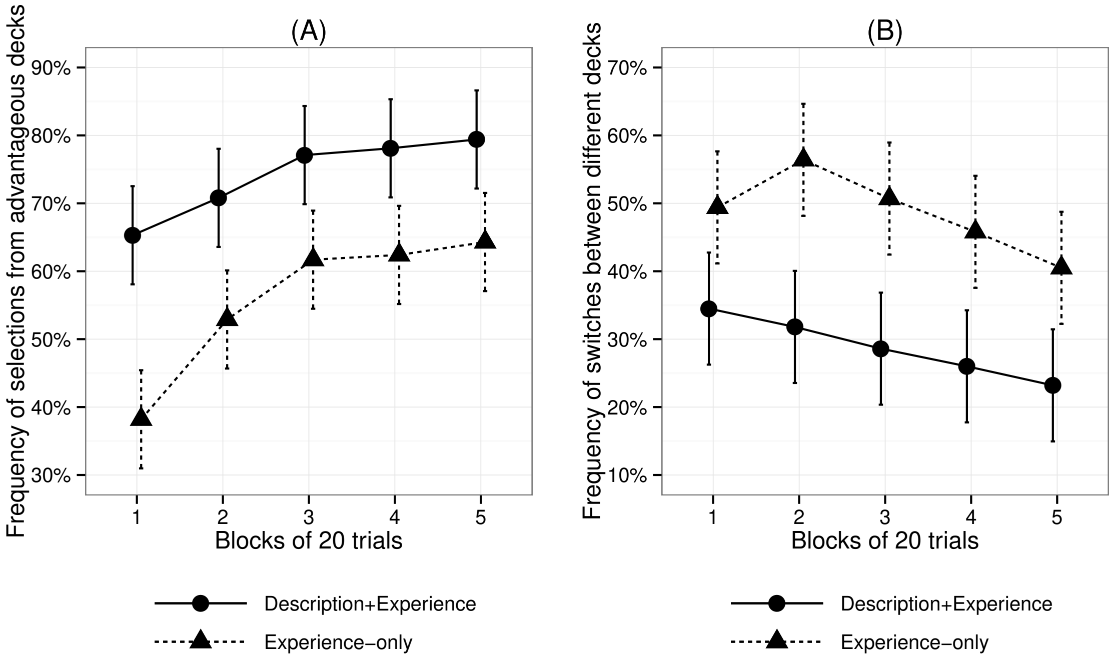

Figure 2. Experiment 1. (A) Evolution of the average frequency of selection of advantageous decks (Decks C + D), as a percentage of total for each block. (B) Evolution of the average frequency of switches between different decks, as a percentage of total for each block. Each block contains 20 trials. Error bars represent the 95% confidence interval around the mean. Experience-only results are for the traditional IGT, while the description+experience results are for the DIGT.

The main effect of the presence of descriptions was significant, with participants selecting from advantageous decks significantly more frequently in the DE condition than in the E condition, across all block of 20 trials on average, with a large effect size (DE=74.14%, E=55.90%, $\chi^2(1)$=19.78, p<.001, d=0.89). The presence of descriptions helped participants identify the advantageous decks and select from them more often.

The main effect of block was also significant ($\chi^2(4)$=88.65, p<.001). A linear contrast was positive and significant (b=0.49, $t(392)$=8.66, p<.001), indicating a higher selection of advantageous decks over time, while a quadratic contrast was negative and significant (b=–

0.24, $t(392)=3.57$, $p<.001$), indicating that the rate of increase reduces over time, as seen in Figure 2A. Post-hoc analyses with Tukey adjustments showed a significantly higher selection of advantageous decks in the last block compared to the first block (Block1=51.75%, Block5=71.85%, $t(392)=8.01$, $p<.001$, $d=1.13$). This increase in selections from the advantageous decks over time is a result of participants gradually learning the task, and being able to identify which decks are the advantageous ones, as well as learning to avoid the disadvantageous ones, in order to extract higher rewards from the task. Most of the learning however appears in the first two blocks (Block1 vs. Block2: $t(392)=4.02$, $p<.001$; Block2 vs. Block3: $t(392)=3.01$, $p=.02$), with no significant differences when applying sequential pairwise comparisons between the last three blocks (pairwise ps$>.97$). This early stabilization of choice preferences is consistent with previous research (Bechara, Damasio, Tranel, & Damasio, 1997; Ert & Erev, 2007).

The interaction between presence of description and block was not significant ($\chi^2(4)=8.17$, $p=.09$), suggesting that the selection rate of advantageous decks across time is similar between the two conditions. In order to exclude this effect of learning, and to focus on stable behavior, we performed a post-hoc analysis with Tukey adjustments comparing the two conditions at the last block: in the last 20 trials, the presence of descriptions led to a 23% significant increase in the selection of advantageous decks, with a medium effect size (Block5 only: DE=79.40%, E=64.30%, $t(230)=2.91$, $p=.004$, $d=0.58$). Even after participants had had a chance to learn about the task experientially, the presence of descriptions still significantly helped them select from the advantageous decks more often.

Switching rates. In addition to the frequency of deck selection, we also analyzed the switching rates between decks (Figure 2B). A selection was classified as a switch every time a card was picked from a different deck to that from which the previous card had been selected. The same model structure was used as in the previous analysis.

The main effect of description was significant, with switching rates being 41% lower in the DE condition compared to the E condition, in each block of 20 trials (DE=28.82%, E=48.56%, $\chi^2(1)=14.48$, $p<.001$, $d=0.76$). Overall participants seemed more uncertain in the E condition and explored more among the different decks, while in the DE condition they exploited more their preferred options, switching less often.

The effect of block was also significant ($\chi^2(4)$=38.11, p<.001). A linear contrast was negative and significant (b=–0.28, $t(392)$=5.68, p<.001), indicating a reduction in switching rates over time. Post-hoc analyses with Tukey adjustments showed that switching rates were lower in the last block compared to the first block (Block1=41.95%, Block5=31.85%, $t(392)$=4.51, p<.001, d=0.64). As participants gathered more information from the task, they explored less and exploited their preferred choices more.

The interaction between presence of description and block was not significant ($\chi^2(4)$=5.85, p=.21), suggesting that the rate of reduction in switching rates were not influenced by the descriptions. A post-hoc analysis with Tukey adjustments between the two description conditions at the last block showed that participants switched 43% less often in the DE condition, a significant difference (Block5 only: DE=23.20%, E=40.50%, $t(161)$=2.93, p=.004, d=0.59).

Discussion

When presented with descriptions in the DE condition, participants selected from the advantageous decks more often than when they had to rely on experience without descriptions in the E condition. Therefore, the presence of descriptions influenced behavior and helped participants to find the rewarding cards and avoid the loss-generating ones, leading to 66% higher financial bonuses. This difference in behavior is indicative of participants integrating the descriptive information into their decision making processes, as descriptions informed participants about the potential outcomes of their choices, and could be used to identify the advantageous decks. Furthermore, the behavior observed in the E condition (without descriptions, therefore a replication of the traditional IGT paradigm) was similar to that found in previous studies using this task. Frequency of advantageous (good) deck selection across all trials, M=56%, was similar to a weighted mean from a meta-analysis of 39 studies covering 1,427 healthy participants, M=57% (Steingroever et al., 2013, Table 5).

Better performance across all trials can be partially explained by the availability of additional descriptive information from the beginning of the task in the DE condition, which provided an advantage to participants, and could be used to make an initial informed choice among the available options. Those in the E condition lacked any information about the composition of cards within each deck, so their first selection was necessarily a random uninformed choice between the four decks available. We can attempt to remove this advantage by comparing the behavior after it has stabilized. Even after many trials, in the last block of 20, deck selection still differed significantly between the E and DE conditions, with participants in the DE condition choosing 23% more often from the advantageous decks than in the E condition. At this point, choice behavior had mostly stabilized.

We also expected exploration to reduce when descriptions were available, and this was observed with lower switching rates in the DE experimental condition. Switching rates can be seen as proxies for exploration (Ert, Erev, & Roth, 2011), as individuals who exclusively exploit their preferred option would not need to switch between the options. In decisions from experience, without descriptions, participants must learn about the decision environment through exploration and feedback. If descriptions are available, by providing additional information about the available options, they offer an alternative avenue for comparing them and finding the most attractive one, reducing the need for exploration. Exploration still remains however, as uncertainty is not fully eliminated, and participants still need to confirm that descriptions are true throughout the task. In addition, participants might be exploring to avoid the boredom of selecting the same alternative repeatedly, or to select a mix strategy across their preferred alternatives.

Overall, the presence of descriptions influenced behavior in a complex task such as the DIGT. However, we were concerned that the usage of a pre-determined fixed schedule of outcomes, as in the original IGT, was not being truly represented by the descriptions. While the descriptions were a true representation of the frequency of the cards within each deck, there was no mention of the actual sequence in which the cards appeared, which might have led participants to believe that cards were shuffled and their order was random. While the original pre-determined sequence is one of the many potential sequences in which the cards would appear if the outcomes were randomized, since they were actually previously known and fixed, descriptions could have also shown participants the sequence of cards. In order to make descriptions a truer representation of the experience, in the next experiment we replaced the fixed schedule with a randomized ordering of cards within each deck.

Experiment 2

Method

Design. The aim of Experiment 2 was to replicate the results found in Experiment 1, and confirm that the presence of descriptions influence behavior in complex tasks with congruent descriptions. As before, there were two conditions: experience-only (E) and description-plus-experience (DE). The only alteration to the paradigm was in the ordering of cards within each deck. Instead of using the original pre-determined sequence of cards from Bechara et al. (1994), we used a pseudo-randomized approach within blocks of 40. This approach should make the experience a truer representation of the descriptions. Since the actual order of cards was not known until the computer randomized it, information about the sequence could not have been provided in the descriptions to participants.

Participants. We recruited 100 participants on-line using Amazon's Mechanical Turk service (42 females; age: M=36.7 years, SD=11.8 years), 49 in the experience-only (E) condition and 51 in description-plus-experience (DE). Participation was restricted to individuals whose location was defined as in the United States. No participants were excluded from the analysis. Participants were paid a fixed amount of US$ 0.25 for participating and an additional bonus according to the outcomes of the choices they made during the experiment, with an average bonus of US$ 0.52 (SD=US$ 0.33)[3].

Task. The task was a replication of Experiment 1, and closely followed the original IGT (Bechara et al., 1994). Participants in the experience-only (E) condition did not receive any additional information, while participants in the description-plus-experience (DE) condition were shown a description of cards underneath each deck. The only alteration from Experiment 1 was in the ordering of cards within each deck. In Experiment 1, the order was always fixed and known beforehand, following the set sequence from the original IGT study, with each participant observing the same ordering of cards. In Experiment 2, we abandoned this pre-determined fixed ordering and opted for a pseudo-randomized approach (Camilleri & Newell, 2011). The frequency of cards within each deck was the

[3]Participants in the DE condition received a significantly 78% higher bonus than participants in the E condition, according to an asymptotic Wilcoxon rank sum test (DE=US$0.66, E=US$0.37, W=1922, z=4.49, p<.001). In addition, there was no significant difference in bonus between Experiments 1 and 2 (W=5122.5, z=0.72, p=.76)).

same as in Experiment 1 (Table 1), but their order was shuffled. Each participant observed a newly randomized ordering of cards. Pseudo-randomization was used to ensure that within each set of 40 cards, participants experienced the same frequency of cards as that in the descriptions, but in a random order. For example, for Deck D, there would always be 36 cards of +50 and 4 cards of –200 points within each set of consecutive 40 cards, with a newly randomized order in each set. The actual sequence was not known until the computer randomized it. The reason for choosing 40 cards is to replicate the original IGT which was also based on sets of 40 cards. This approach is similar to using a deck of 40 cards, which is initially shuffled and revealed in order without replacement. Once all 40 cards of a deck have been shown, the computer would re-shuffle and start again. The task was self-paced and was completed on average in 8.82 minutes (SD=4.12).

Results

As before, the main dependent variable was the frequency of cards selected from the advantageous decks, calculated as the total number of cards selected from decks C and D for each sequential block of 20 choices (Figure 3). They were analyzed with a linear mixed-effects model as in Experiment 1, with the same fixed and random components.

The main effect of the presence of descriptions was significant, with participants selecting from advantageous decks significantly more frequently in the DE condition than in the E condition, across all block of 20 trials on average, with a large effect size (DE=74.67%, E=54.02%, $\chi^2(1)=16.75$, $p<.001$, $d=0.82$). The presence of descriptions helped participants identify the advantageous decks and select from them more often.

The main effect of block was also significant ($\chi^2(4)=18.51$, $p<.001$), with an increase in selection of advantageous decks over time (Block1=58.00%, Block5=69.55%, $t(392)=4.03$, $p<.01$, $d=0.57$), albeit with a smaller effect size when compared to Experiment 1. A linear contrast was significant with a positive slope ($b=0.27$, $t(392)=4.19$, $p<.001$). The interaction between presence of description and block was not significant ($\chi^2(4)=2.01$, $p=.73$).

