

## RISK ASSESSMENT UNDER PERCEPTUAL AMBIGUITY AND ITS IMPACT ON CATEGORY LEARNING

Rubi Hammer

Interdepartmental Neuroscience Program, Northwestern University, Evanston, IL

Moran Cerf

MIT Media Lab, MIT, Cambridge, MA

**Proper item categorization often requires prolonged learning of the value (gain or loss probabilities) associated with different items. We tested individual differences in learning strategies where the need to discriminate between two perceptually similar categories co-occurred with the need to learn the mean value of items within each category. In the primary experimental condition, participants had to learn that items in one category were always associated with a positive monetary value (+7¢ in 100% of the trials), whereas items in another, “risky” category were associated either with a larger monetary gain (+10¢ in 75% of the trials) or with a substantial loss (-27¢ in 25% of the trials). We found that even after hundreds of learning trials, participants often preferred an item from the second, risky category over an item from the always-positive-value category, despite a respective mean-value ratio of approximately 1:9 between the two categories. On the other hand, in the two other experimental conditions, where the value of the items in one category was always higher than the value of the items in the other category, participants quickly learned to identify which category of items had greater mean value, and preferred those items, as expected. The substantial overvaluing of items from a risky category (with lower mean item-value) over the always-positive-value category (with much higher mean item-value) indicates that under conditions that involve risk-taking and the transfer of knowledge between tasks, experience-based categorization decisions may substantially deviate from Bayesian inference driven by experience.**

**Keywords:** Category learning; Payoff variability; Perceptual expertise; Prospect theory; Risk assessment; Reinforcement learning.

### INTRODUCTION

People are often confronted with ambiguous sensory input and uncertainty regarding the course of actions that may maximize utility. For example, a novice mushroom hunter who is searching the woods for high-quality *Russula Xerampelina* mushrooms may find this task to be challenging. The *Xerampelina* is characterized by a wine-red color that differentiates it from the lower-quality, yet edible, red-brown *Russula Cyanoxantha* mushroom. However, nature is not without a sense of irony, as there is a third subtype of mushroom, the red *Russula Emetica*, that is nearly identical to the desirable *Xerampelina*, but is poisonous. Only by tasting the mushroom can one decisively determine if it is a *Xerampelina* or an *Emetica*. In the latter case, tasting the mushroom would be an aversive experience. If half (4/8) of the mushrooms in the forest are *Cyanoxantha*, 3/8 are *Xerampelina*, and 1/8 are

*Emetica*, we can posit that half of the mushrooms are “safe” red-brown mushrooms with mild flavor, whereas the remaining are “risky” (either high-quality or poisonous) red mushrooms. This example raises a few fundamental questions about the learning and decision-making processes under perceptual ambiguity: (i) How effective would a novice mushroom hunter, who does not know these priors, be in learning the differences between the mushroom subtypes? (ii) Would persistent learning by trial and error (which would involve continuous risk-taking) be a more favorable strategy than a conservative one, which would involve gathering only the safer, distinctly red-brown mushrooms? (iii) As the learner gains substantial experience with the three mushroom subtypes, will the respective preference for red mushrooms over red-brown mushrooms gradually increase or decrease?

The above exemplifies many real-life scenarios of exploratory learning by trial and error, where categories of interest are a-priori not clearly defined by their perceptual characteristics and the value of the items within each category. In such scenarios, there are two interrelated and possibly conflicting challenges: the first is to become better at telling apart items from perceptually similar categories by using the items' value (e.g., edibility/flavor) as supervisory information; the second challenge is to become better at making predictions regarding the value of an item (i.e., is it edible and tasty) based on specific perceptual features (e.g., shape and color). When learning about a category of risky items, the lack of a consistent category-to-value mapping poses a distinct challenge. This is unlike scenarios where all items with given perceptual attributes are associated with an identical, or similar, value. Risky items have two (or more) fundamentally different values with which visible attributes can be associated (e.g., red mushrooms that may either have excellent flavor or be poisonous). Currently, little is known about category learning and decision-making strategies in scenarios where perceptual differences between categories are relatively hard to detect and where there is additional uncertainty about the expected utility of categorization decisions. The current study aims to characterize the extent of individual differences in learning and decision-making in such scenarios.

The prospect theory and later related studies provide useful insights about decision-making where people are informed about the a-priori (descriptive) probabilities of different outcomes. These studies show that when informed about the a-priori probabilities, people tend to overestimate the impact of losses in respect to equivalent gain (Benartzi & Thaler, 1995; Kahneman & Tversky, 1979; Tversky & Kahneman, 1992). However, as in the mushroom example, various scenarios offer probabilities of negative outcomes that are not readily available and instead require gradual learning based on experience. Indeed, earlier studies show that experience-based decisions often have limited correspondence with decisions that are based on explicit a-priori probabilities. Specifically, it has been shown that people are less averse to losses and may even prefer risky option when decisions are made based on subjective experience (Barron & Erev, 2003; Barron & Ursino, 2013; Erev & Barron, 2005; Erev, Ert & Yechiam, 2008; Jessup, Bishara & Busemeyer, 2008; Weber, Shafir & Blais, 2004; see Hertwig & Erev, 2009 for a review). This latter kind of deviation from an optimal decision-making strategy may occur due to the payoff-variability effect. This effect characterizes scenarios in which the same action (e.g., categorizing red mushrooms as likely to be tasty) may result in substantially distinct outcomes. Outcomes are herein often measured as gains or losses.

In such scenarios, the individual's assessment of the mean utility of each type of decision may be compromised due to the increased difficulty in logging to the subjective history of gains and losses (and the associated perceptual information) following multiple previous decisions and actions. Therefore, each decision is likely to be biased by the more recent subjective history (or perhaps a subset of events that left a stronger impression). As the probability of negative outcomes decreases, the likelihood of risk-taking behavior increases. This may be true even if the negative outcomes may involve extreme losses (Myers & Sadler, 1960).

Moreover, people may often fail to detect which past events truly resemble an event presently being assessed (Erev & Roth, 2014; Lynn et al., 2015). We suggest that in some scenarios, this may be due to perceptual ambiguity or high cognitive load that may result in a difficulty to process part of the task-relevant information. For example, a real-estate agent may fail to determine if buying a given property would yield future profit, not because of a poor capacity to trace past gains and losses, but due to a limited ability to account for all the factors impacting a given real-estate market. This would compromise the capability to confidently classify a given property as being a safe or a risky investment. Similarly, radiologists, radar operators or product quality-control officers are required to make critical choices under perceptual ambiguity and/or high cognitive load, often while also having to learn from their experience the utility coupled with their different possible choices (for related discussions, see Brooks, Norman & Allen (1991), Hammer (2015) and Palmeri, Wong & Gauthier (2004)).

There is much value in studying learning processes and decision-making dynamics in scenarios where the ability to differentiate between risky and non-risky items involve high cognitive or high perceptual load, specifically when this ability is acquired alongside the learning of the probabilities of respective utilities (for related suggestions, see Mehlhorn et al., 2014 and Schonberg, Fox & Poldrack, 2011). We primarily wanted to assess the degree to which perceptually challenging category learning is compromised by decision-making biases associated with risk assessment. We also consider the possibility that the two cognitive challenges (perceptual discrimination and item value assessment) have mutual impact. That is, where perceptual challenges may also compromise the capacity to assign a specific value to attributes of the categorized items that are most relevant for categorization.

