Journal of Experimental Psychology: Learning, Memory, and Cognition

How Environmental Regularities Affect People's Information Search in Probability Judgments From Experience

Janine Christin Hoffart, Jörg Rieskamp, and Gilles Dutilh

Online First Publication, July 19, 2018. http://dx.doi.org/10.1037/xlm0000572

AMERICAN PSYCHOLOGICAL ASSOCIATION

Journal of Experimental Psychology: Learning, Memory, and Cognition

0278-7393/18/$12.00

http://dx.doi.org/10.1037/xlm0000572

RESEARCH REPORT

How Environmental Regularities Affect People's Information Search in Probability Judgments From Experience

Janine Christin Hoffart, Jörg Rieskamp, and Gilles Dutilh
University of Basel

In everyday life, people encounter smaller rewards with higher probability than larger rewards. Do people expect this reward–probability regularity to hold in experimental settings? To answer this question, we tested whether people's behavior in probability judgment tasks is affected by the correlation between reward size and reward probabilities. In Study 1, we asked people to judge reward probabilities under uncertainty. In line with the ecological reward–probability correlation, people assumed that larger rewards were less likely than smaller rewards. In Study 2, we tested the prediction that people's information search and integration depend on the representativeness of the environment. Participants performed an experience-based probability judgment task in which they sampled outcomes from unknown gambles until they felt confident to estimate the probabilities of the gambles' outcomes. We manipulated the reward–probability relationship of the gambles in 3 experimental groups. Rewards and reward probabilities were negatively correlated, positively correlated, or not correlated at all. A negative correlation mimics the ecological reward–probability relationship often present in real life. We analyzed people's search effort and whether they integrated sample-based uncertainty into their judgments. We found that people sampled fewer outcomes in the ecologically representative condition than in the other 2 conditions. However, people did not integrate sample-based uncertainty in their judgments: In all conditions people treated the observed outcomes as representative of the underlying outcome distribution. People's prior beliefs about regularities in environments provides a potential explanation of why people often rely on small sample sizes when making judgments and decisions from experience.

Keywords: probability updating, sample size, probability judgments, decisions from experience, representative design

How people search for information, when they stop searching, and how they integrate information are central questions of cognitive psychology (e.g., Busemeyer & Rapoport, 1988). Recently, there has been growing interest in how people make decisions from experience. In experience-based tasks, people sample outcomes from unknown gambles to learn about the gambles' outcomes and outcome probabilities (Hertwig, Barron, Weber, & Erev, 2004). Several factors have been identified that influence information search and integration in these tasks (e.g., Hau, Pleskac, Kiefer, & Hertwig, 2008; Hertwig et al., 2004; Rakow, Demes, & Newell, 2008; Roth, Wänke, & Erev, 2016). Yet in this area, past work has mostly neglected how initial beliefs and expectations about the structural regularities of risky gambles influence search effort and behavior (but see Lejarraga, Hertwig, & Gonzalez, 2012; Mehlhorn, Ben-Asher, Dutt, & Gonzalez, 2014). This is surprising when considering that regularities in the structure of experimental stimuli influence behavior in a range of tasks such as causal reasoning and probability judgments. In particular, when people's expectations are congruent with the labeling of stimuli, they learn relationships faster (e.g., Busemeyer, Byun, Delosh, & McDaniel, 1997). If the environment contradicts prior beliefs, for instance, those formed on the basis of real-life experience, people need strong evidence to overcome their beliefs (e.g., Alloy & Tabachnik, 1984).

On the basis of real-life experience, people can, for instance, expect that larger rewards will be less likely than smaller rewards (Pleskac & Hertwig, 2014). Interestingly, specific experimental designs can also lead to regularities between properties of gambles, such as correlations between returns and risk or between rewards and reward probabilities. In some cases, correlations are representative of real-life correlations, in others not. Generally, researchers may not be aware of how specific regularities of their stimulus material influence people's behavior (Fiedler, 2000). More concretely, it has not been well studied whether regularities of gambles, particularly correlations between rewards and reward proba-

Janine Christin Hoffart, Jörg Rieskamp, and Gilles Dutilh, Economic Psychology, Department of Psychology, University of Basel.