We also performed a post-hoc analysis with Tukey adjustments comparing the two experimental conditions at the last block: in the last 20 trials, the presence of descriptions led to a 33% significantly higher selection of advantageous decks, with a medium effect size

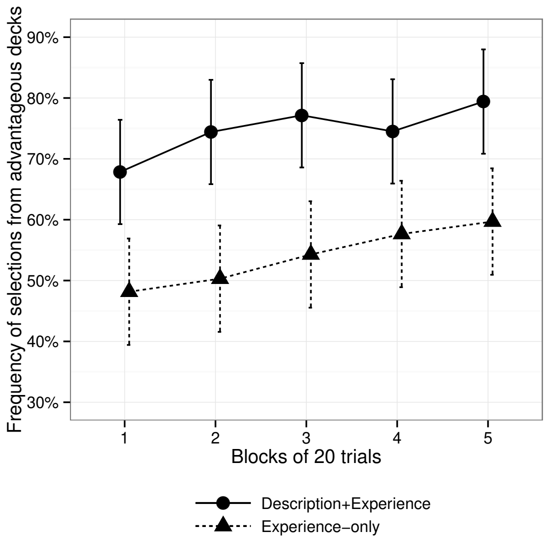

Figure 3. Experiment 2. Evolution of the average frequency of selection of advantageous decks (Decks C + D), as a percentage of total for each block. Each block contains 20 trials. Error bars represent the 95% confidence interval around the mean. In Experiment 2 the order of the cards within each deck was pseudo-randomized, the only change in the task in comparison to Experiment 1.

(Block5 only: DE=79.41%, E=59.69%, $t(211)=3.17$, $p=.002$, $d=0.63$). Even after participants had had a chance to learn about the task experientially, the presence of descriptions still significantly helped them select from the advantageous decks more often.

Discussion

We replicated the findings from Experiment 1 in Experiment 2, with very similar effect sizes for results comparing the presence of descriptions[4]. Participants presented with descriptions selected from advantageous decks more often, and obtained higher financial bonuses. In this experiment the descriptions were a truer representation of the experience, since the actual order of cards was not previously known, until the computer shuffled and randomized them. Only the frequency of the cards within each deck was known, but not their ordering. Because of pseudo-randomization, the frequency described was an exact

[4]We combined the data from Experiments 1 and 2 into one single analysis to evaluate the main effect of Experiment as a proxy for the ordering of cards. There was no significant effect of Experiment ($\chi^2(1)=0.04$, $p=.83$). There was no significant interaction between Experiment and presence of descriptions ($\chi^2(1)=0.14$, $p=.71$). Therefore, the use of a pre-determined schedule or a pseudo-randomized order of cards had no influence on the selection of advantageous decks overall (Exp1=65%, Exp2=64%), or on the impact of descriptions (E: Exp1=56%, Exp2=54%, $t(196)=0.41$, $p=.68$; DE: Exp1=74%, Exp2=75%, $t(196)=0.12$, $p=.91$).

representation of the actual experience, within each set of 40 cards revealed by participants. While in Experiment 1 a true description could have included the actual sequence of cards, this was not possible in Experiment 2.

Across the first two experiments, which were complex decisions-from-experience tasks based around the IGT, the presence of congruent descriptions influenced behavior and helped participants, whether the sequence of card was pseudo-randomized or followed the original fixed schedule. These findings initially appear to go against previous studies using simpler tasks that have shown no influence of congruent descriptions on behavior (Lejarraga & Gonzalez, 2011; Weiss-Cohen et al., 2016; Yuviler-Gavish & Gopher, 2011). We propose that it was the increased complexity of the task, with its four options and multiple outcomes, that led to descriptions being taken into account by participants in our experiment, while in previous studies the tasks were simpler, using two options with fewer outcomes. In the next experiment, we sought to analyze how complexity influences behavior in a more controlled experimental set-up, by creating a task that reconciles our results in Experiments 1 and 2 with those in previous research of description-plus-experience.

Experiment 3

Method

Design. In the first two experiments we observed the influence of descriptions in a relatively complex task, and considered the contrast between our results and those obtained in previous research using simpler tasks. In this experiment we manipulated complexity directly. The aim was to start with simple tasks, similar to those used in earlier description-plus-experience research, and then to increase the complexity within the same experimental framework, therefore directly observing how task complexity moderates the influence of descriptions on behavior. To achieve this, we modified Experiment 2 by manipulating the complexity of the task while maintaining the same basic set-up of selecting cards from different decks with and without descriptions throughout. The task followed a $3 \times 3 \times 2$ between-subjects experimental design. We controlled task complexity across two different dimensions: the number of decks of cards available for participants to choose, which was 2, 4, or 6; and the number of potential outcomes within each choice (i.e., the number of different types of card that composed each deck), which was also 2, 4, or 6. This created a matrix of 3×3 tasks (see Figure 5). Within each cell of this matrix, different participants were given either an experience-only task (E), with no descriptions, or a description-plus-experience task (DE), with descriptions.

When comparing the complexity of this experiment with the previous ones, Experiments 1 and 2 closely matches the central cell of the new experimental matrix of Experiment 3. The previous experiments, based on the IGT, had a total of 14 potential outcomes across its 4 different choices, each choice having an average of 3.5 outcomes. The central cell in the new experiment has a total of 16 outcomes split across 4 different choices with 4 outcomes each (see Table 2). Therefore the new experiment creates both a simpler task (2 choices $\times$ 2 outcomes) and a more complex task (6 choices $\times$ 6 outcomes) in comparison to Experiments 1 and 2 (4 choices $\times$ 3.5 outcomes on average). The simplest task of the new experiment, with 2 choices and 2 outcomes is similar to previous research in the field of description-plus-experience (Lejarraga & Gonzalez, 2011; Weiss-Cohen et al., 2016). The most complex task, with 6 choices and 6 outcomes is considerably more complex than what has been researched before in this field.

The reason for expanding the experiment into highly complex tasks is because we believe that the relationship between task complexity and the influence of descriptions is non-monotonic. As observed in earlier research, in simple tasks, descriptions have no perceptible influence on behavior (Lejarraga & Gonzalez, 2011; Weiss-Cohen et al., 2016). In Experiments 1 and 2 we noticed that by increasing task complexity, descriptions provide useful information to participants and assist behavior, because the task is now more complex and learning experientially is no longer trivial. However we also believe that when the task becomes overly complex, the descriptive information becomes too extensive and therefore also difficult to decipher. In this case, we expected overall performance to suffer in the experience-only condition, and also did not expect much improvement due to the addition of description.

Participants. We recruited 540 participants on-line using Amazon's Mechanical Turk service (239 females; age: M=33.2 years, SD=10.1 years), 30 in each experimental condition. Participation was restricted to individuals whose location was defined as in the

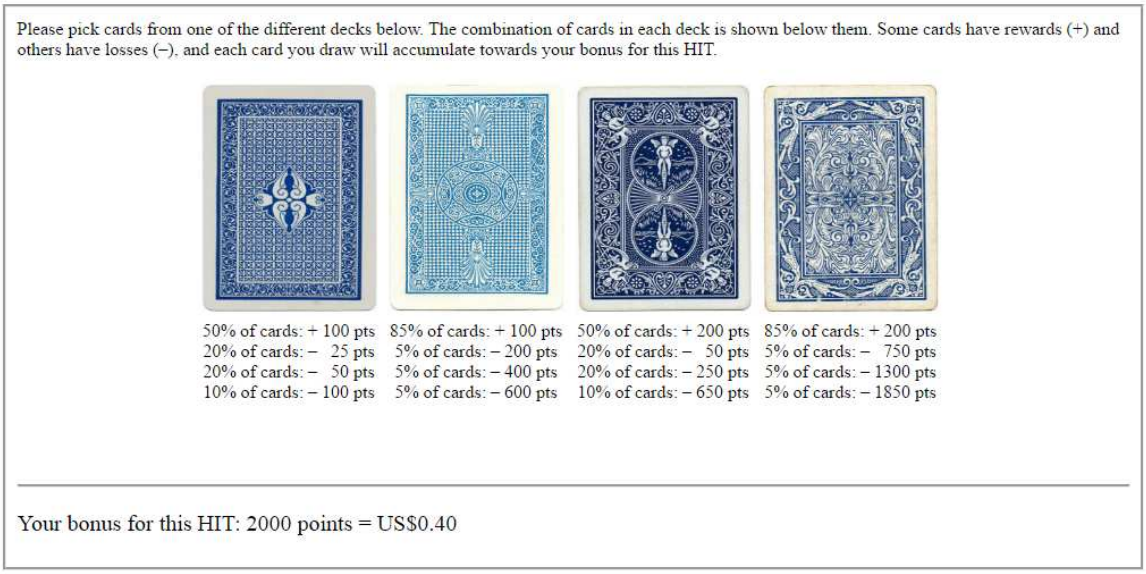

Figure 4 . Screenshot of the first trial of Experiment 3 in the 4-deck × 4-outcome condition with description-plus-experience (DE). In the 2-deck condition only the two middle card positions were used, while in the 6-deck condition an additional two cards were shown in the leftmost and rightmost empty spaces. The order of the decks and the patterns on the back of each deck were both randomized. In this example, from left to right, the decks are C, D, A, and B. Descriptions were not shown in the E condition, and the second sentence in the title was also changed to "Each deck contains a different combination of cards".

United States. No participants were excluded from the analysis. Participants were paid a fixed amount of US$ 0.25 for participating and an additional bonus according to the outcomes of the choices they made during the experiment, with an average bonus of US$ 0.44 (*SD*=US$ 0.32)[5].

Task. The task closely followed Experiment 2, with randomized outcomes, apart from changing the number of options available and number of potential outcomes from each option. Each participant was allocated to a single experimental condition across number of decks, number of outcomes, and presence of descriptions. Participants were presented with either 2, 4, or 6 decks of cards. We created a total of 6 decks of cards, with decks A, B, C, and D relatively similar to their IGT counterparts (Table 2). Participants in the 2-deck condition were presented with decks A and C; participants in the 4-deck condition were presented with decks A, B, C, and D; and all decks were presented to participants in the 6-deck condition. The order of presentation of the decks was randomized, as well as

[5]Participants in the DE conditions received a significantly higher bonus than participants in the E condition, according to an asymptotic Wilcoxon rank sum test (DE=US\$0.50, E=US\$0.39, W=44825, z=4.47, p<.001).

the patterns on the back of the decks. Decks of cards were shown side by side, with the 2- and 4-deck conditions using only the central 2 and 4 positions, respectively (Figure 4). To ensure that participants could see all the decks at the same time, the size of the window used was recorded, and no participant had a window size smaller than the minimum required.

Table 2

Schedule of outcomes used in Experiment 3, written as pairs of "probability: points". In the 2-choice conditions, decks A and C were used. In the 4-choice conditions, decks A, B, C and D were used. In the 6-choice conditions, all decks were used. The actual description text presented to participants followed that of Experiment 1, in the form of "—% of cards: — pts" (see Figure 4). The expected value for each individual card in decks A, P, and B was –25 points (the disadvantageous decks), and in decks C, Q, and D it was +25 points (the advantageous decks).

2 outcomes					
Deck A	Deck P	Deck B	Deck C	Deck Q	Deck D
50%: +200	70%: +200	85%: +200	50%: +100	70%: +100	85%: +100
50%: –250	30%: –550	15%: –1300	50%: –50	30%: –150	15%: –400
4 outcomes					
Deck A	Deck P	Deck B	Deck C	Deck Q	Deck D
50%: +200	70%: +200	85%: +200	50%: +100	70%: +100	85%: +100
20%: –50	10%: –150	5%: –750	20%: –25	10%: –50	5%: –200
20%: –250	10%: –550	5%: –1300	20%: –50	10%: –150	5%: –400
10%: –650	10%: –950	5%: –1850	10%: –100	10%: –250	5%: –600
6 outcomes					
Deck A	Deck P	Deck B	Deck C	Deck Q	Deck D
20%: +255	30%: +250	40%: +280	20%: +135	30%: +125	40%: +140
20%: +200	30%: +200	25%: +200	20%: +100	30%: +100	25%: +100
10%: +90	10%: +50	20%: +40	10%: +30	10%: +25	20%: +20
20%: –50	10%: –150	5%: –750	20%: –25	10%: –50	5%: –200
20%: –250	10%: –550	5%: –1300	20%: –50	10%: –150	5%: –400
10%: –650	10%: –950	5%: –1850	10%: –100	10%: –250	5%: –600

Each deck had either 2, 4, or 6 potential outcomes according to the experimental condition. In contrast to Experiment 2 in which each deck had a different number of outcomes, resembling the payoff schedule of the original IGT, all decks in Experiment 2 had the same number of outcomes within each condition (either 2, 4, or 6 outcomes). The outcomes within each deck were adapted from Chiu and Lin (2007) and are shown in Table 2. Decks A, B, and P have a negative expected value of –25 points for each card, while decks C, D and Q have a positive expected value of +25 points for each card. Hence decks A, B and P were considered the disadvantageous decks, and decks C, D and Q were considered the advantageous decks. Similarly to the IGT, decks A and C were balanced in terms of wins and smaller losses, decks B and D had frequent wins and infrequent but larger losses, and the new decks P and Q were a compromise between them. The schedule of outcomes was pseudo-randomized in sets of 20: within each 20 cards, there was a full representation of all cards for that deck, in the correct proportions, but in randomized order. For example, for Deck A with 2 outcomes, there would always be 10 cards of +200 and 10 cards of –250 within each set of consecutive 20 cards, with a newly randomized order in each set. Participants were either given description-plus-experience or experience-only. Each participant made 100 selections and the task was completed on average in 8.4 minutes (SD=4.1).