Many of the refeed above studies reported substantial individual differences in people loss aversion. Although these studies did not primarily account to mechanisms

underlying the observed individual differences, other studies imply that such variability in loss aversion may be driven by both executive and lower-level procedural and 'affective' processes. For example, it has been shown that the individual loss aversion can be predicted by frontal (executive functions), ventral striatum (procedural learning) and limbic (affective processes) fMRI responses, and by gray matter volume in a structural amygdala–thalamus–striatum network (Canessa et al., 2013; Tom et al., 2007). Some behavioral studies show that variability in loss aversion can be reduced when subjects are trained to diminish affective processes (e.g., Sokol-Hessner et al., 2009). Currently there is no clear understanding of the exact role of each of those brain networks in loss aversion. However, it is possible that learning to categorize risky items is an exceptionally more complex process since it requires a proper exchange of information in an elaborated brain network. Some brain modules in this network are associated with the assessment of positive events, or at least value-neutral events, whereas others are primarily associated with the processing of aversive events. Thus, scenarios that are characterized by non-monotonic mapping between the item value and the item perceptual features that are relevant for categorization may require greater information exchange between several brain networks for processing the same category of items. Such non-monotonic feature-value mapping characterizes scenarios that involve risk, as exemplified above—mushrooms with shades of red are associated with either high positive values (great taste) or negative values (piousness), whereas mushrooms with shades of red-brown are associated with intermediate positive values (edible, but not tasty). This increases the complexity in forming a decisive categorical decision when one is trying to answer the question “which color of mushrooms I should like most?”. That is, risk assessment relies on non-monotonic feature-value mapping that may require a more elaborated brain network than otherwise required for simpler reinforcement learning. We suggest that this increased processing complexity may be associated with greater difficulty to transfer experience from one task to another, and with greater individual differences.

Here we employed a novel experimental paradigm to test how people learn to differentiate items from two visually similar categories that differ from one another in the monetary value of the items they encompass. Such a scenario is with great interest since it exemplifies a complex interaction between lower-level perceptual processes, procedural processes and higher-level executive processes associated with the long-term assessment of gain and loss (Hammer, Sloutsky & Grill-Spector, 2015). Following a learning phase, during which participants had to categorize one item at a time (and were provided with each item's monetary value as

feedback), we tested participants' decision-making in a task where they had to decide which of two items is expected to have a greater value. The participants' expectation in the decision-making task should have been based on their experience during the learning phase. We suggest that such scenarios, where one should learn the value of individual items prior to having to choose among several items are most common. In the primary experimental condition, the participants had to learn that items in one category were always associated with a fixed monetary gain (+7¢ in 100% of the trials), whereas items in the other category were associated with either a larger monetary gain (+10¢ in 75% of the trials) or an even larger but less frequent monetary loss (-27¢ in 25% of the trials). That is, the second object-category included risky items. We refer to this experimental condition as the Gain–Risk condition.

The payoff scheme in the Gain–Risk condition insured that the objective mean utility difference between the two categories was substantial, and the occurrence of negative outcomes was not too rare (1:3 negative/positive outcome frequency ratio). It was accepted that this should have made it possible for participants to make a reasonable assessment of each category value, even if expecting human learning and decision-making to somewhat deviate from an optimal strategy. That is, by employing a rational (even if imperfect) learning and decision-making strategy, this payoff scheme should have resulted in an aversion towards the risky category.

Performance in the Gain–Risk condition was compared with two experimental conditions that did not involve risk-taking. However, in those other two experimental conditions (see Methods for details), as well, one category involved a complex category-to-value mapping where items in one category had two distinct values, as it was in the Gain–Risk experimental condition. However, unlike the Gain–Risk condition, here the category-to-value mapping was monotonic.

In one of those two conditions (referred as the Gain–Smaller-Gain condition) the participant had to learn to choose between a category of items with larger positive value (+7¢ in 100% of the trials) and a category of items with smaller mean positive value (+3¢ in 75% of the trials and 0¢ in 25% of the trials). This condition simulated a scenario where the mean item value of the two categories was the same as in the Gain–Risk condition, but without involving risk taking (there was no money loss associated with any item). In the third condition (referred as the Gain–Loss condition) the participant had to learn to choose between a category of items with a positive value (+7¢ in 100% of the trials) and a category of items with a negative value (-9¢ in 75% of the trials and 0¢ in 25%

of the trials). This condition simulated a scenario in which the preference to the first category over the other is expected to be most obvious.

There are no early studies who tested scenarios like those tested here and we initially had several hypotheses regarding what would be the most prevalent behavior to be observed given the current setup. Thus, this should be considered an exploratory study intended at providing an initial valuation of risk assessment under perceptual ambiguity and its impact on category learning. Nevertheless, as discussed above, we expected that at least for some people risk taking scenarios will be likely to be associated with compromised cognitive processing, more than scenarios that do not involve risk. Accordingly, we hypothesized the following possible outcomes: (1) Given the substantial difference between the mean item value of the two categories in all three experimental conditions, most participants will show similar learning capacity in all three conditions. This will be evident in a good capacity to differentiate between the two categories and in a clear preference towards items with greater mean utility (i.e., employing a close-to-optimal categorization, where the mean value of an item determines its category identity). (2) In the Gain–Risk condition, people may exhibit greater idiosyncrasy even under conditions in which the computation (or at least a gross yet sufficiently good assessment) of the mean utility should be feasible. This would be due to greater difficulty to trace the value of multiple highly-similar exemplars from each category, in scenarios where items from one of the two categories have distinct contrasting values (non-monotonic feature-value mapping). In such a case, compromised learning may be evident as: (i) a limited capacity to distinguish between perceptually confusable items across the two categories following the learning phase; and/or (ii) compromised capacity to assess the items’ value and the risk associated with each item category (due to a poor assessment of the mean category utility). While (i) would be associated with suboptimal choices affecting mostly perceptually confusable items from both categories (items closer to the between-categories boundary), (ii) would impact decisions concerning all items, including perceptually distinct items from the two categories.