This research was supported by a grant from the Swiss National Science Foundation (SNF 143854) to Jörg Rieskamp and Gilles Dutilh.

Correspondence concerning this article should be addressed to Janine Christin Hoffart, Economic Psychology, Department of Psychology, University of Basel, Missionsstrasse 62a, 4055 Basel, Switzerland. E-mail: janine.hoffart@unibas.ch

bilities, affect behavior in decision-making experiments. Furthermore, it is not known whether the representativeness of these correlations influences behavior. We address these questions in two studies. In Study 1, we examined whether people have specific expectations about how rewards and reward probabilities of risky gambles are correlated. In Study 2, we examined whether people search for less information and integrate information differently in ecologically representative environments where rewards and reward probabilities of risky gambles are negatively correlated as compared to unrepresentative environments.

The Structure of the Environment and Ecological Rationality

Natural environments are characterized by certain statistical regularities. In financial decision-making situations, expected returns are typically positively correlated with risk (measured as the variance of returns). For instance, investments in stocks usually offer higher returns than investments in bonds but involve larger risk. This risk–return trade-off forms the basis of Markowitz's (1952) mean-variance model and standard portfolio theory (e.g., Sharpe, 1964). Whereas risks and returns are positively correlated, rewards and reward probabilities are negatively correlated across many situations: Pleskac and Hertwig (2014) analyzed the underlying regularities of several real-world domains including roulette, horse racing, life insurance, artificial insemination, and scientific publications. In essence, they showed that across all domains, larger gains occurred with lower probabilities than smaller gains. For instance, the larger the prize of a national lottery is, the smaller the probability of winning this prize.

Relying on regularities of the environment can simplify judgment processes and help people to make accurate decisions and judgments (e.g., Gigerenzer, Todd, & ABC Research Group, 1999; Todd & Gigerenzer, 2012). For instance, in environments where different sources provide redundant information, people search for little information and focus on the most valid information when making decisions. But in environments with little information redundancy, people search for more information and integrate all information when making decisions (Dieckmann & Rieskamp, 2007). In summary, people often adapt their behavior to the environment by selecting decision strategies that are appropriate for the structure of the environment (e.g., Mata, Schooler, & Rieskamp, 2007; Pohl, 2006; Rieskamp & Otto, 2006). People's adaptation to natural environments has been recently demonstrated with the risk–reward heuristic.

The Risk–Reward Heuristic

Work on the risk–reward heuristic describes how people estimate reward probabilities from potential costs and benefits of risky options (Pleskac & Hertwig, 2014). Pleskac and Hertwig (2014) offered participants an opportunity to gamble at the cost of $2. Participants had a chance to win a monetary reward whose magnitude was known and varied between participants. Importantly, participants had to estimate the unknown reward probability. In line with the ecological negative correlation between rewards and reward probabilities, people estimated the chances of winning smaller rewards as being higher than the chances of winning larger rewards. Pleskac and Hertwig (2014) concluded that under complete uncertainty, people heuristically inferred reward probabilities from the ratio of the reward magnitude and the costs of gambling. Yet in many experimental situations, people have some information about the probabilities of possible outcomes. For instance, in experience-based tasks, people sample outcomes to estimate outcome probabilities. Such tasks involve what Knight (1921) called *statistical probabilities*, which lie on the continuum between complete uncertainty and risk (Camilleri & Newell, 2013; Hau, Pleskac, & Hertwig, 2010; Hertwig & Erev, 2009).