Results

The main dependent variable was the frequency of cards selected from the advantageous decks, calculated as the total number of cards selected from decks C, Q and D for each sequential block of 20 choices (Figure 5). It was analyzed with a linear mixed-effects model as in Experiments 1 and 2. The fixed effects were the number of decks (2, 4, or 6), the number of outcomes in each deck (2, 4, or 6), the presence or absence of descriptions (E or DE), the blocks of 20 choices each (with polynomial contrasts), and their interactions. The model also contained a random intercept for each participant. Post-hoc analyses were Tukey adjusted.

The main effect of the presence of descriptions was significant, with participants selecting from advantageous decks significantly more frequently in the DE condition than in the E condition, on average across each block of 20 trials (DE=63.47%, E=55.90%, $\chi^2(1)$=14.75, p<.001, d=0.33). The presence of descriptions helped participants identify the advantageous decks and select from them more often. There was also a significant main effect of number of decks ($\chi^2(2)$=24.43, p<.001), indicating that an increase in task complexity, as measured by number of decks, led to a decrease in selection from advantageous decks, as the task became harder and more difficult to identify the advantageous decks (2-decks: 66.42%, 4-decks: 57.59%, 6-decks: 55.04%). There was no significant main effect of number of outcomes ($\chi^2(2)$=4.71, p=.09), with no influence of number of outcomes on selections from advantageous decks (2-outcomes: 61.69%, 4-outcomes: 60.65%, 6-outcomes: 56.72%). There

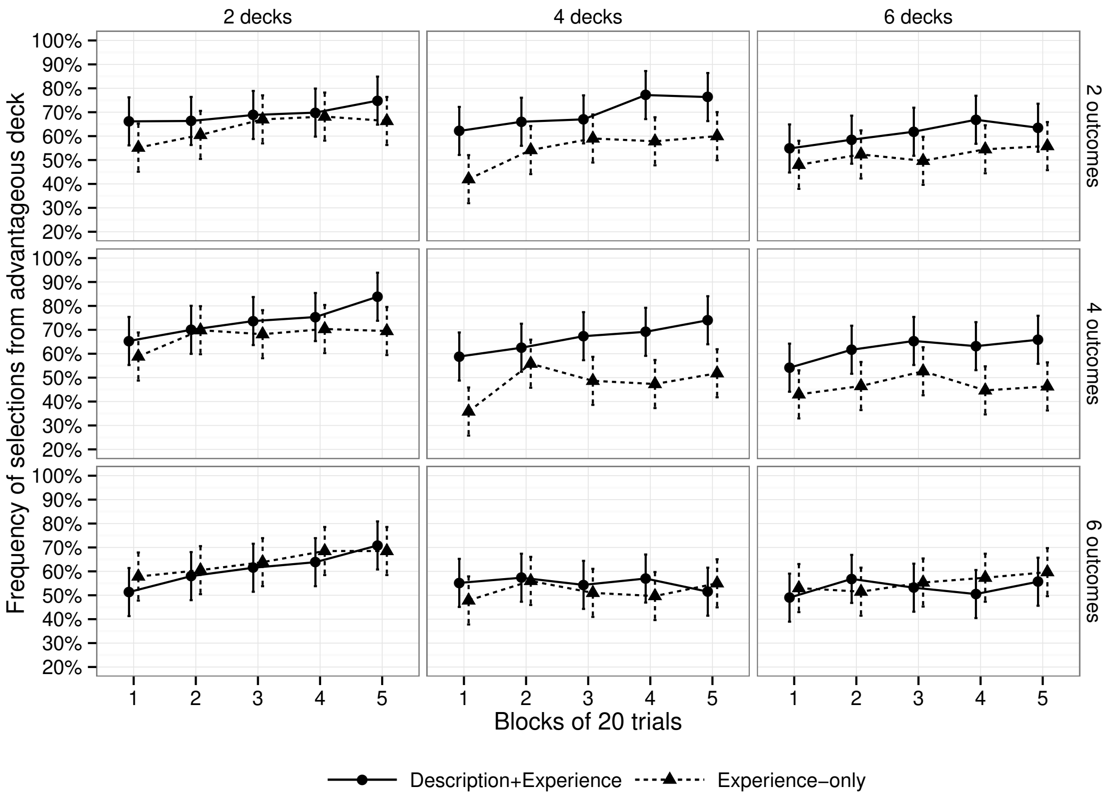

Figure 5. Experiment 3. Evolution of the average frequency of selection of advantageous decks (Decks C + D + Q), as a percentage of total for each block. Each block contains 20 trials. Each data point represents an average of 30 participants, for a total of 60 participants in each cell of the matrix. Error bars represent the 95% confidence interval around the mean.

was a learning effect, as seen by the main effect of block ($\chi^2(4)$=105.68, p<.001), confirmed by a post-hoc analysis that showed higher selection of advantageous decks in the last block when compared to the first block (Block1=53.25%, Block5=63.85%, $t(2088)$=9.64, p<.001, d=0.59), with a significant positive linear contrast (b=0.24, $t(2088)$=9.68, p<.001), and significant negative quadratic contrast (b=–0.08, $t(2088)$=2.61, p=.009).

To further elucidate what drives the influence of descriptions, we analyzed the 2-way interactions between descriptions and the two complexity manipulations separately. The interaction between number of outcomes and presence of description was significant ($\chi^2(2)$=9.15, p=.01). In a post-hoc analysis for number of outcomes, we observed the largest difference in selection of advantageous decks between E and DE in the middle 4-

outcome condition, a smaller but still significant difference in the 2-outcome condition, and no significant difference in the 6-outcome condition (DE–E difference in each condition: 2-outcomes=9.98%, $t(522)=2.92$, $p=.004$, $d=0.44$; 4-outcomes=13.39%, $t(522)=3.92$, $p<.001$, $d=0.58$; 6-outcomes=–0.63%, $t(522)=0.18$, $p=.85$, $d=0.06$). There was no significant interaction between number of decks and presence of description ($\chi^2(2)=3.58$, $p=.17$). However, in a post-hoc analysis for number of decks, we observed the same pattern of a larger difference in selection of advantageous decks between E and DE in the middle 4-deck condition (DE–E difference in each condition: 2-decks=3.11%, $t(522)=0.91$, $p=.36$, $d=0.14$; 4-decks=12.26%, $t(522)=3.59$, $p<.001$, $d=0.53$; 6-decks=7.37%, $t(522)=2.16$, $p=.03$, $d=0.32$).

None of the other 2-, 3-, and 4-way interactions were significant (all ps $\geq .11$). As in Experiment 1, we did not observe an interaction between presence of description and block ($\chi^2(4)=3.90$, $p=.42$), which would indicate no difference in learning due to the presence or absence of descriptions. In other words, changes in selections of the advantageous options progressed similarly across blocks regardless of the presence or absence of descriptions.

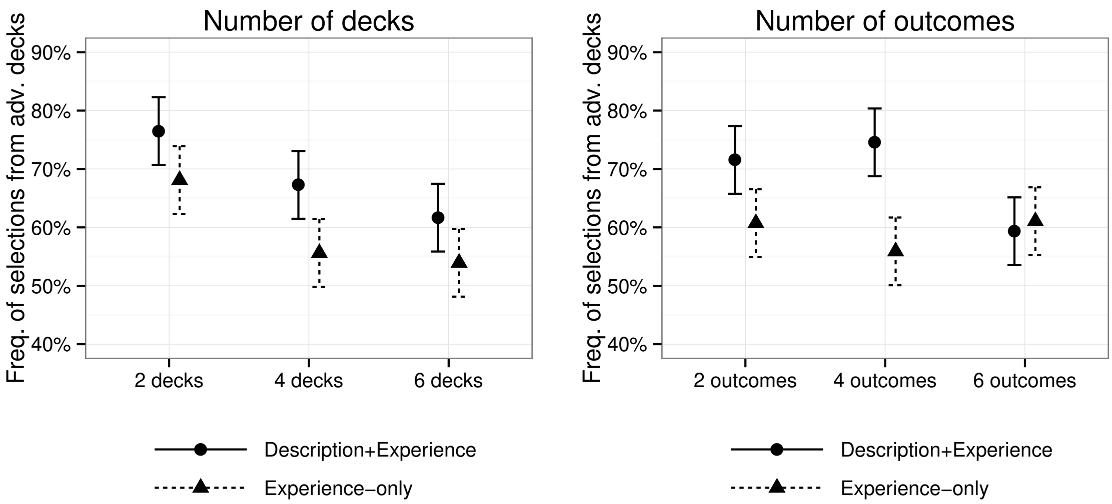

Figure 6. Experiment 3. Average frequency of selection of advantageous decks (Decks C + Q + D) in the last 20 trials (Block 5). Error bars represent the 95% confidence interval around the mean.

To exclude the effect of learning, and to focus on more stable behavior, we performed a post-hoc analysis at the last block only (Figure 6), with Tukey adjustments. In the last 20 trials, the presence of descriptions led to an overall 16% significant increase in the selection of advantageous decks, with a medium effect size (Block5 only: DE=68.48%, E=59.22%, $t(1101)=3.84$, $p<.001$, $d=0.33$). The small difference and low effect size hide the underlying interactions between the complexity manipulations and the presence of description. The influence of description was highest in the middle condition for number of outcomes (Block5 only, DE–E difference in each condition: 2-outcomes=10.83%, $t(1101)=2.59$, $p=.01$, $d=0.39$; 4-outcomes=18.67%, $t(1101)=4.46$, $p<.001$, $d=0.67$; 6-outcomes=–1.72%, $t(1101)=0.41$, $p=.68$, $d=0.06$). The influence of description was also highest in the middle condition for number of decks (Block5 only, DE-E difference in each condition: 2-decks=8.39%, $t(1101)=2.01$, $p=.045$, $d=0.30$; 4-decks=11.67%, $t(1101)=2.79$, $p=.005$, $d=0.42$; 6-decks=7.72%, $t(1101)=1.85$, $p=.07$, $d=0.28$). We observed a non-monotonic inverted U-shaped pattern of influence of description across both complexity manipulations (Figure 7).

Discussion

We observed the same overall influence of the presence of descriptions in Experiment 3 as in Experiments 1 and 2. As before, participants who were presented with descriptions selected from the advantageous decks 14% more often than participants who did not receive descriptions. These figures are lower than the comparable results in Experiments 1 and 2 because they hide the intricate underlying relationship between task complexity and the influence of descriptions, which followed a non-monotonic inverted U-shaped pattern (Figure 7), as predicted.

In simple tasks, in which the payoffs were relatively easy to learn experientially, participants did not benefit from the presence of descriptions. They performed well with experience alone, finding the advantageous decks, and performance did not improve significantly by adding descriptions. This lack of influence from descriptions in simple tasks replicated findings in previous similar research using two alternatives (Lejarraga & Gonzalez, 2011; Weiss-Cohen et al., 2016). In very complex tasks, the payoff structure was not only considerably more difficult to learn experientially, but it also required very long and verbose descriptions. Reading, analyzing, and deciphering the written descriptions very likely presented a substantial cognitive challenge. In these situations, participants were not

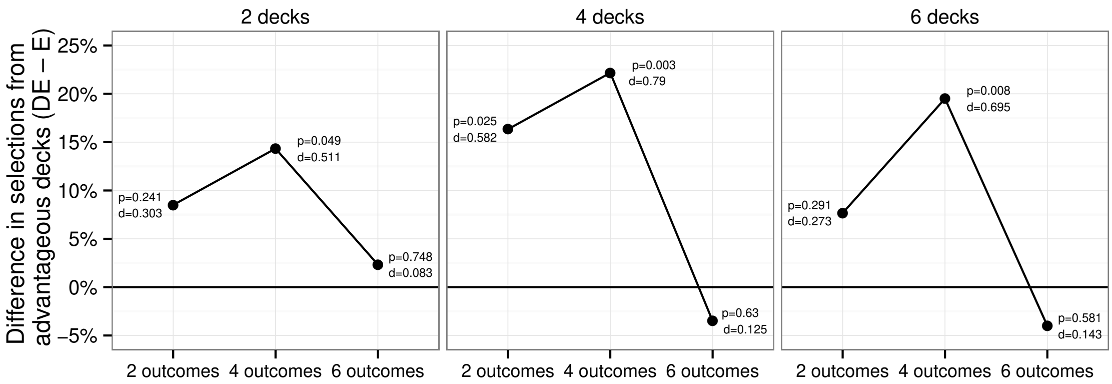

Figure 7. Experiment 3. Difference in frequency of selection of advantageous decks (Decks C + Q + D) in the last 20 trials (Block 5) due to the presence of description (DE–E). The values for the significance level (p-values) and effect size (Cohen's d) for each difference are shown next to each data point. The highest influence of description can be observed in the middle condition: 4-decks x 4-outcomes.

as proficient in finding the advantageous decks as in the simpler tasks, leading to lower performance via experience alone, and complicated descriptions did not provide any significant assistance.