## METHODS

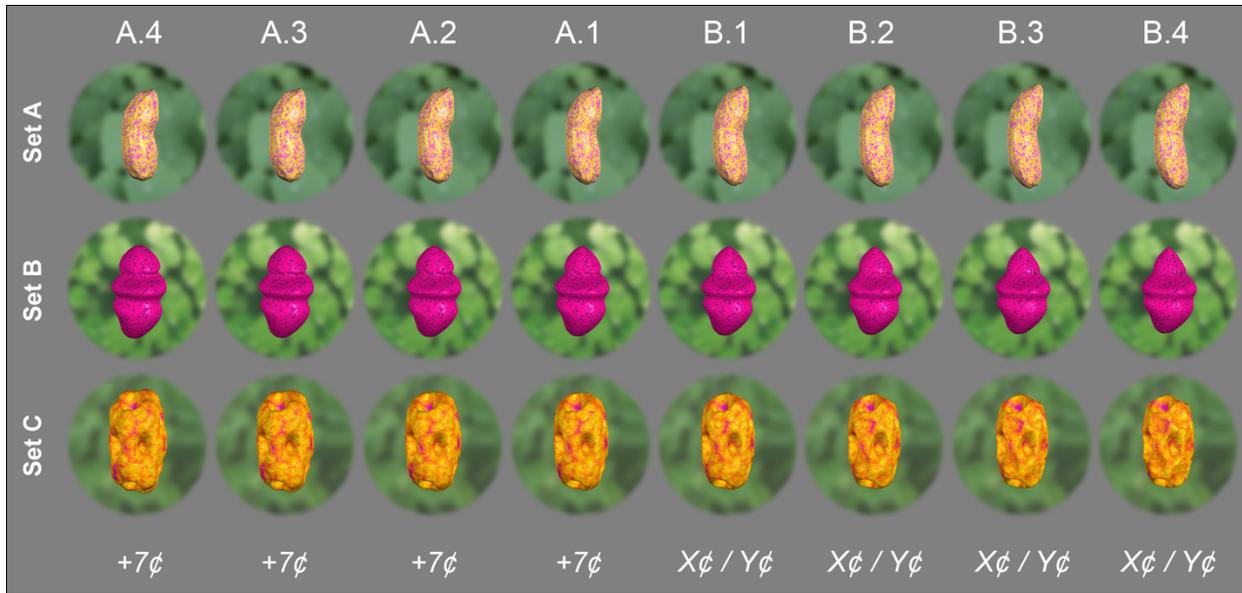
**Participants.** Twenty adults with normal or corrected-to-normal vision and with no history of neural or psychiatric disorders participated in the experiment. One participant was excluded from the analysis for not following instructions, and another was excluded due to poor performance in all three experimental conditions. The remaining 18 participants (10 females; ages 19–45; mean

$\pm$  standard deviation:  $27.9 \pm 5.6$  years) had a mean number of years of education of 19.4 (SD =  $\pm 4.0$ ; second-year undergraduate or greater). Participants were paid \$45–\$90 for their participation in the three experimental sessions (exact payment depended on task performance). This study was reviewed and approved by the Institutional Review Board (IRB) at Northwestern University and was conducted in accordance with the IRB policies. At the beginning of the first experimental session, participants were debriefed about the study, gave their consent and signed a consent form.

**Materials.** Multiple distinct stimulus sets of novel computer-generated images of “exotic fruits” with generic “organic” characteristics of real fruits were produced for this study. Each stimulus set included eight fruit-like stimuli, linearly morphed between two extreme exemplars (Figure 1). The multiple stimulus sets were required due to the within-subject experimental design, where each participant completed three Visual Category Learning (VCL) tasks. Each task was based on a distinct set of stimuli. Furthermore, before starting the experimental tasks, the participants were introduced to the experimental procedure by performing a training task, using a stimulus set that was not used in the experimental tasks. Stimulus sets were counterbalanced across participants and experimental conditions. Each set was used uniformly in one of the three conditions (counterbalanced across participants).

E-Prime ® 2.0 (Psychology Software Tools, Inc) was used for stimuli presentation and for recording participant responses. The experiment was run on a Lenovo™ 14” laptop with a 1600×900 pixels display. Each stimulus occupied 600×600 pixels. Stimuli were presented on a gray background. Participants were seated with their head within 60–90 centimeters from the computer monitor, and responded by using the “arrow” keys.

**Procedure.** Each participant completed three experimental sessions (a within-subject design). Each session was conducted on a different day. In each session, the participant performed a single VCL task in one of the three experimental conditions. The order of the conditions was counterbalanced across participants. In each of the three experimental conditions, the participant had to learn to distinguish between two categories of novel exotic fruits. In all three experimental conditions, one category of fruits, Category A, was characterized by a fixed item value of +7 cents (¢). The other category (Category B) had fruits that could have had one of two values, which differed among the three experimental conditions (see Table 1). In the Gain–Loss condition, a Category A item value was always +7¢, whereas the value of a Category B item was either -9¢ or 0¢. In the



**Figure 1.** Stimuli from three representative sets (labeled here as set A, B and C; five sets were generated for the study). Each set included eight stimuli, linearly morphed between two extremes (A.4 and B.4). In each set, the participants had to learn to differentiate between two categories (Category A with stimuli A.1 to A.4 and Category B with stimuli B.1 to B.4). Categories were characterized by a distinct monetary value pertaining to each category member items. In all conditions, the value of items in one category (referred to throughout the manuscript as Category A) was +7 cents, whereas the two values of items in the other category (Category B) were assigned likelihoods of 75% and 25% (see Table 1 for a breakdown of the items’ monetary value). In each stimulus set, the assignment of stimuli to Categories A or B was counterbalanced across participants.

Condition	Gain—Loss	Gain—Smaller-Gain	Gain—Risk
Category A	+7¢ (100%)	+7¢ (100%)	+7¢ (100%)
Category B	-9¢ (75%)	0¢ (75%)	+10¢ (75%)
	0¢ (25%)	+3¢ (25%)	-27¢ (25%)
Category B (average)	-6.75¢	+0.75¢	+0.75¢

**Table 1.** Item-value probabilities (in parentheses) in the three experimental conditions. The participants had the opportunity to learn these probabilities based on the item–value associations provided as feedback in the learning phase.

Gain—Smaller-Gain condition, a Category A item value was always +7¢, whereas the value of a Category B item was either 0¢ or +3¢. In the Gain—Risk condition, a Category A item value was always +7¢, whereas the value of a Category B item was either +10¢ or -27¢. Note that the mean value of Category B items in the Gain—Risk and Gain—Smaller-Gain conditions was identical (+0.75¢). The assignment of stimuli to Category A or Category B was counterbalanced across participants. Participants were not informed about the number of item categories that had to be learned in each task.

In each experimental session, there were five tasks: (i) pre-learning perceptual- discrimination task; (ii) pre-learning item-preference task; (iii) learning phase; (iv) post-learning perceptual-discrimination task; (v) post-learning item-preference task.

The pre- and post-learning perceptual-discrimination tasks were used to assess pre- to post-learning changes in perceptual-discrimination capabilities, regardless of the learned item preference based on monetary value. In each trial of the perceptual-discrimination tasks, two items (“exotic fruits”) were simultaneously presented for a duration of 2500 milliseconds, during which the participant had to decide if the two were perceptually identical (“down arrow” key) or different (“up arrow” key). That is, relating to the two stimuli, in each trial the participant had to answer the question “Are these two items identical?”. Items were paired as follows: (i) A.1 vs. B.1 (highly similar items from different categories); (ii) A.1 vs. A.1, or B.1 vs. B.1 (identical same-category items at the between-categories boundary); (iii) A.3 vs. A.4, or B.3 vs. B.4 (different items from the same category). (iv) A.3 vs. A.3, A.4 vs. A.4, B.3 vs. B.3, or

B.4 vs. B.4 (identical items from the same category). Left-screen/right-screen item presentation was counterbalanced. Each of the two (one pre- and one post-learning) discrimination tasks included 64 trials.