Rewards and Reward Probabilities Are Correlated in Experience-Based Tasks

Interestingly, the payoff structure of the most frequently used gambles in decisions-from-experience (dfe) studies (e.g., Hau et al., 2008; Hertwig et al., 2004; Rakow et al., 2008) resembles the negative reward–reward probability correlation observed outside the laboratory: Larger rewards are less likely than smaller rewards. Figure 1 displays gains and probabilities of gambles used in dfe studies (e.g., Hau et al., 2008; Hertwig et al., 2004; Rakow et al., 2008). Rewards and reward probabilities are negatively linked across competing gambles (gray dashed lines link competing gambles) and across all gambles (black dotted line). While the negative correlation across gambles is particular to the gamble pairs we analyzed, a similar correlation as for the competing gambles has been reported elsewhere (Pleskac & Hertwig, 2014). A reason for this correlation can be deduced from the way researchers design experiments and specially the pairs of gambles they use: For choices to be informative for researchers, competing gambles typically have similar expected values (Rieskamp, 2008). Consider a pair of gambles where each gamble offers a reward with some probability and zero otherwise. When the first gamble's reward size is larger than the second gamble's reward size, the first gamble's reward probability needs to be smaller to lead to similar

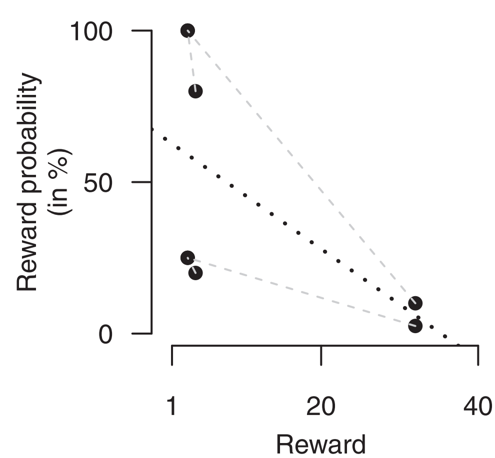

Figure 1. Relationship between reward probabilities and gains of gambles that have been used in many decisions-from-experience studies. Each dot represents one gamble that consists of two outcomes (reward or zero outcome). Gray dashed lines link competing gambles. Gambles that are linked with more than one other gamble were used in more than one choice trial. The black dotted line indicates the regression line between rewards and probabilities across all gambles.

expected values of both gambles. Using gambles with similar expected values is required to avoid trivial, noninformative choices. However, the observation that reward sizes and reward probabilities are negatively correlated in experience-based tasks may help to explain unresolved questions about how people search for and integrate information in experienced-based tasks.

In experience-based tasks, people often search less than would be required for an unbiased representation of the gambles' outcome distribution (e.g., Hau et al., 2008; Hertwig et al., 2004; Rakow et al., 2008). Reasons why people rely on small sample sizes range from sampling costs such as time and opportunity costs (Hau et al., 2008) to memory constraints that make it difficult to deal with large sample sizes (e.g., Hertwig et al., 2004; Rakow et al., 2008, but see Plonsky, Teodorescu, & Erev, 2015). Lejarraga et al. (2012) proposed that people in an experience-based task learn over the course of the experiment how gambles are structured and exploit this knowledge by sampling less. Here, we extend the hypothesis that learning influences behavior in experience-based tasks and investigate how the structure of gambles and people's expectations influence search effort.

Outline of the Studies and Predictions

In Study 1 we investigated whether people expect a negative correlation between rewards and reward probabilities in decision-making tasks. This study was an extension of Pleskac and Hertwig's (2014) work on the risk–reward heuristic. There, participants had two cues (i.e., reward probability and costs of gambling) to inform probability judgments. In choice studies, typically no explicit costs are associated with gambling. Therefore, in Study 1, we assessed people's probability expectations for different reward amounts when no costs were associated with gambling.

In Study 2 we investigated people's search effort and accuracy in an experience-based probability judgment task. Briefly, we asked participants to repeatedly judge reward probabilities for two-outcome gambles (rewards of varying magnitudes or zero outcome). To make their judgments, participants first sampled outcomes from a gamble's outcome distribution. Using a between-subjects design, we manipulated the reward–reward probability correlation. We contrasted a representative condition in which rewards and reward probabilities were negatively correlated (representative negative correlation, RNC) with two nonrepresentative conditions in which rewards and reward probabilities either were not correlated (nonrepresentative no correlation, NRNC) or were positively correlated (nonrepresentative positive correlation, NRPC). We hypothesized that a correlation between rewards and reward probabilities affects search effort and/or information integration.