It was in the medium complexity tasks that descriptions most helped participants. In the middle cell of Experiment 3 (4 choices × 4 outcomes), descriptions led to a 39% increase in overall selections of advantageous decks (43% in the last block), and an 82% increase in the final financial bonus, the highest across all conditions and similar to the results from Experiments 1 and 2, which closely matches this task in terms of overall complexity. These medium complexity tasks were too complex to be learned via experience alone, but not overly complex to be described succinctly and for these descriptions to be analyzed by and provide useful information to participants. Descriptions are most influential in decision making when they provide additional information to participants that cannot be as easily and efficiently gathered in a timely fashion via experience alone.

Cognitive modeling

In order to theoretically model the role of descriptions in decisions from experience, and their interaction with task complexity, a set of cognitive computational models was fitted to the experimental data. Our aim was to evaluate how descriptions influence traditional reinforcement learning models, and how the descriptive information is represented and integrated into the decision-making process.

We start by fitting a reinforcement learning (RL) model to the observed human behavior in the experience-only conditions of Experiment 3. RL models that rely solely on experience via feedback have been extensively and successfully used to explain behavior in decisions from experience in the past (Erev & Barron, 2005; Yechiam & Busemeyer, 2008; Yechiam & Ert, 2007), and the IGT is paradigm commonly explored with these models (Ahn, Busemeyer, Wagenmakers, & Stout, 2008; Busemeyer & Stout, 2002; Dai, Kerestes, Upton, Busemeyer, & Stout, 2015; Worthy, Pang, & Byrne, 2013; Yechiam & Busemeyer, 2005). Based on their past performance in similar tasks, we would expect RL models to fit our behavioral data well in the experience-only conditions of the current experiments. With regards to tasks combining descriptions and experience, Lejarraga and Gonzalez (2011) have shown that a simple RL model can also explain behavior well in simple tasks with descriptions, and we expect to replicate their findings here in our simple tasks. However, given the observed difference in behavior when description was added to these paradigms in more complex tasks, and our previous modeling efforts (see Weiss-Cohen et al., 2016), we predict that traditional experience-only RL models will perform poorly in the description-plus-experience conditions. As shown here by the observed empirical results, descriptions can sometimes influence behavior and can provide additional useful information for participants to perform better in their tasks. These situations should be conducive for a model that combines descriptions and experience.

Below, we present a description-plus-experience model that combines both descriptive and experiential information, which should help explain the observed differences in behavior in the description-plus-experience conditions. Our previous attempt at a description-plus-experience model combined the two sources of information with different weights, and the weights determined the importance given to each source depending on the experimental condition (see Weiss-Cohen et al., 2016). Since the influence of descriptions in Experiment 3 appeared to have been moderated by the complexity of the task, we will vary these weights according to a complexity measure, based on entropy. In comparison with traditional experience-only RL models, we expect the description-plus-experience model to provide a better fit for the observed human behavior in description-plus-experience tasks, in particular in more complex tasks, with little or no difference in the simpler tasks.

The models

The aim of fitting a cognitive model to the data was to assess and formalize how the two sources of information, descriptive and experiential, are combined. We fitted two models to the behavioral data: an *experience-only* prospect-valence learning (PVL) model and a *description-plus-experience* adaptation of that model (D-PVL). The PVL model is a reinforcement-learning model that relies on experiential information alone, using the feedback provided after each trial (Ahn et al., 2008; Fridberg et al., 2010). The D-PVL model built upon that, combining the experience-only RL component from the PVL model with a representation of the descriptive information. The descriptive component was calculated as the expected value of the information presented to participants underneath each choice. The two sources of information were combined using a weight, which was determined via entropy, a proxy for task complexity. Crucially, the experiential part of the D-PVL model was based around the same RL model as the PVL model. Therefore the D-PVL model added descriptions to a traditional experience-based RL model, and we are particularly interested in how this integration was performed. We start by describing the PVL model, which formed the basis of both models.

Experience-only model (PVL)

We build our models upon one successful RL model from the literature, a prospect-valence learning (PVL) model using a prospect-theory utility function and a delta-learning rule. This model has been extensively and efficiently used in the decisions-from-experience literature, in particular using the IGT, and shown to perform better than competing models when fitting experimental data to simulated participants, which is the approach we used here (Ahn et al., 2008; Ahn, Krawitz, Kim, Busemeyer, & Brown, 2011; Dai et al., 2015; Fridberg et al., 2010; Worthy et al., 2013; Yechiam & Busemeyer, 2005, 2006).

Firstly, observed payoffs are evaluated by a prospect-theory type of utility function (Kahneman & Tversky, 1979), $U(\cdot)$, defined as:

$$U(r_j(t)) = \begin{cases} (r_j(t)/100)^{\alpha}, & \text{if } v_j(t) > 0. \\ -\lambda(-r_j(t)/100)^{\beta}, & \text{if } v_j(t) < 0. \end{cases}$$

where $r_j(t)$ is the payoff received from selection option j at time t. Payoff values were divided by 100 to reduce the magnitude of the observed feedback and realign them closer to their monetary payoffs. The free parameters α and β, both ranging between 0 and 2, determine the curvature of the value function for positive and negative payoffs, respectively. Lower values of α and β reduce the distance between extreme values of payoffs, while higher values magnify the distances. The loss aversion parameter, λ, is the free parameter ($0 \leq \lambda \leq 10$) that determines higher sensitivity to losses in comparison to gains. The higher the value of λ, the higher the importance given to losses over gains.

Secondly, expectancies for the value of rewards for each option are formed via a learning rule, which integrates the experienced feedback after each trial. The learning rule used was a delta rule, which uses a learning rate that determines how much the new information gathered via feedback, in the form of prediction error, influences the updating of the expectancies at each trial (Speekenbrink & Konstantinidis, 2015; Konstantinidis, Ashby, & Gonzalez, 2015; Sutton & Barto, 1998; Yechiam & Busemeyer, 2005, 2006). Feedback observed is integrated after each trial, to arrive at the experienced expectancy $E_j(t)$ for option j at time t:

$$E_j(t) = E_j(t-1) + \phi \cdot \delta_j(t) \cdot [U(v_j(t)) - E_j(t-1)],$$

where ϕ is the free learning rate parameter ($0 \leq \phi \leq 1$), which is a weight given to new information observed, with lower values resulting in slower learning. The variable $\delta_j(t)$ is a dummy variable, which is equal to one if option j was chosen on trial t, and zero otherwise. The model only updates the value $E_j(t)$ of an option when that option has been selected and its feedback has been observed. When the option has not been selected, the $E_j(t)$ remains unchanged. The initial value for E_j was set to zero[6].

[6]Attempts to change this to the value of descriptions, D_j, in the description-plus-experience model, led to worse fitting models, as it resulted in more constant behavior over time with a flatter learning curve. This is likely due to the participants exploring their options in the beginning of the task even when descriptions were present, a behavior that would have been suppressed by a model with a non-zero starting E_j.

Finally, the model-predicted probability of selecting a given option j at time t is determined by a time-dependent Softmax rule (Sutton & Barto, 1998) that combined the expected values E_j across all options J:

$$\hat{P}_{jt} = \frac{e^{\Theta \cdot E_j(t)}}{\sum_J e^{\Theta \cdot E_J(t)}},$$

where Θ is the choice sensitivity. If $\Theta = 0$, the model randomly guesses between the options regardless of their expectancies, while higher values of Θ will lead to more deterministic maximization behavior. Θ itself is time-dependent and varies according to t, and is determined by the free parameter θ, $(0 \leq \theta \leq 2)$:

$$\Theta = (t/10)^{\theta},$$

this allows choice sensitivity to increase over time, making selections more random in the beginning and more deterministic as time progresses, to reflect the natural tendency of individuals to explore more in the beginning of tasks and less as the task progresses and they have gathered more information from the environment. Values of θ below 1 make the shape of the choice sensitivity over time concave, while values above 1 make it convex, and linear when $\theta = 1$.

Description-plus-experience model (D-PVL)

In the description-plus-experience (D-PVL) model, a representation of descriptions for each choice j, D_j, was combined with the experience, E_j, at each trial t, as follows:

$$ED_j(t) = \omega_c \cdot D_j + (1 - \omega_c) \cdot E_j(t).$$

The experience component, $E_j(t)$, was calculated using the same PVL approach as in the experience-only model, although new parameters were fitted. A representation of the descriptive information is included via D_j as the subjective expected value of the descriptive information for choice j, calculated using cumulative prospect theory (CPT), based on the descriptions provided to participants underneath each alternative. According to Tversky and Kahneman (1992), the CPT value is calculated using a value and a probability-weighting function, $W(\cdot)$ and $U(\cdot)$ respectively,

$$D_j = \sum_m W(p_{jm})U(v_{jm}),$$

where p_{jm} are the probabilities and v_{jm} are the potential values for each outcome m of option j. $U(\cdot)$ is the same function as defined above for the PVL model, using the same parameters. $W(\cdot)$ is the probability weighting function, defined as:

$$W(p) = \frac{p^\gamma}{(p^\gamma + (1-p)^\gamma)^{\frac{1}{\gamma}}},$$

where γ is the free parameter ($0 \leq \gamma \leq 2$) that determines the sensitivity to probabilities via the curvature of the probability weighting function. Values of γ below 1 lead to overweighting of rare events, while values above 1 lead to underweighting of rare events.

Experience, $E_j(t)$, and description, D_j, are combined using ω_c which determines the weight given to description, and its compliment given to experience (Weiss-Cohen et al., 2016). The ω_c weight changes according to experimental condition c, and is calculated as follows:

$$\omega_c = 1 - e^{(-\xi/S_c)},$$

where ξ is a free parameter which determines the strength of the weight given to descriptions ($0 \leq \xi \leq 3$), divided by the entropy S_c, for each condition c, which was calculated according to the choices available to participants. Entropy has been used before to quantify task complexity, with higher entropy associated with higher complexity (Swait & Adamowicz, 2001). According to the weighting formula used, the weight given to descriptions decreases when entropy increases, therefore ω_c is higher in simpler experimental conditions such as 2×2 and lower in more complex ones such as 6×6. An exponential relationship was used to ensure that the weight ω_c remained bounded between 0 and 1, regardless of the values of S_c.

Entropy for condition c, denoted by S_c, was calculated in two different ways, and the two approaches will be compared in the results section. Entropy was initially defined according to the probabilities displayed to participants on the descriptions:

$$S_c = -\sum_{jm}[p_{jm} \cdot log_2(p_{jm})],$$

for all probabilities p_{jm} for every outcome m of option j for that condition.

However while this approach provided a relatively good fit, it did not provide the best fit to the observed behavioral results. We believe this is because the basic RL model with partial feedback already captures some of the idiosyncrasies of having different numbers of alternatives from which to choose. The more alternatives available, the less an individual learns about the environment after each selection, since only information about one choice is revealed at each trial. These findings will be discussed in more detail in the results section below.

We therefore propose an alternative approach to calculating entropy. We divided the total entropy for each condition by the number of alternatives, or number of decks of cards, in that condition, denoted as A_c, which resulted in an average entropy per alternative:

$$S'_c = -\frac{1}{A_c}\sum_{jm}[p_{jm} \cdot log_2(p_{jm})].$$

Alternatively, S'_c can be considered as the entropy of one of the alternatives, chosen randomly between the ones that were available. The model using this averaging entropy approach will be denoted as D-PVL$'$. As an example, the value of S'_c for the simplest condition in Experiment 3, condition 2 $\times$ 2, which had two options and each option had two outcomes with 50% probability each, was $S'_c = -1/2 \cdot (0.5 \cdot log_2(0.5) + 0.5 \cdot log_2(0.5) + 0.5 \cdot log_2(0.5) + 0.5 \cdot log_2(0.5)) = 1.0$. In comparison, entropy for the middle condition, 4 $\times$ 4, was $S'_c = 1.3$, and for the most complex condition, 6 $\times$ 6, it was $S_c = 2.3$. S'_c was mostly influenced by the number of alternatives. In comparison, the values for the total S_c were higher, and increased much faster as task complexity increased.