The pre- and post-learning item-preference tasks were used to assess pre- to post-learning changes in the preference for items from Category A versus items from Category B. In each trial of these tasks, two items were presented for a duration of 2500 milliseconds, during which the participant had to decide if she prefers the item presented on the left (“left arrow” key) or the one on the right (“right arrow” key). That is, relating to the two stimuli, the participant had to answer the question “Are these two items are of the same kind?” Items were paired as follows: (i) A.1 vs. B.1 (highly similar items from different categories); (ii) A.2 vs. B.2 (different-category items of medium similarity); (iii) A.4 vs. B.4 (different-category items of low similarity); (iv) catch trials with A.2 vs. A.4, or B.2 vs. B.4 (same-category items of medium similarity). Left- screen/right-screen item presentation was counterbalanced. Each item-preference task consisted of 72 trials.

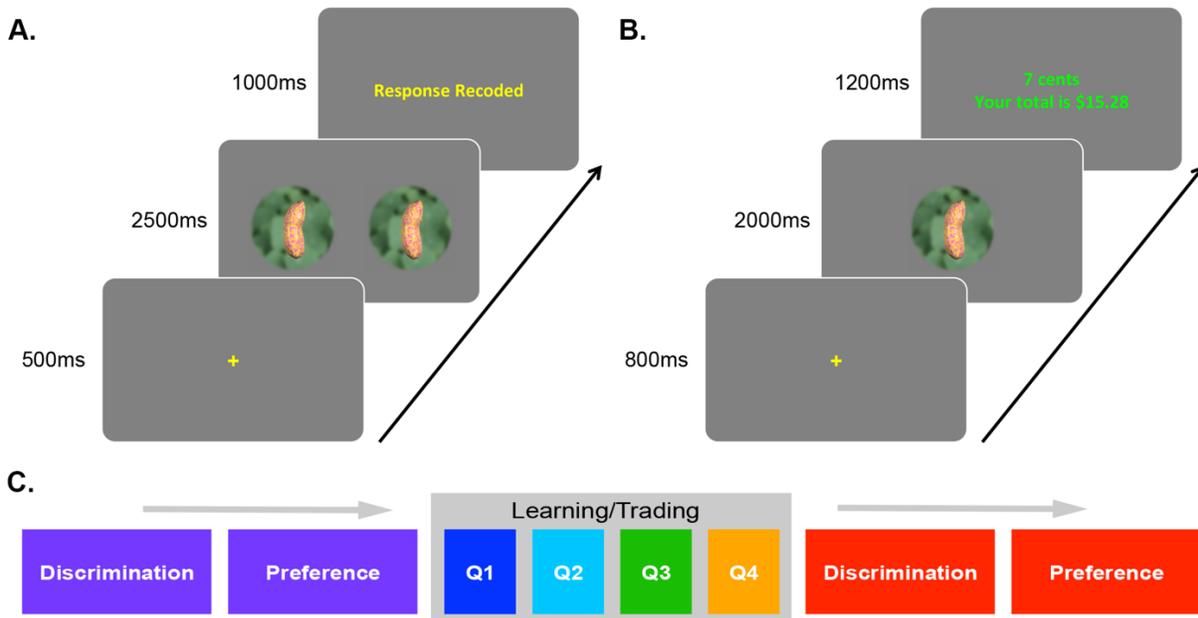
During the perceptual-discrimination tasks and item-preference tasks, there was no trial- by-trial feedback (for each trial, participants were only told whether they responded in a timely manner). Participants were informed about their overall task performance at the end of the experimental session (Figure 2A).

At the beginning of the learning phase, each participant was informed that she has \$13. The participant was then told that her performance during the learning phase would impact the amount of money she would receive at the end of the experimental session. Participants could make up to \$20 in addition to the \$13 they initially received (the actual sum depended on the experimental

condition and participant performance). It was possible, yet highly unlikely, for a participant to finish the experimental session with less than the \$13 initially received.

In each learning trial, a single “exotic fruit” was shown for 2000 milliseconds, during which the participant had to decide if she thought the fruit was profitable (and “select” it by pressing the “down arrow” key) or not (and “reject” it by pressing the “up arrow” key). That is, relating to the presented stimulus, the participant had to answer the question “Would you like to keep this fruit?” Following her decision, the participant received feedback indicating the monetary gain/loss associated with her decision, and the total amount of money she gained so far for the task. In trials where the participant earned money, the feedback was presented in green, and in trials where she lost money, it was presented in red (Figure 2B). For example, in the Gain–Risk condition, selecting an item (“down arrow” key) from Category A always resulted in a gain of 7¢. Not selecting the item (“up arrow” key) had no impact on the participant’s total earnings thus far (see item-value description in Table 1). In the Gain–Risk condition, selecting an item (“down arrow” key) from Category B could have resulted in a gain of 10¢ ( $p = 0.75$ ) or loss of 27¢ ( $-27¢$ ;  $p = 0.25$ ).

The learning task included 384 trials and was divided into four quarters with 96 trials each. At the end of each quarter, the participant could take a short break. The perceptually distinct between-categories stimuli were presented fewer times (items A.4 and B.4; 24 times each) than the perceptually confusable between-categories stimuli (A.1 and B.1; 88 times each). In each quarter of the learning phase, items in Category B were associated with one reward value in 75% of the trials and another reward value in the other 25% of the trials (see Table 1).



**Figure 2.** (A) Illustration of a trial composition for the pre- and post-learning discrimination and item-preference tasks, during which no trial-by-trial feedback was given to indicate decision accuracy or item value. (B) Illustration of a trial composition for the learning phase, during which trial-by-trial feedback indicated the monetary value of the item shown earlier in the trial. (C) The composition of an experimental session: (1) pre-learning perceptual-discrimination task; (2) pre-learning item-preference task; (3) learning phase; (4) post-learning perceptual-discrimination task; (5) post-learning item-preference task.

## RESULTS

Given the within-subject repeated-measure design, all the analyses below are based on repeated-measure ANOVAs and post-hoc t-tests. In all ANOVAs, the reported p-values are Greenhouse–Geisser corrected (sphericity is not assumed). The primary objective of the study was to test how the learning of each item-value affects categorization capabilities and the preferences of some items over others. We further investigated how item-value learning altered perceptual capabilities, and to characterize differences between the three experimental conditions in post learning behavior. The following analyses are aimed at addressing these questions.