How Can a Correlation Influence Learning and Search Effort in Experience-Based Tasks?

As Lejarraga et al. (2012) has proposed, learning can influence search effort. To test this hypothesis, we compared how much people sampled between conditions. If general learning influences sample size, participants should sample less when they have learned the relationship between rewards and reward probabilities. To determine if participants learned the correlation, we prompted their probability expectations at the end of the experiment for different rewards without allowing them to draw any outcomes. If participants have learned how rewards and reward probabilities are correlated, their probability estimates in this control block should resemble the correlation experienced during the previous phases. Generally, people can learn the relationships in both correlation conditions, implying that people should draw fewer outcomes in both correlated conditions. However, previous research has shown that people learn relationships between cues and criteria faster when the relationship is congruent with their expectations (e.g., Busemeyer et al., 1997). If people expect rewards and probabilities to be negatively correlated, they may learn the reward–probability correlation in the representative condition faster. In this case, we expect that sample sizes will be smaller and people's probability estimates in the control block will more closely resemble the previously learned correlation in the representative condition than in the nonrepresentative condition.

We further analyzed the impact of learning by comparing how sample sizes evolved over trials. If learning impacts search effort and the control block reveals that people learned the relationships in both correlated conditions, sample sizes should decrease more in the two conditions with correlated rewards and reward probabilities than in the condition where rewards and reward probabilities are not correlated.

Alternatively, it could be that it is not general learning but familiarity with the search environment that influences search effort. In line with the notion that human cognition is adapted to natural environments (e.g., Brunswik, 1947, 1955; Gigerenzer, Hoffrage, & Kleinbolting, 1991; Gigerenzer & Hug, 1992), we hypothesized that people are adapted to navigating in environments where smaller rewards occur with higher probabilities than larger rewards. If participants' search effort is selectively influenced by representative correlations, we would expect first that they will draw fewer outcomes in the representative condition than in both nonrepresentative conditions and second that in both conditions where rewards and reward probabilities are correlated, people's responses in the control block will resemble the reward–probability structure observed in the previous blocks.

How Can a Correlation Between Rewards and Reward Probabilities Influence Accuracy in Experience-Based Tasks?

We propose that in correlated environments, people can exploit their knowledge about correlations and make reasonably accurate judgments based on relatively small sample sizes. In our experiment, we showed participants the outcomes (i.e., zero and a positive reward) before they started sampling. This implies that in the correlated conditions, participants could form prior beliefs about the reward probabilities based on the reward magnitudes.

However, previous research showed that people often integrate sample-based uncertainty insufficiently. Instead, they simply rely on the relative observed outcome frequency when they make judgments (e.g., Griffin & Tversky, 1992; Tversky & Kahneman, 1971). The "natural-mean" heuristic describes such a strategy and has been successfully applied in modeling decisions from experience (Hertwig & Pleskac, 2008).

We contrasted the two hypotheses stated above by comparing two models: The Bayesian updating model (M_B) assumes that people's probability judgments are based on their prior belief about the reward probability and a Bayesian updating process that relies on the sampled information. The model makes the assumption that people's prior beliefs follow the objective correlation present in the experimental environment. The second model is a variant of the natural-mean heuristic and captures the idea that people attend only to the observed relative outcome frequencies when judging probabilities (M_{NM}). The models are described in detail in Appendix A.

Study 1: Do People Believe That Rewards and Probabilities Are Correlated in the Laboratory?

To investigate whether people believe that reward magnitudes and probabilities are correlated across gambles in experimental tasks, we asked people in two experiments to infer reward probabilities from reward magnitudes when no costs were associated with gambling.