The same Softmax choice rule from the PVL model is used for the D-PVL models, replacing E_j with ED_j, although as before, new parameters are fitted.

Model fitting

Data sets containing 100 simulated participants were generated for each of the 9 experimental conditions in Experiment 3 (number of decks: 2, 4, or 6 $\times$ number of outcomes: 2, 4, or 6), with the same pseudo-randomized methodology used to generate actual data sets for the experiments in blocks of 20.

We started by fitting the experience-only PVL model to the observed human behavior in the 9 different experience-only (E) experimental conditions of Experiment 3. This model was not allowed to take into consideration the descriptive information, as the participants also did not have access to any descriptions. A total of 900 modeled simulated participants were confronted with 270 observed human participants. All simulated participants across all underlying experimental conditions shared the same set of free parameters. The best fit parameters were found by minimizing the multinomial log-likelihood (LL) between the average observed proportions of choice from each deck and the average model-predicted proportions for each of the individual conditions separately, with each condition receiving the same weight (Erev & Barron, 2005):

$$LL_c = -2 \sum_{jt} ln \left(\frac{N_c!}{\prod n_{jt}!} \prod (\hat{P}_{jt})^{n_{jt}} \right)$$

where N_c is the total number of participants in each condition, n_{jt} is the number of participants who chose option j at trial t, and $\hat{P}_{jt}$ is the model-predicted probability of choosing option j at trial t.

To allow for behavioral differences between the E and DE experimental conditions, we also fitted the PVL model to the DE conditions in Experiment 3. Any changes in the parameters could be explained by a different approach that individuals might have taken towards the task when description was available. However this model still does not allow for the descriptive information itself to be integrated into the decision-making process. We check if the descriptive information was used by participants by fitting the two alternative D-PVL models (with total entropy and with average entropy) against the description-plus-experience (DE) experimental conditions of Experiment 3. The same 100 simulated participants were used as above, but now the model was also allowed to take into account the descriptive information, since this was available to participants. The D-PVL models were fitted in the same way as the PVL model above against the observed human behavior, minimizing the LL.

Because of the different number of parameters between the models, the Bayesian Information Criterion (BIC), which penalizes for additional parameters, was calculated to compare the models, $BIC_c = LL_c + f \cdot ln(N)$, where f is the number of free parameters and N is the number of fitted observations for each evaluation (100 trials). Lower BIC values represent better fitting models. We report the mean BIC_c which is the mean across all 9 conditions.

Model evaluation and results

Three models were evaluated, with four sets of parameters fitted in total: the experience-only PVL model, which did not account for the influence of descriptive information, with five free parameters, was fitted twice, against human behavior in the E (called PVL$_e$) and the DE conditions (PVL$_{de}$), separately, allowing for two sets of different parameters; and the two description-plus-experience (D-PVL) models, one with total entropy and one with the alternative average entropy approach (D-PVL$'$), which combined both descriptive and experiential information, both with seven free parameters, were fitted against human behavior in the DE conditions only.

The results of the model fitting analysis were in line with the behavioral results (Figure 8). The PVL$_e$ model fitted against the human behavior in the E experimental conditions proved a relatively good fit (mean BIC by condition $M_{BIC} = 1,283$). This model was considerably better than a null model, which randomly selects decks of cards at each trial among the available options, returning an $M_{BIC} = 1,484$ in the E conditions. As expected, when comparing the PVL$_e$ model to the human behavior in the DE conditions, the fit was substantially worse overall ($M_{BIC} = 1,403$). This is because of the behavioral differences observed in the experiment, likely a result of the introduction of descriptive information, while the model was not allowed to integrate that new information. It was still considerably better than a random behavior null model in the DE conditions ($M_{BIC} = 1,624$). The higher random BIC for the random null model in the DE conditions indicate

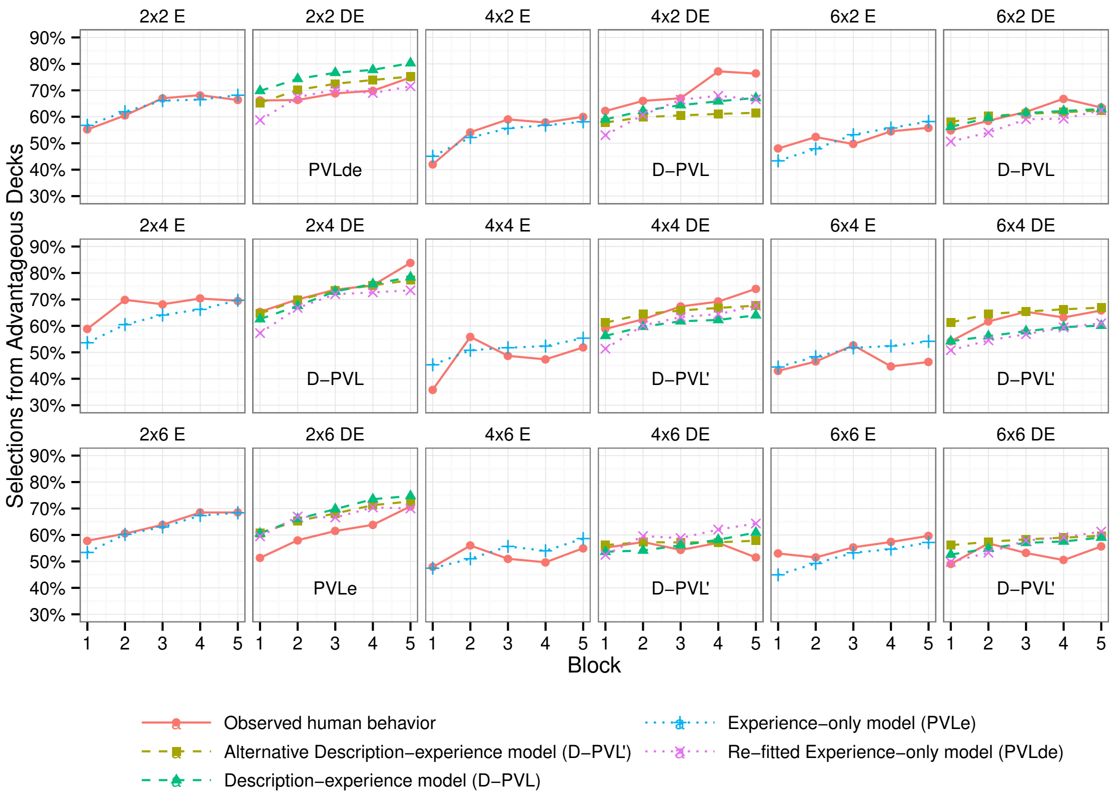

Figure 8. Comparison of observed human data, and the three cognitive models fitted against them, for each of the 18 experimental conditions. The experimental conditions above each cell can be identified as number of decks × number of outcomes followed by E for experience-only and DE for description-plus-experience. The text within the DE cells identify which was the best fitting model for that particular condition (only one PVL_e model was fitted against the E conditions). The best-fitting model overall was the alternative D-PVL′ model using the average entropy approach, in particular in the higher complexity conditions. In medium-complexity conditions, the total entropy model D-PVL provided a better fit. The PVL models fitted better in one condition each, both simpler conditions with 2 decks.

that participants were behaving less randomly when description was present, and therefore a model that predicts random behavior is a poorer predictor of human behavior in the DE condition, but a better predictor in the E conditions, when participants were behaving closer to random.

The fit results were substantially improved by refitting the experience-only model to the behavior in the DE conditions (PVL_{de}), with new parameters ($M_{BIC} = 1,316$), as shown in Table 3. The new PVL_{de} model still did not include any descriptive information, and relied on experience alone. While the original parameters fitted against the E conditions would provide a poor prediction for behavior in the DE conditions, the newly fitted parameters accommodate for some of the differences in observed behavior. However as we will confirm later in a generalization test against Experiments 1 and 2, this is likely a result of overfitting.

Finally, by fitting the D-PVL model that was allowed to take descriptions into account, and using total entropy to moderate the weights given to description, there was an additional improvement in the fit against the DE conditions, with a 2% reduction in BIC ($M_{BIC} = 1,294$), compared to the PVL$_{de}$ model. The alternative model D-PVL$'$, using the average entropy approach, resulted in an even better fit ($M_{BIC} = 1,271$), with an additional 2% reduction in BIC. While the parameters did not change across experimental conditions, we can split the results according to them, and verify in which conditions each model performed better. The D-PVL models were the best performing models with lower BICs in 7 out of the 9 experimental conditions, in particular the conditions with higher task complexity. The two conditions in which they were outperformed by the PVL models were both 2-deck conditions (with 2 and 6 outcomes). In these relatively simple 2×2 and 2×6 conditions, the experience-only PVL model proved a better fit for the observed behavior, replicating the finding from Lejarraga and Gonzalez (2011), who also showed that a simple experience-only RL model without descriptions provided a good fit for human behavior in simple tasks. It is only in more complex tasks, where we predicted that descriptions would be more useful for participants, that the D-PVL models outperformed the PVL models.

The alternative model using average entropy, D-PVL$'$, returned considerably better fits in comparison to the total entropy D-PVL model. Average entropy can be interpreted as the entropy of a single alternative, or deck of cards, selected at random from the ones available. Because all the alternatives in our paradigms contained the same number of potential outcomes, their average entropies did not differ considerably. Comparing the models using Schwarz weights (Wagenmakers & Farrell, 2004) showed a strong preference for D-PVL$'$, with $w(BIC_{DPVL'}) > .9999$, which can be interpreted as the probability that this is the best model among the models presented here (Lewandowsky & Farrell, 2011). We believe that employing average entropy as a moderator of weights given to descriptions yielded better fitting models because the experience component of our models (E_j) already

Table 3

Best fit parameters of the three cognitive models, PVL (fitted twice, against E and DE observed human data in Experiment 3), D-PVL and D-PVL′ (fitted against DE data only), and mean BICs. Lower BICs represent better fits. n.a.=not applicable.

Free parameter	Exp. only E data (PVL_e)	Exp. only Refit DE data (PVL_{de})	Description+ Experience (D-PVL)	Alternative Descr.+Exp. (D-PVL′)
α (curvature of pos. values)	1.23	0.51	1.60	1.26
β (curvature of neg. values)	0.44	0.47	0.53	0.60
λ (weight of neg. values)	1.83	1.82	9.58	2.73
γ (curvature of probabilities)	n.a.	n.a.	0.92	0.84
ϕ (learning rate)	0.31	0.27	0.02	0.06
θ (choice sensitivity)	0.14	0.06	0.22	0.20
ξ (description's weight)	n.a.	n.a.	0.59	2.12
No. of free parameters	5	5	7	7
Mean BIC	1,403	1,316	1,294	1,271
No. of conditions best fit	1	1	3	4

incorporated the deleterious influence of additional alternatives, but it is not influenced by the number of outcomes. The performance of traditional experience-only RL models already deteriorates considerably when dealing with a larger number of alternatives (Konstantinidis et al., 2015). This is specially the case when only partial feedback is available, as the RL model can only update the expectancy of the most recently selected alternative, for which feedback was presented. This results in a smaller reduction of uncertainty about the environment after each trial when more alternatives are available, slowing down the differentiation between the alternatives and, consequently, the ability to identify the better ones. Similar patterns were shown by Ashby et al. (2017). When there are only two options, with each selection, half of the alternatives are updated. With six options, only one sixth is updated. Therefore it takes longer to reduce uncertainty with a traditional RL model when there are more alternatives. There is no similar mechanism for number of outcomes from each alternative, with the RL model incorporating new information in the same manner regardless of the number of outcomes. Traditional RL models as the ones used here are therefore sensitive to number of alternatives but insensitive to number of outcomes. Therefore we believe that by adding an entropy measure that is mostly related to the number of outcomes, not alternatives, we have breached this remaining gap in our D-PVL models, without taking into consideration the effect of number of outcomes twice, which is the case when total entropy is used.

Since we observed a non-monotonic U-shaped curve in the relationship between presence of description and task complexity in the human behavior from Experiment 3, a cognitive model that captures human behavior appropriately should also replicate that finding. We compared the modeled predictions for the last block of trials for the PVL_e and the D-PVL′ models (Figure 9). The inverse U-shaped pattern that was observed in the human behavior was also replicated with cognitive models. Increasing the number of outcomes increased the entropy monotonically, but influence of description, moderated by a weight determined by the inverse of entropy, was non-monotonic.