**Performance in the learning phase.** In the learning phase, differences among the three experimental conditions were assessed as change in item preference between the first and fourth learning quarters (first 96 learning trials [Q1], last 96 learning trials [Q4]). This analysis is based on the participants' selection rate of items A.1, A.2, B.1 and B.2 (the perceptually more similar items, closer to the between-categories boundary), since these were the only items also presented in Q4 of the learning task. The dependent measure, the item selection rate, is the number of trials in which an item (A.1, A.2, B.1 or B.2) was selected by the participant divided by the total number of trials in which

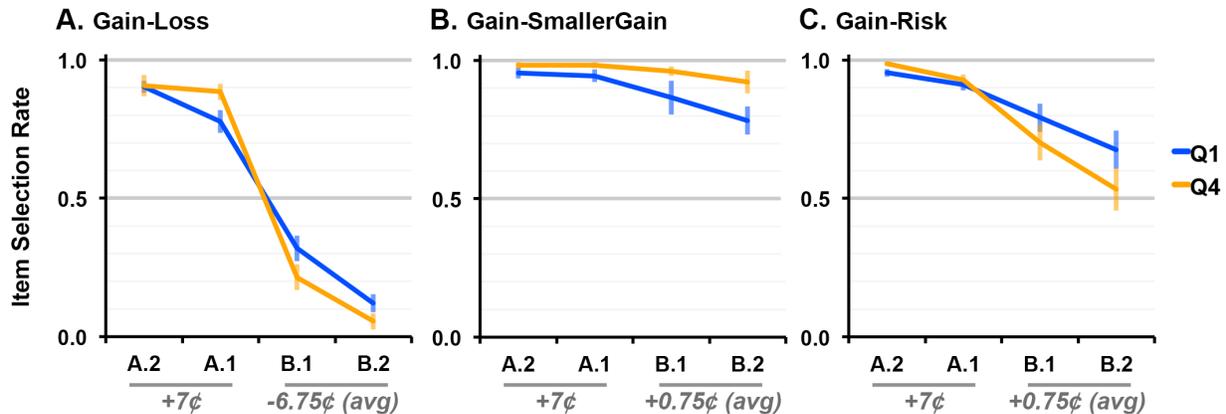
it was shown. In the learning phase, the item selection rate indicates whether the participant perceived the value of items from a given category as positive (that is, profitable and thus the item is expected to be frequently selected) or negative (associated with lose and thus will be rarely selected).

A  $3 \times 2$  ANOVA shows a significant interaction between conditions (Gain–Loss; Gain–Smaller-Gain; Gain–Risk) and stages of learning (Q1 versus Q4),  $F(2, 34) = 4.48$ ,  $p < 0.03$ ,  $\eta_p^2 = 0.21$ . This interaction is primarily driven by the distinct performance pattern in the Gain–Loss condition, where the participants quickly learned (already in Q1) to prefer items from Category A (with a fixed value of  $+7\text{¢}$  for all items) and to avoid items from Category B (with value of  $-9\text{¢}$  in 75% of the trials and  $0\text{¢}$  in 25% of the trials). This preference pattern was greater in Q4 (Figure 3).

A less expected and potentially more meaningful finding is the difference between the item-preference pattern in the Gain–Smaller-Gain and Gain–Risk conditions, where the mean item value was identical both in Category A ( $+7\text{¢}$ ) and in Category B ( $+0.75\text{¢}$ ). Accounting only to these two conditions, a  $2 \times 2$  ANOVA shows a significant interaction between condition (Gain–Smaller-Gain; Gain–Risk) and learning phase (Q1; Q4),  $F(1, 17) = 6.94$ ,  $p < 0.02$ ,  $\eta_p^2 = 0.29$ . This effect was associated with a

significant condition-by-phase-by-stimuli linear contrast,  $F(1, 17) = 6.67$ ,  $p < 0.02$ ,  $\eta_p^2 = 0.28$  (this linear contrast analysis accounts for item discriminability, i.e., the item distance from the between-categories boundary). Despite the two conditions having an identical average item-values in both categories, in the Gain–Smaller-Gain

condition, the selection rate of items from Category B increased from Q1 to Q4, whereas in the Gain–Risk condition, it decreased. This indicates an emerging aversion to items in Category B in the Gain–Risk condition, even though items in Category B were, on average, associated with a (small) positive value. This finding is consistent with the loss aversion account



**Figure 3.** Mean item selection rate in the learning phase in Q1 (blue) vs. Q4 (orange). The analysis is limited to items A.1, A.2, B.1 and B.2. Error bars are standard error of the mean. Avg value = average monetary value of a category item.

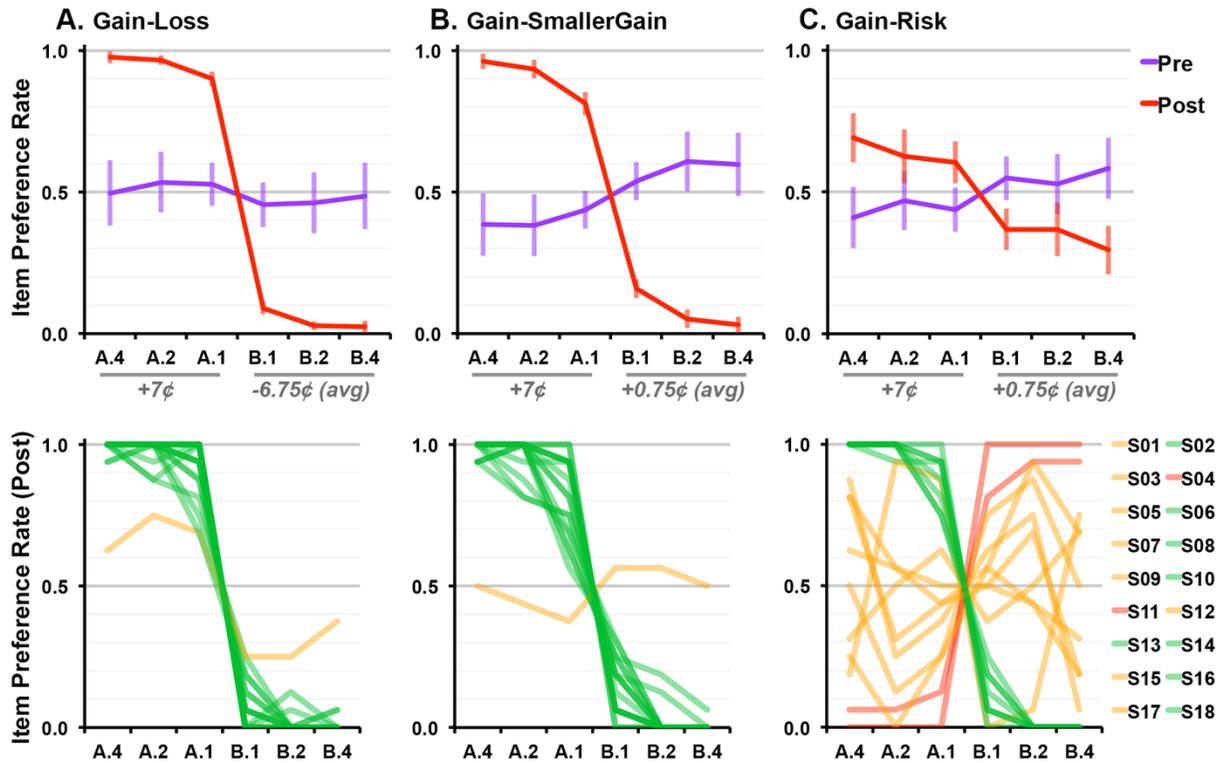
**Performance on the pre- and post-learning item-preference tasks.** For the item-preference tasks, differences between the three experimental conditions were assessed as the difference in item preference between pre-learning and post-learning. This analysis is limited to items A.1, A.2, A.4, B.1, B.2 and B.4, and it only includes trials in which the participant had to decide between a Category A and a Category B item (A.1 vs. B.1, A.2 vs. B.2, or A.4 vs. B.4). Catch trials (where both items were from the same category) are excluded from this analysis. The dependent measure in this analysis is the item-preference rate, which reflects how many times each item was preferred by the participant over an item from the other category. This preference rate shows which category was perceived as more valuable and the participants' ability to differentiate between items from the two categories.