Method Experiment 1

Fifty-seven undergraduates from the University of Basel volunteered to participate in a paper-and-pencil survey (49 women, 8 men; M_{age} = 22.21 years, SD = 3.67, range = 18–35). Participants estimated reward probabilities for two-outcome gambles. The outcomes were zero or a specific known reward (reward magnitudes [in Swiss francs]: CHF 2.40, 4.00, 4.70, 12.00, 16.00, 28.00). These rewards were taken from past studies conducted at the Center for Economic Psychology and thus represented realistic outcomes of psychological experiments. The amounts were displayed randomly in increasing or decreasing order. For all analyses in this article, we used the software R (R Core Team, 2014).

Results and Discussion of Experiment 1

Participants' probability estimates varied as a function of reward amount: Participants assigned higher probabilities to smaller rewards than to larger rewards (see Figure 2). This finding is supported by a Bayesian hierarchical beta regression (see Appendix B for model description). With vague priors, the median posterior estimate for the overarching slope coefficient was $\beta_1 = -0.1$ (95% highest density interval [$HDI_{95\%}$]: −0.11 to −0.09). Described in terms of odds of winning, the results imply that for every CHF 10 decrease of reward amount, participants believed they were 2.72 ($HDI_{95\%}$: 2.6 to 3) times more likely to win. The 95% HDI of the slope coefficient excluded zero, indicating a reliable negative correlation between participants' probability estimates and rewards. The results were similar for both orders of reward presentation: slope: $\beta_{1,increasing} = -0.1$, $HDI_{95\%}$: −0.12 to −0.08 and $\beta_{1,decreasing} = -0.09$, $HDI_{95\%}$: −0.11 to −0.08; intercept (converted to the probability scale): $\beta_{0,increasing} = 0.47$, $HDI_{95\%}$: 0.33 to 0.6 and $\beta_{0,decreasing} = 0.51$, $HDI_{95\%}$: 0.4 to 0.63.

Experiment 1 supports the findings of Pleskac and Hertwig (2014). Participants assumed larger rewards were less likely to occur than smaller rewards. However, the fact that we always presented rewards in either increasing or decreasing order may be a potential confound. To control for the possibility that our results are an artifact of the presentation order of rewards, we replicated Experiment 1 with a fully randomized order of rewards.

Experiment 2

In the second experiment, we followed the protocol of Experiment 1 with one exception: We randomized the order in which we presented the rewards for every participant.

Method for Experiment 2. We predefined the number of subjects with a power analysis using a simulation approach. We searched for the minimum number of required subjects for finding a reliably negative slope coefficient when the data-generating slope coefficient was half as large as the median coefficient of Experiment 1 in 90% of the simulation. The power analysis revealed that we needed to test 42 subjects to reach the desired power. We collected data from 44 psychology students (38 women, 4 men, 2 not reported; M_{age} = 25.3 years, SD = 5.2, range = 20–53) at the University of Basel.

Results of Experiment 2. Similar to that in Experiment 1, participants' probability estimates varied as a function of reward magnitude: Participants estimated higher probabilities for smaller rewards than for larger rewards (see Figure 2). A beta regression revealed that the median posterior estimate for the overarching slope coefficient was $\beta = -.09$ ($HDI_{95\%}$: –.11 to –.08). This means that for every CHF 10 decrease of reward amount, participants believed they were 2.54 ($HDI_{95\%}$: 2.16 to 3) times more likely to win. Again, the 95% HDI of the slope coefficient excluded zero,

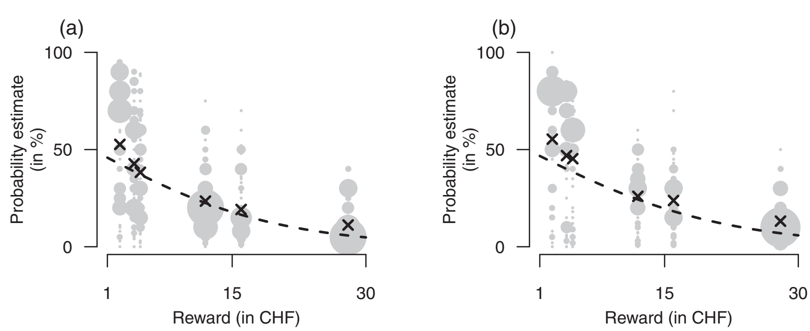

Figure 2. Probability estimates as a function of rewards in (a) Experiment 1 and (b) Experiment 2. The dots display individual responses; the dot size indicates response frequency. The crosses display the mean probability estimates. The dashed lines show the regression lines based on the median posterior group estimates of the coefficients of the beta regression. CHF = Swiss francs.