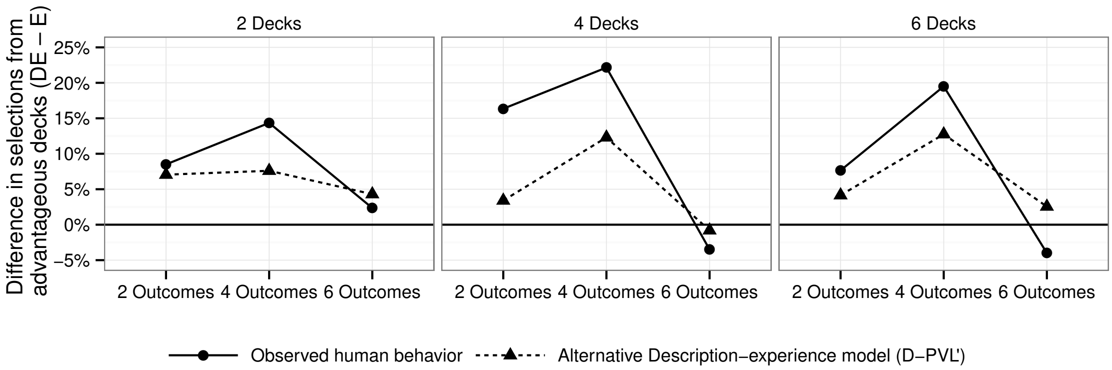

Figure 9. Difference between selection of advantageous decks as predicted by the cognitive models and human behavior, for each experimental condition, in the last block of 20 trials. The modeled difference is the prediction of the experience-only model subtracted from that of the description-plus-experience model.

The best fit parameters were relatively consistent across the different models (Table 3). A few parameters changed in reaction to the presence of descriptions. In particular, the learning rate ϕ was lower in the D-PVL′ model ($\phi_{DPVL'} = 0.06$) compared to the PVL_e model ($\phi_{PVLe} = 0.27$). As participants had access to descriptive information, they did not need to learn as much from feedback after every trial, and updated their expectancies more slowly in the reinforcement learning component of the model. We also observed that the weight given to negative rewards λ, was higher in the D-PVL′ model ($\lambda_{DPVL'} = 2.73$) than in the PVL_e model ($\lambda_{PVLe} = 1.83$). We believe that this is due to the increased relevance of losses when constantly presented in textual descriptions, as they appear more salient and ever-present than the occasionally observed feedback, similar to the "mere presentation" effect (Erev, Glozman, & Hertwig, 2008). The choice sensitivity parameter θ was higher in the D-PVL′ model ($\theta_{DPVL'} = 0.20$) than in the PVL$_e$ model ($\theta_{PVLe} = 0.14$). Higher choice sensitivity translates into more deterministic behavior, and it is likely that the presence of descriptions reduced uncertainty, and allowed participants to be more secure in their decisions, behaving less randomly in their choices. Finally, the weights given to description for the D-PVL′ model, based around the best fitting ξ parameter and calculated according to the ω_c formula above with average entropy, varied between 0.57 for the most complex conditions and 0.92 for the simplest conditions.

Model generalization test

A good computational model should not only be able to agree with the observed data, but also facilitate *a priori* predictions and generalizations into new environments and under altered circumstances (Shiffrin, Lee, Kim, & Wagenmakers, 2008). In order to generalize well, such models must avoid overfitting of task-specific effects, idiosyncratic strategies and heuristics adopted in relation to the experiment against which they were originally calibrated (Konstantinidis, Speekenbrink, Stout, Ahn, & Shanks, 2014), and be able to reliably predict behavior in a new experimental design.

We conducted a generalization analysis of the models, calibrating their parameters with the human data from Experiment 3 (as per the previous sections), and then testing the fitted results against the human data in Experiments 1 and 2, using the *generalization criterion* methodology in Busemeyer and Wang (2000). We simulated new model predictions with the fitted parameters from Table 3 using the two IGT paradigms from Experiments 1 and 2 as the new generalization designs. Two hundred simulated participants were created using the outcomes from the IGT, half using a fixed schedule and half a random schedule. According to Busemeyer and Wang (2000), the best model is the one that produces the smallest discrepancy between model predictions and human data, in our case using mean log-likelihood (M_{LL}), also used for generalizing by Erev and Haruvy (2005).

The PVL$_e$ model was a relatively good predictor for the E conditions in Experiments 1 and 2 ($M_{LL} = 2,284$), which can also be observed in the two E plots in Figure 10. This is not surprising as this model has been extensively used before to predict behavior in the traditional IGT. As in Experiment 3, the PVL$_e$ model was not a good predictor for the DE

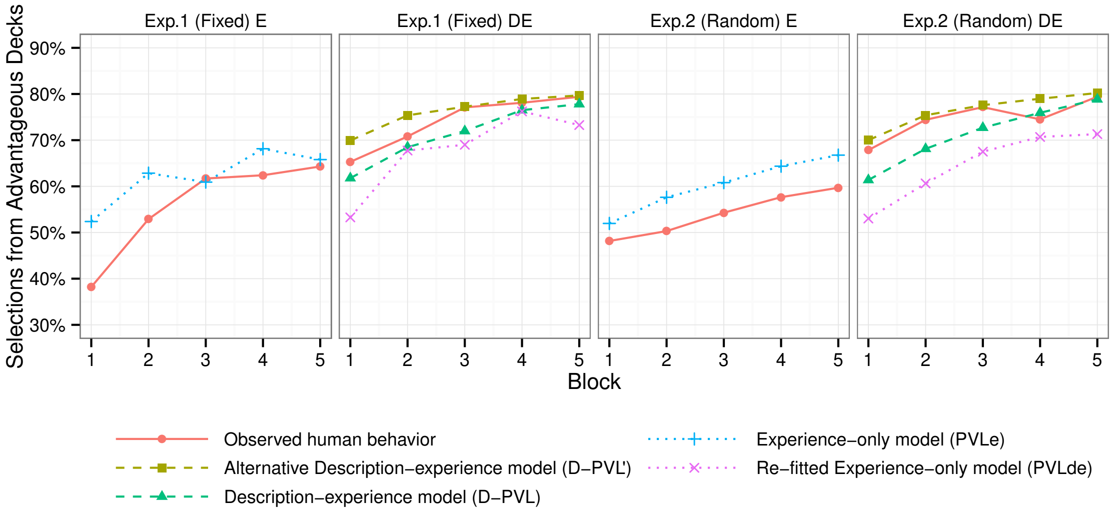

Figure 10. Model predictions tested against the IGT simulated data from Experiments 1 and 2, both with fixed and random schedule of outcomes, using the parameters calibrated with human data from Experiment 3. E refers to experience-only (traditional IGT) and DE to description-plus-experience (DIGT). When descriptions were available, the D-PVL′ model returned much better predictions, ahead of the D-PVL and PVL models.

conditions, with a considerably worse fit ($M_{LL} = 2,492$). The re-fitted PVL$_{de}$ model showed an improvement against the human behavior in the DE experimental conditions, with the new parameters capturing some of the effects of the presence of descriptions ($M_{LL} = 2,238$). The D-PVL models were again better predictors of the behavior in the DE conditions of Experiments 1 and 2, both outperforming the experience-only PVL$_{de}$ model, by much larger margins than in Experiment 3 (DE plots in Figure 10). As before, the alternative model D-PVL′ provided the best fit across all models ($M_{LL} = 1,885$), while the D-PVL was also better than PVL$_{d}e$, but with a smaller improvement ($M_{LL} = 2,038$). Overall, the D-PVL′ model was a much better predictor of human behavior in the DIGT of Experiments 1 and 2 than the PVL model, which was expected given the complexity of this task and the empirical differences in human behavior observed between the E and DE conditions. The models returned a much better fit than the random null model against the IGT, which performed worse in the E conditions ($M_{LL} = 2,778$), and considerably much worse in the DE conditions ($M_{LL} = 3,968$), given how much participants' behavior diverged from random under the latter paradigms. The generalization criterion methodology used here showed that the D-PVL′ model generalized well into a new task design, from Experiment 3 into Experiments 1 and 2, showing a good model fit, despite the added complexity of this model. The simpler PVL_{de} did not generalize as well, despite its relatively good performance earlier, likely as a result of specious overfitting of the PVL_{de} parameters against the specific behavior in Experiment 3.

Discussion on cognitive modeling

Overall, the description-plus-experience D-PVL′ model proved a better fit of the observed human behavior, in particular in the higher complexity experimental conditions. The model also generalized well, from being calibrated against the design of Experiment 3, into the different new designs of Experiments 1 and 2. While previous attempts to model description-plus-experience tasks with congruent information had shown that a traditional experience-only model PVL could be used to explain behavior relatively well (Lejarraga & Gonzalez, 2011), we believe that this was only the case because the tasks used were simple. Simple tasks should be easy to learn experientially, and the addition of descriptive information did not influence behavior, or modeling predictions significantly. This finding has been replicated here in our simplest experimental condition 2×2 (and also 2×6), with the experience-only PVL models providing good fits for the observed human behavior and little difference between predicted model results between PVL and D-PVL (see Figure 9). However when complexity is increased, then descriptions can provide useful information for participants, helping them make better decisions. This was shown empirically in the behavior data, and confirmed with cognitive modeling. If descriptions had not been taken into account by participants, then the best fitting parameters for the PVL_e model from the E conditions should predict behavior in the DE conditions well, which was not the case. Even the re-fitted experience-only model PVL_{de} did not provide a much improved fit: this might have been a result of over-fitting, as shown by the generalization analysis of the models against the IGT tasks in Experiments 1 and 2.

General discussion

Previous research has shown that introducing congruent descriptions to decisions from experience did not influence behavior (Lejarraga & Gonzalez, 2011; Weiss-Cohen et al., 2016). However we believe that this was because the tasks used previously were relatively simple, and that descriptions might only be taken into consideration by individuals when it is advantageous to do so, given the higher cognitive cost associated with processing them, compared to the easier processing of experiential information (Glöckner et al., 2012; Lejarraga, 2010). The aim of our current study was to show that congruent descriptions influence behavior in more complex tasks, where the addition of descriptions is advantageous given the higher cognitive effort required to decipher the task. The presence of descriptions helped participants perform better in the described DIGT in Experiments 1 and 2, enabling them to choose the advantageous decks more often and sooner. In our current research, given the complexity of the task, it was cognitively advantageous for participants to use descriptions, which in turn influenced behavior. This did not occur in previous research, where the task was relatively simple and relying on experience alone was sufficient.

In Experiment 3, we showed that the influence of descriptions on decisions from experience is moderated by task complexity. Descriptions helped participants' performance the most in tasks of medium complexity, where the experience is relatively too complex to be learned easily and efficiently, but descriptions are still relatively simple and can be processed without too much additional effort. When the task was very simple, participants were able to learn about the task experientially, which requires lower cognitive effort than analyzing the descriptions. Participants mostly seemed to have neglected descriptions, replicating results observed in previous research conducted with similarly simple tasks (Lejarraga & Gonzalez, 2011; Weiss-Cohen et al., 2016). Increasing the task complexity too much, however, led to a situation in which both experience and descriptions were overly complex. Learning via experience in complex tasks is more difficult, but processing the complicated written information required to describe such a complex task is also taxing and demanding. In these very complex tasks, the addition of descriptions did not help participants' performance. This created a non-monotonic inverted U-shaped pattern for the relationship between task complexity and influence of descriptions: highest for tasks of middle complexity, and lower in both extremes of low and high complexity (see Figure 7). Similarly shaped relationships between task complexity and decision performance had been observed before in other domains (Eppler & Mengis, 2004; Hwang & Lin, 1999; Streufert & Driver, 1965). They suggest that too much information can lead to cognitive overload, with deleterious influences on performance, similar to what was observed here. A limitation of our current experiments is that descriptions were always present on screen and constantly available to participants, which might have contributed to information overload. As a future avenue of research, the timing and frequency of the presence of descriptions should be manipulated.

A cognitive model that combined representations of both descriptive and experiential information was also fitted to the behavioral data in Experiment 3. The combined description-plus-experience model provided better fitting results than a more traditional experience-only model that relied on experience alone, and did not consider any additional descriptive information, in particular in the more complex experimental conditions. In the simpler conditions of Experiment 3, the combined model was no better than the traditional model, a result that was previously shown both in Lejarraga and Gonzalez (2011) and Weiss-Cohen et al. (2016), where an experience-only model fit the behavioral data relatively well for simple tasks. We observed that as task complexity increased, the addition of descriptive information into the model led to better fitting results. The combined description-plus-experience model also returned the same inverted U-shaped pattern for the relationship between task complexity and performance improvement due to the addition of descriptions, as observed behaviorally. The description-plus-experience combined model fitted against Experiment 3 also provided good predictions for the behavior observed in the DE conditions of Experiments 1 and 2 in a generalization test. An analysis of the best fitting parameters showed that individuals pay more attention to losses, learn more slowly and choose less randomly when descriptions are available to them. The proxy for complexity used in the model was entropy, and while an overall task entropy was envisaged initially, the best fitting model resulted from the use of an average entropy for each alternative, which varied mostly due to the change in the number of outcomes. We believe this to be the case because traditional reinforcement learning models using partial-feedback already indirectly take into account the number of alternatives available, since only one can be updated at a time.