One-way ANOVA for the pre-learning item-preference task shows no significant differences between the three conditions, all  $p > 0.20$ . That is, prior to learning the participants did not prefer items from one category over the other, as expected. However, most participants showed consistency in preferring some items over others (but across all participants, no item was preferred more frequently than others).

One-way ANOVA for the post-learning item-preference task shows a significant condition main effect  $F(1, 17) = 9.87$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.37$ , and a corresponding significant condition-by-stimuli linear contrast,  $F(1, 17) = 12.84$ ,  $p < 0.003$ ,  $\eta_p^2 = 0.43$ . This linear contrast

indicates that items preference was affected by both monetary value and perceptual discriminability. Post-hoc ANOVAs show significant differences between the Gain–Risk condition and the two other conditions, where the Gain–Risk condition differed from the other two conditions at  $p < 0.01$  (Bonferroni-corrected), with no significant differences between the Gain–Loss and Gain–Smaller-Gain conditions,  $p > 0.20$  (Figure 4).

While in the Gain–Loss and Gain–Smaller-Gain conditions almost all participants showed a clear preference for items from Category A (average item value =  $+7\phi$ ) to items from Category B (average item value of  $-6.25\phi$  and  $+0.75\phi$ , respectively), the Gain–Risk was characterized by substantial individual differences, where only 8 out of 18 participants consistently preferred items from Category A ( $+7\phi$ ). Two participants showed a clear preference for items from Category B (average value =  $+0.75\phi$ ), and the rest (8/18) showed mixed preferences (Figure 4C, bottom).



**Figure 4.** Top row: Mean item-preference rate for the pre- and post-learning item-preference tasks (limited to trials where participants had to choose between items A.1 vs. B.1, A.2 vs. B.2 or A.4 vs. B.4; error bars are standard error of the mean). Bottom row: individual performances for the post-learning item-preference tasks. Participants with a clear preference for Category A are marked in green, participants with a clear preference for Category B are marked in red, and participants with mixed preferences are marked in orange. See Appendix 1 for an additional evaluation of the individual differences reported here.

In a debriefing conducted at the end of the Gain–Risk experimental session, 15 out of the 18 participants reported that they believed it was within their ability, or at least that they thought it was possible, to differentiate between  $+10ϕ$  and  $-27ϕ$  items in Category B. Only three participants claimed that they suspected that the item-value assignment for items in Category B was random. Two of these participants had a distinct preference for items from Category A. That is, even after a few hundred learning trials, most participants had a false belief that they can determine which of the items from Category B would yield a profit. A Chi-squared test indicate that this false-believe rate (15/18) is greater than chance (9/18),  $\chi^2(1) = 4.50$ ,  $p < 0.05$ .

**Performance in the pre- and post-learning perceptual-discrimination tasks.** Pre- and post-learning perceptual-discrimination abilities were assessed using two measures: (i) a correct response rate when the participant was asked if two items are identical (A.1. vs. A.1, or B.1 vs. B.1) or different (A.1 vs. B.1) for the perceptual-discrimination tasks; (ii) the rate at which item A.1 (larger average value) was preferred over item B.1 (smaller average value) for the item-preference tasks.

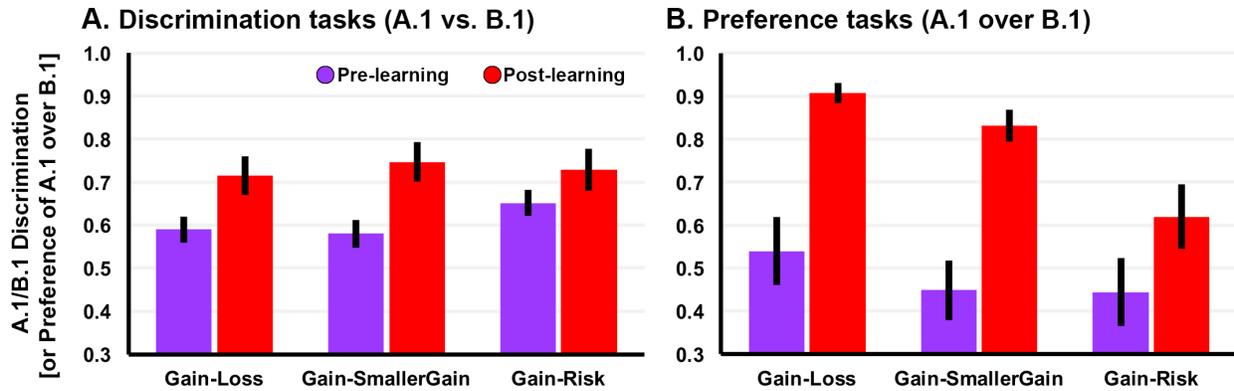
Note that for both tasks, the trial composition (the presented stimuli and trial duration) was identical, and the only difference was the framing of the question to the participant: “Are these two items identical?” (perceptual-discrimination tasks) vs. “Which one of these two items do you prefer?” (item-preference tasks). That is, the two measures differed only in participant motivation.

A  $2 \times 3$ , phase (pre-learning vs. post-learning) by condition (Gain–Loss, Gain–Smaller-Gain, and Gain–Risk), ANOVA for the perceptual-discrimination tasks shows only a significant main effect, indicating a pre- to post-learning improvement in participant ability to discriminate between items A.1 and B.1 (most similar between-category items),  $F(1, 17) = 20.23$ ,  $p < 0.0001$ ,  $\eta_p^2 = 0.54$  (Figure 5A).

Similarly, a  $2 \times 3$ , phase (pre-learning vs. post-learning) by condition (Gain–Loss, Gain–Smaller-Gain, and Gain–Risk), ANOVA for the item-preference tasks shows a significant main effect, indicating a pre- to post-learning increase in the preference for item A.1 over B.1,  $F(1, 17) = 70.50$ ,  $p < 0.0001$ ,  $\eta_p^2 = 0.81$ . However, this

ANOVA also shows an experimental condition main effect,  $F(1, 17) = 3.32$ ,  $p < 0.05$ ,  $\eta_p^2 = 0.16$ . Post hoc t-test indicate significant differences in the Gain–Loss and Gain–Smaller–Gain conditions (both  $p < 0.005$ ), with a

small and non-significant preference for item A.1 over item B.1 in the Gain–Risk condition ( $p > 0.1$ ). There was no significant phase-by-condition interaction,  $F(1, 17) = 1.63$ ,  $p > 0.2$  (Figure 5B).