indicating a reliable effect of the reward magnitudes on participants' probability estimates.

Discussion of Study 1

Experiment 2 confirmed the findings of Experiment 1: People judged larger rewards as less likely than smaller rewards also when the rewards were presented in random order. In both experiments, we put our participants in a situation of total uncertainty: They had to judge reward probabilities without any context or guidance on how they should solve the task. Participants used the reward magnitude as the only available cue to inform their probability judgments. Thus, participants seemed—under uncertainty—to treat rewards and reward probabilities as correlated variables.

Study 2: Does a Correlation Between Rewards and Reward Probabilities Influence Behavior in an Experienced-Based Probability Judgment Task?

With Study 2, we explored whether a correlation between rewards and reward probabilities, and in particular the direction of correlation, influences behavior in a probability judgment task.

Method

We asked people to repeatedly judge from experience the probabilities for two-outcome (zero or positive outcome) gambles. Between subjects, we manipulated the mapping between rewards and reward probabilities. The range of possible rewards over the gambles was the same in each condition; however, in the RNC condition, rewards and reward probabilities were negatively correlated. In the NRPC condition they were positively correlated. In the NRNC condition they were not correlated. Figure 3 graphically represents these correlations.

Participants. Ninety subjects from the student pool of the University of Basel participated in the study (49 women, 41 men; M_{age} = 26 years, *SD* = 5, range = 19–45). The experiment was approved by the ethics committee of the University of Basel, and all participants signed informed consent. Participants decided between a show-up fee of CHF 15 (roughly $15 at the time of the experiment) or course credit. Additionally, they received a performance-dependent bonus payment (CHF 0 to CHF 10). To determine the bonus, in each trial of the test block, participants were endowed with CHF 0.20. From this endowment, we subtracted the squared deviation between response and objective reward probability, expressed as a percentage. If the deviation was larger than 9.99, the participant did not get any bonus for this trial. Participants received the summed bonuses of all trials of the test block (*M* = CHF 4.85, *SD* = 1.06).

Materials and procedure. The experiment had a learning block, a test block, and a control block.

Learning block. In the learning block, participants estimated reward probabilities for nine randomly presented two-outcome (reward or zero) gambles (see Table 1). In each trial, participants first saw the reward amount and then drew 15 random outcomes from the gamble's outcome distribution by clicking a button on a computer keyboard. We chose the sample size of 15 for two reasons: First, it ensured that participants experience a range of probabilities in the learning block, which is only possible when the sample size is not too small. Second, a sample size of 15 was observed as the median sample size per trial in Hertwig et al. (2004). Rewards were displayed in white numerals in the middle of a black circle. Zeros were displayed in black numerals in the middle of a white circle. After participants had drawn all 15 required outcomes, we prompted their reward–probability estimate. Participants then got feedback about the objective reward probability and how much their estimate deviated from this probability. The purpose of the learning block was twofold. First, it allowed people to adapt to the reward–probability relationship in this study. Importantly, this means that prior beliefs about possible correlations could be corrected in the NRNC condition. Second, it ensured that participants understood that the draws were random.