Continuous exploration in the presence of descriptions

The traditional experience-only IGT is a task based on trial and error, as participants explore the options available to them to gather information about the composition of the cards within each deck, thereby reducing the uncertainty about the unknown characteristics of the task. By adding additional information in the form of descriptions to the task, the need for exploration is considerably reduced. An ideal agent who wanted to maximize their financial gains in our task could analyze the descriptive information available at the beginning of the task, decide which is their preferred alternative, and exploit that option across all trials, thus avoiding costly exploration. This was not observed and, while participants switched less often between decks when descriptions were present, they did not completely eliminate the switching behavior. This exploratory behavior, even after a considerable amount of information had been collected about the task, has been observed before, with many underlying explanations proposed.

Exploration is well adapted to dynamic environments, where full exploitation can leave the individual ignorant of changes in the reward structures of non-explored options (Knox, Otto, Stone, & Love, 2012; Speekenbrink & Konstantinidis, 2015). In such dynamic situations, descriptions are actually detrimental (Rakow & Miler, 2009). While our current task was static and not dynamic, perhaps human cognitive mechanisms for gathering information are better adapted to dynamic situations, which are more often encountered in real life, and not to static environments, which normally characterize artificial laboratory experiments. Shanks, Tunney, and McCarthy (2002) proposed that this sub-optimal type of exploratory behavior could be due to boredom, as participants did not want to choose the same option repeatedly, despite of the costs associated with diverging from optimal behavior. Alternatively, participants might have been selecting a mixed strategy where their preferred selection pattern was to diversify across different decks with a certain frequency, rather than having a single favorite option (Ashby et al., 2017; Konstantinidis et al., 2015). In decisions from description-plus-experience, exploration might be partly a result of participants' need to confirm the veracity of descriptions via direct personal experience. Perhaps this was driven by limited trust in the descriptions, and further research is necessary to establish the relationship between the influence of descriptions and the trustworthiness of their sources.

Implications for research on warnings

This research has important implications in the field of creating more effective warnings. Warnings can be seen as descriptions introduced to well-experienced situations. And while extensive research on warnings has investigated their physical characteristics such as color, placement, shape, size, tone, and symbology, very limited research has looked at the interaction between warnings and personal experience, and even those that did focused on familiarity and prior, not concurrent or posterior, experiences (for reviews, see Argo & Main, 2004; Rogers, Lamson, & Rousseau, 2000). We propose that the lack of usefulness of descriptions in overly complicated tasks might be due to the complexity of the descriptions themselves which are used to describe these tasks, making them hard to interpret. Perhaps if simpler descriptions could have been provided, even in complex tasks, then these would have influenced behavior more. Analogously, simpler warnings should be more efficient and increase compliance, while in reality it seems that warnings are becoming longer and more complex over time, as shown by the example of patient information leaflets distributed with medications which provide too much information that can be difficult to understand (Bandesha, Raynor, & Teale, 1996; Bradley, McCusker, Scott, & Li Wan Po, 1995). And while most real-life tasks are considerably more complex than the experiments presented here, it might be that individuals perceive certain tasks to be simpler than they are, perhaps by habituation, reducing their acceptance for and compliance with additional information in the form of descriptions or warnings. Further research on description-plus-experience should investigate how complexity influences the effectiveness of warnings in the form of descriptions interacting with experiences, and attempt to manipulate the two dimensions separately, for example by introducing simple descriptions into complex tasks and vice-versa. Description-plus-experience paradigms, such as the DIGT employed here, are useful tools for measuring how much impact descriptions (and by extension warnings) have on behavior, while assessing individual differences and their moderators.

Appendix A

Analyses for individual decks

Most of the research on the Iowa Gambling Task typically combines the two "advantageous" decks together (decks C+D) into a single metric for analysis, according to their positive long-term outcome potential, but disregarding individual deck characteristics (see Steingroever et al., 2013, for a review). The same approach was used here in our studies. However this hides any variation of choices among the individual decks. IGT studies show that participants tend to prefer the decks with less frequent losses, regardless of long-term rewards ("frequency-of-losses effect": Ahn et al., 2008). A consequence of this effect is that they tend to select from deck B more frequently than expected, given it has negative EV, the highest volatility, and the largest losses ("prominent deck B effect": Lin, Chiu, Lee, & Hsieh, 2007), and select from deck C less frequently than expected, given it has positive EV, the lowest volatility, and the smallest losses ("sunken deck C effect": Chiu & Lin, 2007). In this appendix, we have examined whether the same effects were present in our studies by re-analyzing the data at the level of individual decks and investigated how the addition of descriptions might have disrupted these earlier findings (Figure A1).

We re-analyzed the data from Experiments 1 and 2, together with the 4-deck $\times$ 4-outcome conditions of Experiment 3. These conditions were selected to make them comparable, as the total of 16 outcomes in the latter conditions closely resembles the former two, both with 14 outcomes each. The choice proportions of selections for each of the four decks within each sequential block of 20 trials were analyzed in four individual linear mixed-effects models. The fixed effects were the presence or absence of descriptions (E or DE), the blocks of 20 trials each, and their interaction. The models also contained a random intercept for each participant. A single model combining all decks was also generated for pairwise comparisons between decks. Post-hoc pairwise comparisons were Tukey adjusted.

In the experience-only conditions, the deck-level findings typically associated with IGT were also observed (Figure A1, rightmost panels), such as the prominent deck B effect, with participants selecting from deck B significantly more than all other decks (pairwise ps$<.006$). We also observed the sunken deck C effect, with participants selecting from deck C as frequently as from deck D (pairwise p=.11).

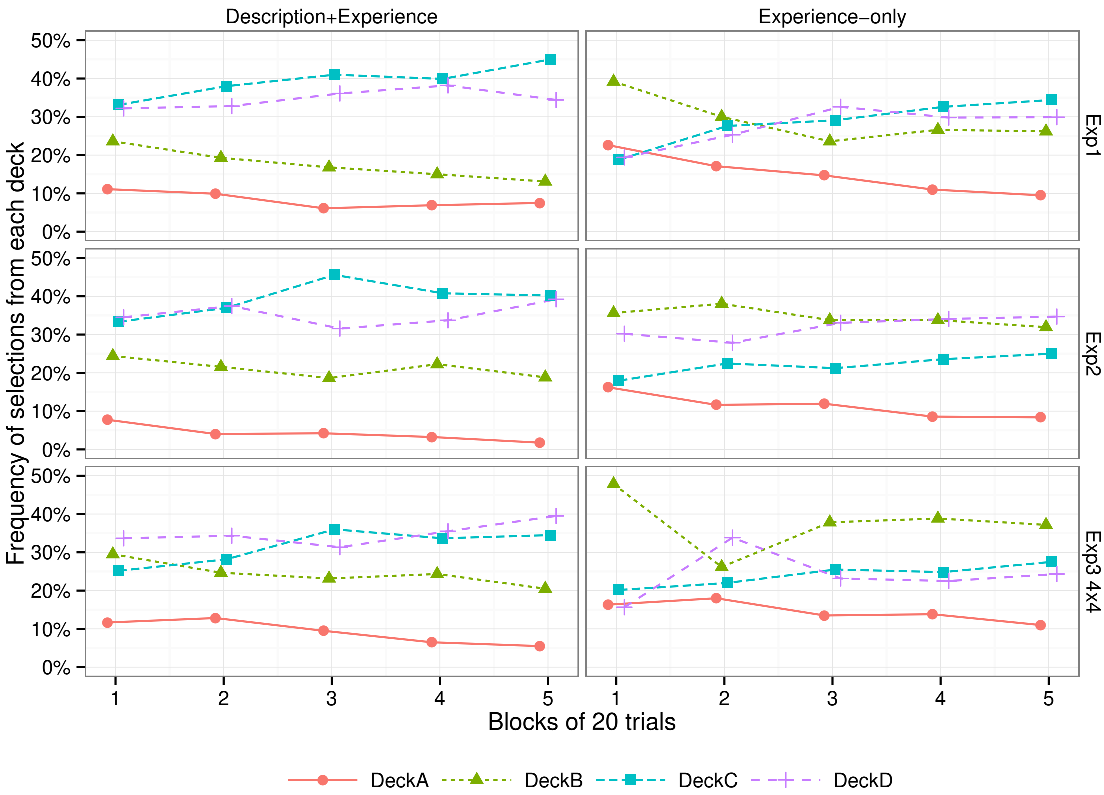

Figure A1. Evolution of the average frequency of selection from each individual deck, as a percentage of total for each block of 20 trials, for each experiment. The schedule of outcomes was fixed in Experiment 1 and pseudo-randomized in Experiments 2 and 3. Experiment 3 only shows the comparable 4-deck × 4-outcome condition. Decks A and B were disadvantageous with negative EV and decks C and D were advantageous with positive EV. Decks B and D had frequent wins and infrequent (but larger) losses, while decks A and C were more balanced with 50% wins and 50% losses.

The presence of descriptions creates a much clearer separation between selections from the advantageous and disadvantageous decks (Figure A1, leftmost panels). Descriptions helped participants identify deck A as the most disadvantageous sooner, and significantly reduced the overall selections from that deck (Means and 95% C.I. shown: E=13.5% [12.1,14.9], DE=6.9% [5.5,8.3], $t(258)=6.48$, $p<.001$, $d=0.80$). Descriptions also greatly reduced selections from deck B (E=33.2% [29.6,36.7], DE=20.5% [16.7,24.1], $t(258)=4.96$, $p<.001$, $d=0.62$), which were no longer the most often selected deck, thus eliminating the "prominent deck B effect". Despite an increase in selections from deck C (E=25.0% [20.1,29.9], DE=37.6% [32.7,42.5], $t(258)=3.58$, $p<.001$, $d=0.44$), the effect size was smaller, and the “sunken deck C effect” still remained with descriptions. This effect was exacerbated by the even smaller influence of descriptions on deck D (E=28.3% [24.0,30.7], DE=35.0% [30.7,39.3], $t(258)=2.15$, $p=.03$, $d=0.27$), with a small effect size. It seems that participants are still affected by the frequency-of-losses effect, which would lead them to select deck D more often than expected, because of its low frequency of (large) losses, in the end selecting relatively equally often from the two advantageous decks ($p=.29$). The main effect of block was significant for all decks, in the direction expected, with selections increasing over time for the two advantageous decks, and decreasing over time for the two disadvantageous decks (all ps $< .005$), but not influenced by the presence of descriptions, with non-significant interactions (all ps $\geq .18$).

Overall the main contributions of descriptions at individual deck level were to strongly reduce the “prominent deck B” effect, as deck B was no longer the most often selected deck, and help participants identify the disadvantageous deck A, with relatively smaller increases in selections from decks C and D, which sustained a lingering “sunken deck C” effect.

Appendix B

Supplementary material

The data from all studies have been made publicly available for download via the Open Science Framework and can be accessed at https://osf.io/rzadn/.

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. W. (2011). A model-based fMRI analysis with hierarchical Bayesian parameter estimation. *Journal of Neuroscience, Psychology, and Economics*, *4*, 95–110. doi: 10.1037/a0020684

Argo, J. J., & Main, K. J. (2004). Meta-analyses of the effectiveness of warning labels. *Journal of Public Policy & Marketing*, *23*, 193–208. doi: 10.1509/jppm.23.2.193.51400

Ashby, N. J., Konstantinidis, E., & Yechiam, E. (2017). Choice in experiential learning: True preferences or experimental artifacts? *Acta Psychologica*, *174*, 59–67. doi: 10.1016/j.actpsy.2017.01.010

Bandesha, G., Raynor, D. K., & Teale, C. (1996). Preliminary investigation of patient information leaflets as package inserts. *International Journal of Pharmacy Practice*, *4*, 246–248. doi: 10.1111/j.2042-7174.1996.tb00875.x

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, *16*, 215–233. doi: 10.1002/bdm.443

Barron, G., Leider, S., & Stack, J. (2008). The effect of safe experience on a warnings' impact: Sex, drugs, and rock-n-roll. *Organizational Behavior and Human Decision Processes*, *106*, 125–142. doi: 10.1016/j.obhdp.2007.11.002

Bates, D., Maechler, M., Bolker, B., & Walker, S. (2014). *lme4: Linear mixed-effects models using Eigen and S4.* R package version 1.1-6. Retrieved from `http://cran.r-project.org/package=lme4`

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

Bechara, A., Damasio, H., Damasio, A. R., & Lee, G. P. (1999). Different contributions of the human amygdala and ventromedial prefrontal cortex to decision-making. *The Journal of Neuroscience*, *19*, 5473–5481.