**Figure 5.** Mean between-category perceptual discrimination indicating the ability to discriminate between items A.1 and B.1 (most similar between-category items) for the perceptual-discrimination tasks, and as the preference for item A.1 (an item with a larger average value) over item B.1 (an item with a smaller average value) for the item-preference tasks (error bars are the standard error of the mean). (A) Perceptual-discrimination tasks. (B) Preference tasks.

## DISCUSSION

Little is known about how risk assessment affect decision-making in scenarios where learning about two perceptually similar categories co-occurs with learning about the mean value of items from each category (i.e., experience-based utility learning). The current study tested individual differences in learning and decision-making in such scenarios. In the Gain–Loss and Gain–Smaller–Gain conditions, almost all participants learned to differentiate Category A from Category B, and to prefer items from Category A (+7¢), with greater mean value, over items from Category B (average value of -6.75¢ in the Gain–Loss condition and +0.75¢ in the Gain–Smaller–Gain condition). However, in the Gain–Risk condition, most participants did poorly in transferring knowledge acquired in the learning phase to the post-learning item-preference task, where they had to choose which one of two presented items is with a greater value. In the learning phase during the Gain–Risk condition, there was a decrease in the selection rate of items from Category B as the task progressed. This indicates that participants perceived items in Category B as likely to be associated with losses. That is, most participants underestimated the average true value of items in Category B (which was +0.75¢), and perceived those as negative. This is specifically supported by comparing the performance in the Gain–Risk condition with the performance in the Gain–Smaller–Gain condition, where items in Category B had the same average value but where the selection rate of Category B items increased as the task progressed. Nevertheless, in the Gain–Risk item-preference task, where the

participants had to choose which one of two items is with greater value, the preference towards items from Category B was on average the same as for items from Category A.

Differences between the Gain–Smaller–Gain and Gain–Risk conditions in the learning phase (Figure 3) can be attributed to a slight underestimation of the value of items in Category B in the Gain–Risk condition. It seems that participants mistaken it to be negative on average (instead of +0.75¢). This could have resulted from a minor miscalculation of item-value probabilities (which were +10¢ in 75% of the trials and -27¢ in 25% of the trials). However, not showing a clear preference for items from Category A (with a value of +7¢) in the Gain–Risk (post-learning) item-preference task (Figure 4) is a sign of an irrational decision-making strategy (frequently preferring items from Category B despite having less than ninth the average value of Category A). The preference for items from Category B could have been considered as a rational strategy only if there were twelve times more trials with an item value of +10¢ than trials with a value of -27¢. In such a case the average value of a Category B item would have been greater than +7¢ (7.15¢). However, in the current experiment this ratio was much smaller (3:1) with an average value of only +0.75¢ for Category B items. That is, the decision-making of many participants substantially deviated from what should be expected from a rational Bayesian inference. Evidently, most participants failed to approximate the posterior probabilities of item values based on their experience in the learning phase (see Vul et al, 2014 for a related discussion).

The perceptual ambiguity in the experimental tasks likely contributed to the participants' confusion, resulting in most participants believing that they could differentiate between +10¢ and -27¢ items within Category B (Gain–Risk condition). In a debriefing conducted at the end of the Gain–Risk experimental session, 15 out of the 18 participants reported that they believed it was within their ability, or that they thought it was possible, to differentiate between the +10¢ and the -27¢ items in Category B. Only three participants reported that they suspected items in Category B to have a 'random' value assignment. Evidently, the perceptual ambiguity resulted in risk-taking behavior by many participants who had the false belief that they can discriminate between the +10¢ items and the -27¢ items from Category B. Effectively, that was impossible since each Category B stimulus was associated with both +10¢ and -27¢ in a 3:1 ratio. Two participants (S04 and S11; Figure 4) showed a clear preference for Category B items in the post-learning item-preference task, indicating an even greater misjudgment of items' value. The clear preference for items from Category B in the post-learning item-preference task indicates that these two participants had good capacity to perceptually differentiate between the two categories.

In Appendix 1 we present a supplementary analysis suggesting that people can be grossly divided into two subgroups of decision maker. The first is a group of participants that in the Gain–Risk condition showed relatively higher preference for Category B items over Category A items in the post-learning item-preference task (cluster C1 in Panel C). We found that all of those showed lower risk aversion at the later stage of the learning task (Q4). On the other hand, low post-learning preference for Category B items was evident among participants regardless of the degree of their risk aversion in the later stage of the learning task (cluster C2 in Panel C). That is, high preference for Category B items (with lower mean value than Category B items) in the post-learning item-preference task was always associated with low risk aversion during learning. On the other hand, low preference for Category B items in the post-learning item-preference task is likely to be attributed to a more rational, even if imperfect, decision-making. All the participants in cluster C2 likely realized that the value of Category B items is low and substantially lower than Category A items. But some were more capable than the others in determining that the mean value of Category B items was nevertheless positive. This could have resulted in higher preference to Category B items during the learning task (selecting a Category B item when offered with one), but not in the item-preference task (recognizing a Category A item to be a better choice than a Category B item).

In the perceptual-discrimination tasks, where the participants were asked to decide if two stimuli were identical or different, without accounting to the items value, we found that the pre- to-post learning changes in the ability to differentiate between the two categories was about the same in all three experimental conditions (for related discussions, see Hammer & Diesendruck, 2005; Hammer et al., 2008, 2009; and Hammer, Sloutsky & Grill-Spector, 2015). Nevertheless, performance in the post-learning item-preference task indicates that greater monetary incentive was associated with better perceptual-discrimination capabilities. This was evident in a greater preference to item A.1 over item B.1, where the perceived value contrast between the two categories was greater: The lowest discrimination capability was evident in the Gain–Risk condition, whereas a near-perfect discrimination capability was evident in the Gain–Loss condition where participants almost always preferred item A.1 (+7¢) to the perceptually similar item B.1 (with a mean value of -6.75¢). This finding further exemplifies the interactive effect of low perceptual discriminability with the assessment of gain and loss. Here we find that greater perceived value-contrast between two highly similar categories increases people's capacity to differentiate between those categories. That is, while high-similarity between categories may make the process of value-learning more difficult and confusing, limited capacity to assess items' value due to non-monotonic feature-value mapping may compromise perceptual learning capabilities that are required for differentiating between similar categories.

We suggest that the substantial deviation from a rational decision-making pattern associated with the use of Bayesian inference (in the Gain–Risk condition), by many participants, is unlikely to be associated with a payoff-variability effect (Myers & Sadler, 1960), *per-se*. Instead, our findings indicate that suboptimal decision-making results from people failing to detect which past events are like those being presently assessed (Erev & Roth, 2014; Lynn et al., 2015). Here, we showed that such poor assessment may result from perceptual ambiguity along a single visual feature-dimension (an object shape) only whenever there is a non-monotonic feature-value mapping (associated with risk-taking). The minor perceptual differences between stimuli and actual categories increase the odds for people holding to the false belief that they can predict the value of risky items based on some perceptual features, even when this is impossible. It is possible that poor decision-making patterns, like those found in the Gain–Risk condition, may also characterize scenarios where differences between item categories are vague on a more abstract level. For example, items associated with gain or loss may often be differentiated by a complex rule that requires incorporating multiple sources of information or

accounting for multiple perceived features. Learning such a categorization rule may also require persistent trial and error. As we showed here, most people will likely fail to learn it whenever the learning process involves non-monotonic feature-value mapping characterizing risk-taking.