Bechara, A., Damasio, H., Tranel, D., & Damasio, A. R. (1997). Deciding advantageously before knowing the advantageous strategy. *Science*, *275*, 1293–1295. doi: 10.1126/science.275.5304.1293

Bradley, B., McCusker, E., Scott, E., & Li Wan Po, A. (1995). Patient information leaflets on over-the-counter (OTC) medicines: The manufacturer's perspective. *Journal of Clinical Pharmacy and Therapeutics*, *20*, 37–40. doi: 10.1111/j.1365-2710.1995.tb00623.x

Briscoe, E., & Feldman, J. (2011). Conceptual complexity and the bias/variance tradeoff. *Cognition*, *118*, 2–16. doi: 10.1016/j.cognition.2010.10.004

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., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, *100*, 432–459. doi: 10.1037/0033-295X.100.3.432

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

Camilleri, A. R., & Newell, B. R. (2011). Description- and experience-based choice: Does equivalent information equal equivalent choice? *Acta Psychologica*, *136*, 276–284. doi: 10.1016/j.actpsy.2010.11.007

Camilleri, A. R., & Newell, B. R. (2013a). The long and short of it: Closing the description-experience "gap" by taking the long-run view. *Cognition*, *126*, 54–71. doi: 10.1016/j.cognition.2012.09.001

Camilleri, A. R., & Newell, B. R. (2013b). Mind the gap? Description, experience, and the continuum of uncertainty in risky choice. In N. Srinivasan & P. Chandrasekhar (Eds.), *Progress in brain research* (pp. 55–72). Oxford, UK: Elsevier. doi: 10.1016/B978-0-444-62604-2.00004-6

Campbell, D. J. (1988). Task complexity: A review and analysis. *The Academy of Management Review*, *13*, 40–52. doi: 10.2307/258353

Chiu, Y.-C., & Lin, C.-H. (2007). Is deck C an advantageous deck in the Iowa Gambling Task? *Behavioral and Brain Functions*, *3*, 37. doi: 10.1186/1744-9081-3-37

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, *24*, 87–185. doi: 10.1017/S0140525X01003922

Dai, J., Kerestes, R., Upton, D. J., Busemeyer, J. R., & Stout, J. C. (2015). An improved cognitive model of the Iowa and Soochow Gambling Tasks with regard to model fitting performance and tests of parameter consistency. *Frontiers in Psychology*, *6*, 229. doi: 10.3389/fpsyg.2015.00229

Eppler, M. J., & Mengis, J. (2004). The concept of information overload: A review of the literature from organization science, accounting, MIS, and related disciplines. *The Information Society*, *20*, 325–344. doi: 10.1080/01972240490507974

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., Ert, E., Plonsky, O., Cohen, D., & Cohen, O. (2017). From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological Review*, *124*, 369–409. doi: 10.1037/rev0000062

Erev, I., Glozman, I., & Hertwig, R. (2008). What impacts the impact of rare events. *Journal of Risk and Uncertainty*, *36*, 153–177. doi: 10.1007/s11166-008-9035-z

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

Ert, E., & Erev, I. (2007). Replicated alternatives and the role of confusion, chasing, and regret in decisions from experience. *Journal of Behavioral Decision Making*, *20*, 305–322. doi: 10.1002/bdm.556

Ert, E., Erev, I., & Roth, A. E. (2011). A choice prediction competition for social preferences in simple extensive form games: An introduction. *Games*, *2*, 257–276. doi: 10.3390/g2030257

Fantino, E., & Navarro, A. (2012). Description-experience gaps: Assessments in other choice paradigms. *Journal of Behavioral Decision Making*, *25*, 303–314. doi: 10.1002/bdm.737

Fasolo, B., Hertwig, R., Huber, M., & Ludwig, M. (2009). Size, entropy, and density: What is the difference that makes the difference between small and large real-world assortments? *Psychology & Marketing*, *26*, 254–279. doi: 10.1002/mar.20272

Feldman, J. (2000). Minimization of Boolean complexity in human concept learning. *Nature*, *407*, 630–633. doi: 10.1038/35036586

Frey, R., Mata, R., & Hertwig, R. (2015). The role of cognitive abilities in decisions from experience: Age differences emerge as a function of choice set size. *Cognition*, *142*, 60–80. doi: 10.1016/j.cognition.2015.05.004

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

Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review*, *102*, 684–704. doi: 10.1037/0033-295X.102.4.684

Glöckner, A., Fiedler, S., Hochman, G., Ayal, S., & Hilbig, B. E. (2012). Processing differences between descriptions and experience: A comparative analysis using eye-tracking and physiological measures. *Frontiers in Psychology*, *3*, 173. doi: 10.3389/fpsyg.2012.00173

Halford, G. S., Wilson, W. H., & Phillips, S. (1998). Processing capacity defined by relational complexity: Implications for comparative, developmental, and cognitive psychology. *Behavioral and Brain Sciences*, *21*, 803–864. doi: 10.1017/S0140525X98001769

Hasher, L., & Zacks, R. T. (1984). Automatic processing of fundamental information: The case of frequency of occurrence. *The American Psychologist*, *39*, 1372–1388. doi: 10.1037/0003-066X.39.12.1372

Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, *13*, 517–523. doi: 10.1016/j.tics.2009.09.004

Hills, T. T., Noguchi, T., & Gibbert, M. (2013). Information overload or search-amplified risk? Set size and order effects on decisions from experience. *Psychonomic Bulletin & Review*, *20*, 1023–1031. doi: 10.3758/s13423-013-0422-3

Hwang, M. I., & Lin, J. W. (1999). Information dimension, information overload and decision quality. *Journal of Information Science*, *25*, 213–218. doi: 10.1177/0165551994232684

Jacoby, J., Speller, D. E., & Berning, C. K. (1974). Brand choice behavior as a function of information load: Replication and extension. *Journal of Consumer Research*, *1*, 33–42. doi: 10.1086/208579

Jacoby, J., Speller, D. E., & Kohn, C. A. (1974). Brand choice behavior as a function of information load. *Journal of Marketing Research*, *11*, 63–69. doi: 10.2307/3150994

Jessup, R. K., Bishara, A. J., & Busemeyer, J. R. (2008). Feedback produces divergence from prospect theory in descriptive choice. *Psychological Science*, *19*, 1015–1022. doi: 10.1111/j.1467-9280.2008.02193.x

Johnson, E. J., & Payne, J. W. (1985). Effort and accuracy in choice. *Management Science*, *31*, 395–414. doi: 10.1287/mnsc.31.4.395

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–292. doi: 10.2307/1914185

Keller, K. L., & Staelin, R. (1987). Effects of quality and quantity of information on decision effectiveness. *Journal of Consumer Research*, *14*, 200–213. doi: 10.1086/209106

Knox, W. B., Otto, A. R., Stone, P., & Love, B. C. (2012). The nature of belief-directed exploratory choice in human decision-making. *Frontiers in Psychology*, *2*, 398. doi: 10.3389/fpsyg.2011.00398

Konstantinidis, E., Ashby, N. J., & Gonzalez, C. (2015). Exploring complexity in decisions from experience: Same minds, same strategy. In D. C. Noelle et al. (Eds.), *Proceedings of the 37th annual conference of the cognitive science society* (pp. 1177–1182). Austin, TX: Cognitive Science Society.

Konstantinidis, E., Speekenbrink, M., Stout, J. C., Ahn, W.-Y., & Shanks, D. R. (2014). To simulate or not? Comment on Steingroever, Wetzels, and Wagenmakers (2014). *Decision*, *1*, 184–191. doi: 10.1037/dec0000016

Lejarraga, T. (2010). When experience is better than description: Time delays and complexity. *Journal of Behavioral Decision Making*, *23*, 100–116. doi: 10.1002/bdm.666

Lejarraga, T., & Gonzalez, C. (2011). Effects of feedback and complexity on repeated decisions from description. *Organizational Behavior and Human Decision Processes*, *116*, 286–295. doi: 10.1016/j.obhdp.2011.05.001

Lenth, R. V. (2016). Least-squares means: The R package lsmeans. *Journal of Statistical Software*, *69*, 1–33. doi: 10.18637/jss.v069.i01

Lewandowsky, S., & Farrell, S. (2011). *Computational modeling in cognition: Principles and practice.* Thousand Oaks, CA: SAGE Publications.

Lin, C.-H., Chiu, Y.-C., Lee, P.-L., & Hsieh, J.-C. (2007). Is deck B a disadvantageous deck in the Iowa Gambling Task? *Behavioral and Brain Functions*, *3*, 16. doi: 10.1186/1744-9081-3-16

Mathy, F., & Bradmetz, J. (2004). A theory of the graceful complexification of concepts and their learnability. *Current Psychology of Cognition*, *22*, 41–82.

Mathy, F., & Feldman, J. (2012). What's magic about magic numbers? Chunking and data compression in short-term memory. *Cognition*, *122*, 346–362. doi: 10.1016/j.cognition.2011.11.003

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, *63*, 81–97. doi: 10.1037/h0043158

Noguchi, T., & Hills, T. T. (2016). Experience-based decisions favor riskier alternatives in large sets. *Journal of Behavioral Decision Making*, *29*, 489–498. doi: 10.1002/bdm.1893

Paul, S., & Nazareth, D. L. (2010). Input information complexity, perceived time pressure, and information processing in GSS-based work groups: An experimental investigation using a decision schema to alleviate information overload conditions. *Decision Support Systems*, *49*, 31–40. doi: 10.1016/j.dss.2009.12.007

R Core Team. (2014). *R: A language and environment for statistical computing.* R Foundation for Statistical Computing, Vienna, Austria. Retrieved from `http://www.r-project.org/`

Rakow, T., & Miler, K. (2009). Doomed to repeat the successes of the past: History is best forgotten for repeated choices with nonstationary payoffs. *Memory & Cognition*, *37*, 985–1000. doi: 10.3758/MC.37.7.985

Rakow, T., & Newell, B. R. (2010). Degrees of uncertainty: An overview and framework for future research on experience-based choice. *Journal of Behavioral Decision Making*, *23*, 1–14. doi: 10.1002/bdm.681

Reutskaja, E., & Hogarth, R. M. (2009). Satisfaction in choice as a function of the number of alternatives: When "goods satiate". *Psychology & Marketing*, *26*, 197–203. doi: 10.1002/mar.20268

Rogers, W. A., Lamson, N., & Rousseau, G. K. (2000). Warning research: An integrative perspective. *Human Factors*, *42*, 102–139. doi: 10.1518/001872000779656624

Shah, A. M., & Wolford, G. (2007). Buying behavior as a function of parametric variation of number of choices. *Psychological Science*, *18*, 369–370. doi: 10.1111/j.1467-9280.2007.01906.x

Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, *15*, 233–250. doi: 10.1002/bdm.413

Shiffrin, R., 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

Shlomi, Y. (2014). Subjective integration of probabilistic information from experience and description. *Judgment and Decision Making*, *9*, 491–499.

Speekenbrink, M., & Konstantinidis, E. (2015). Uncertainty and exploration in a restless bandit problem. *Topics in Cognitive Science*, *7*, 351–367. doi: 10.1111/tops.12145

Steingroever, H., Wetzels, R., Horstmann, A., Neumann, J., & Wagenmakers, E.-J. (2013). Performance of healthy participants on the Iowa Gambling Task. *Psychological Assessment*, *25*, 180–193. doi: 10.1037/a0029929

Streufert, S., & Driver, M. J. (1965). Conceptual structure, information load and perceptual complexity. *Psychonomic Science*, *3*, 249–250. doi: 10.3758/BF03343120

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

Swait, J., & Adamowicz, W. (2001). The influence of task complexity on consumer choice: A latent class model of decision strategy switching. *Journal of Consumer Research*, *28*, 135–148. doi: 10.1086/321952

Thorngate, W. (1980). Efficient decision heuristics. *Behavioral Science*, *25*, 219–225. doi: 10.1002/bs.3830250306

Tversky, A., & Kahneman, D. (1992). Advances in Prospect-Theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323. doi: 10.1007/Bf00122574

Wagenmakers, E.-J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, *11*, 192–196. doi: 10.3758/BF03206482

Weiss-Cohen, L., Konstantinidis, E., Speekenbrink, M., & Harvey, N. (2016). Incorporating conflicting descriptions into decisions from experience. *Organizational Behavior and Human Decision Processes*, *135*, 55–69. doi: 10.1016/j.obhdp.2016.05.005

Wood, R. E. (1986). Task complexity: Definition of the construct. *Organizational Behavior and Human Decision Processes*, *37*, 60–82. doi: 10.1016/0749-5978(86)90044-0

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. (2005). Comparison of basic assumptions embedded in learning models for experience-based decision making. *Psychonomic Bulletin & Review*, *12*, 387–402. doi: 10.3758/BF03193783

Yechiam, E., & Busemeyer, J. R. (2006). The effect of foregone payoffs on underweighting small probability events. *Journal of Behavioral Decision Making*, *19*, 1–16. doi: 10.1002/bdm.509

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., & 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

Yuviler-Gavish, N., & Gopher, D. (2011). Effect of descriptive information and experience on automation reliance. *Human Factors*, *53*, 230–244. doi: 10.1177/0018720811406725