In summary, we show that category learning based on item-value assessment in conditions that involve risk-taking may result in compromised information gathering, and subsequently suboptimal decision-making that deviates from Bayesian inference. Such deviation may characterize a large portion (approximately half) of highly educated mentally healthy young adults. This is an exploratory study that covered only a fraction of the possible scenarios. To establish a more comprehensive model of the impact of risk assessment on category learning, future studies should account for other scenarios with different respective item values and value probabilities. An example would be a scenario where risky items have a value of -9¢ in 75% of the trials and a value of +30¢ in 25% of the trials (with a mean value of +0.75¢, as the one tested here, but with different probabilities for gain vs. loss). It is not guaranteed (and perhaps unlikely) that in such a scenario people would act the same way as in the Gain–Risk task tested here (see Dai et al., 2008 and Pratt (1964) for related discussions). Testing such a scenario can be valuable in verifying if the effect observed here is an outcome of the non-monotonic item-value mapping, per-se, or if it depends also on the risky items' loss versus gain probabilities. It would be also valuable to test if the poor risk assessment evident here would persist through multiple learning sessions with several thousand learning trials. It is possible that with sufficiently large statistics, people would become more accurate in their item-value assessment. Alternatively, it is possible that the early item-value bias would endure, with little impact from the number of learning trials. To better understand the origins of individual differences in risk assessment, future studies should also test a larger and more diverse sample of participants with greater age diversity.

## REFERENCES

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, 16(3), 215-233.

Barron, G., & Ursino, G. (2013). Underweighting rare events in experience based decisions: Beyond sample error. *Journal of Economic Psychology*, 39, 278-286.

Benartzi, S., & Thaler, R. H. (1995). Myopic loss aversion and the equity premium puzzle. *The quarterly journal of Economics*, 110(1), 73-92.

Brooks, L. R., Norman, G. R., & Allen, S. W. (1991). Role of specific similarity in a medical diagnostic task. *Journal of Experimental Psychology: General*, 120(3), 278.

Canessa, N., Crespi, C., Motterlini, M., Baud-Bovy, G., Chierchia, G., Pantaleo, G., ... & Cappa, S. F. (2013). The functional and structural neural basis of individual differences in loss aversion. *Journal of Neuroscience*, 33(36), 14307-14317.

Dai, X., Wertenbroch, K., & Brendl, C. M. (2008). The value heuristic in judgments of relative frequency. *Psychological Science*, 19(1), 18-19.

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, 112(4), 912-931.

Erev, I., Ert, E., & Yechiam, E. (2008). Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions. *Journal of Behavioral Decision Making*, 21(5), 575-597.

Erev, I., & Roth, A. E. (2014). Maximization, learning, and economic behavior. *Proceedings of the National Academy of Sciences*, 111 (Supplement 3), 10818-10825.

Hammer, R. (2015). Impact of feature saliency on visual category learning. *Frontiers in psychology*, 6.

Hammer, R., & Diesendruck, G. (2005). The role of dimensional distinctiveness in children's and adults' artifact categorization. *Psychological Science*, 16(2), 137-144.

Hammer, R., Bar-Hillel, A., Hertz, T., Weinshall, D., & Hochstein, S. (2008). Comparison processes in category learning: From theory to behavior. *Brain Research*, 1225, 102-118. Hammer, R., Diesendruck, G., Weinshall, D., & Hochstein, S. (2009). The development of category learning strategies: What makes the difference?. *Cognition*, 112(1), 105-119.

Hammer, R., Sloutsky, V., & Grill-Spector, K. (2015). Feature saliency and feedback information interactively impact visual category learning. *Frontiers in psychology*, 6.

- Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, 13(12), 517-523.
- Jessup, R. K., Bishara, A. J., & Busemeyer, J. R. (2008). Feedback produces divergence from prospect theory in descriptive choice. *Psychological Science*, 19(10), 1015-1022.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the econometric society*, 263-291.
- Knight, F. H. (1921) Risk, uncertainty and profit, Houghton-Mifflin.
- Lynn, S. K., Wormwood, J. B., Barrett, L. F., & Quigley, K. S. (2015). Decision-making from economic and signal detection perspectives: development of an integrated framework. *Frontiers in Psychology*, 6.
- Mehlhorn, K., Ben-Asher, N., Dutt, V., & Gonzalez, C. (2014). Observed variability and values matter: Toward a better understanding of information search and decisions from experience. *Journal of Behavioral Decision Making*, 27(4), 328-339.
- Myers, J. L., & Sadler, E. (1960). Effects of range of payoffs as a variable in risk taking. *Journal of Experimental Psychology*, 60(5), 306.
- Palmeri, T. J., Wong, A. C., & Gauthier, I. (2004). Computational approaches to the development of perceptual expertise. *Trends in Cognitive Sciences*, 8(8), 378-386.
- Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica: Journal of the Econometric Society*, 122-136.
- Ramsey, F.P. (1926). Truth and Probability. In Ramsey, 1931, *The Foundations of Mathematics and other Logical Essays*, Ch. VII, p.156-198, edited by R.B. Braithwaite, London: Kegan, Paul, Trench, Trubner & Co., New York: Harcourt, Brace and Company.
- Schonberg, T., Fox, C. R., & Poldrack, R. A. (2011). Mind the gap: bridging economic and naturalistic risk-taking with cognitive neuroscience. *Trends in Cognitive Sciences*, 15(1), 11-19. Sokol-Hessner, P., Hsu, M., Curley, N. G., Delgado, M. R., Camerer, C. F., & Phelps, E. A. (2009). Thinking like a trader selectively reduces individuals' loss aversion. *Proceedings of the National Academy of Sciences*, 106(13), 5035-5040.
- Tom, S. M., Fox, C. R., Trepel, C., & Poldrack, R. A. (2007). The neural basis of loss aversion in decision-making under risk. *Science*, 315(5811), 515-518.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4), 297-323.
- Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cognitive Science*, 38(4), 599-637.
- Weber, E. U., Shafir, S., & Blais, A. R. (2004). Predicting risk sensitivity in humans and lower animals: risk as variance or coefficient of variation. *Psychological Review*, 111(2), 430.

#### ACKNOWLEDGMENT

We thank Miguel Brendl (Kellogg Scholl of Management, Northwestern University) for insightful discussions.

#### AUTHOR CONTRIBUTION

Conceptualization of Research Objectives: R.H.; Method Development: R.H. and M.C.; Data Acquisition and Curation: R.H.; Formal Analysis: R.H.; Visualization: R.H.; Writing – Original Draft: R.H.; Writing – Review & Editing: R.H. and M.C.; Resources Acquisition: M.C. and R.H.

#### FUNDING

This study did not receive any grants from funding agencies in the public, commercial, or not-for-profit sectors.