Test block. In the test block, participants estimated probabilities for 50 gambles (see Figure 3 for an overview of the gambles). The test block trials followed the same procedure as that used in the learning block with two exceptions: Participants decided freely

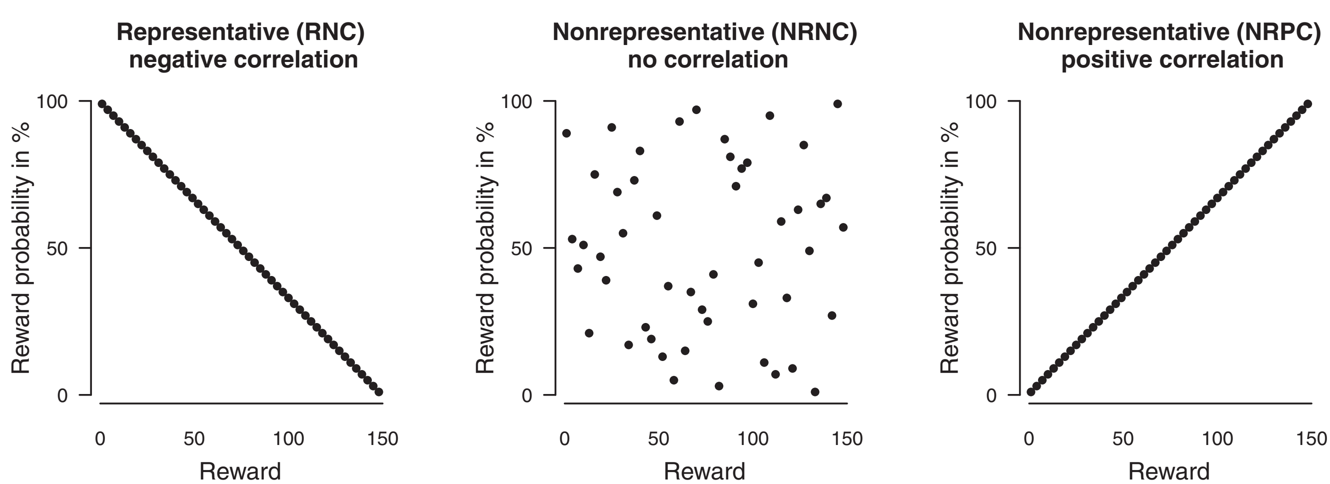

Figure 3. The relationship between the reward probabilities and the rewards in the test block separately for the three conditions. Zero occurs with counter probability. Each gamble is illustrated by a dot.

Table 1
The Nine Gambles of the Learning Block

		Reward probability per condition		
Gamble	Reward	Representative negative correlation	Nonrepresentative no correlation	Nonrepresentative positive correlation
1	2	.99	.52	.03
2	6	.96	.96	.06
3	11	.93	.03	.09
4	68	.55	.93	.49
5	72	.52	.55	.52
6	77	.49	.06	.55
7	137	.09	.09	.93
8	141	.06	.49	.96
9	146	.03	.99	.99

Note. All three conditions had the same rewards and probabilities.

how many outcomes they wanted to draw, and they did not receive performance feedback.

Control block. In the control block, participants estimated probabilities for nine gambles: Three gambles provided small rewards (3, 8, 12); three provided medium rewards (69, 74, 78); and three provided large rewards (138, 143, 147). Participants did not get to sample any outcomes. They saw a gamble's reward and estimated the reward probability based on knowledge acquired in the learning and test blocks without receiving feedback. The purpose of this block was to examine whether participants learned the correlations between rewards and probabilities during the previous two blocks.

Results

Learning block. We removed trials in which participants' probability estimates deviated by more than 50 percentage points from observed relative outcome frequencies (18 of 810 responses). We identified this outlier criterion before running the study on the basis of a pilot study in which some participants reported for a few trials that they erroneously judged the probability of the zero outcome instead of the positive reward. On average, participants' observed relative reward frequencies deviated by 6.3 (SD = 6) percentage points from the objective reward probabilities (M_{RNC} = 6.7, SD_{RNC} = 6.1; M_{NRNC} = 5.7, SD_{NRNC} = 5.5; M_{NRPC} = 6.5, SD_{NRPC} = 6.2). Figure 4 displays probability estimates as a function of objective reward probability and observed relative reward frequencies. As expected, participants'

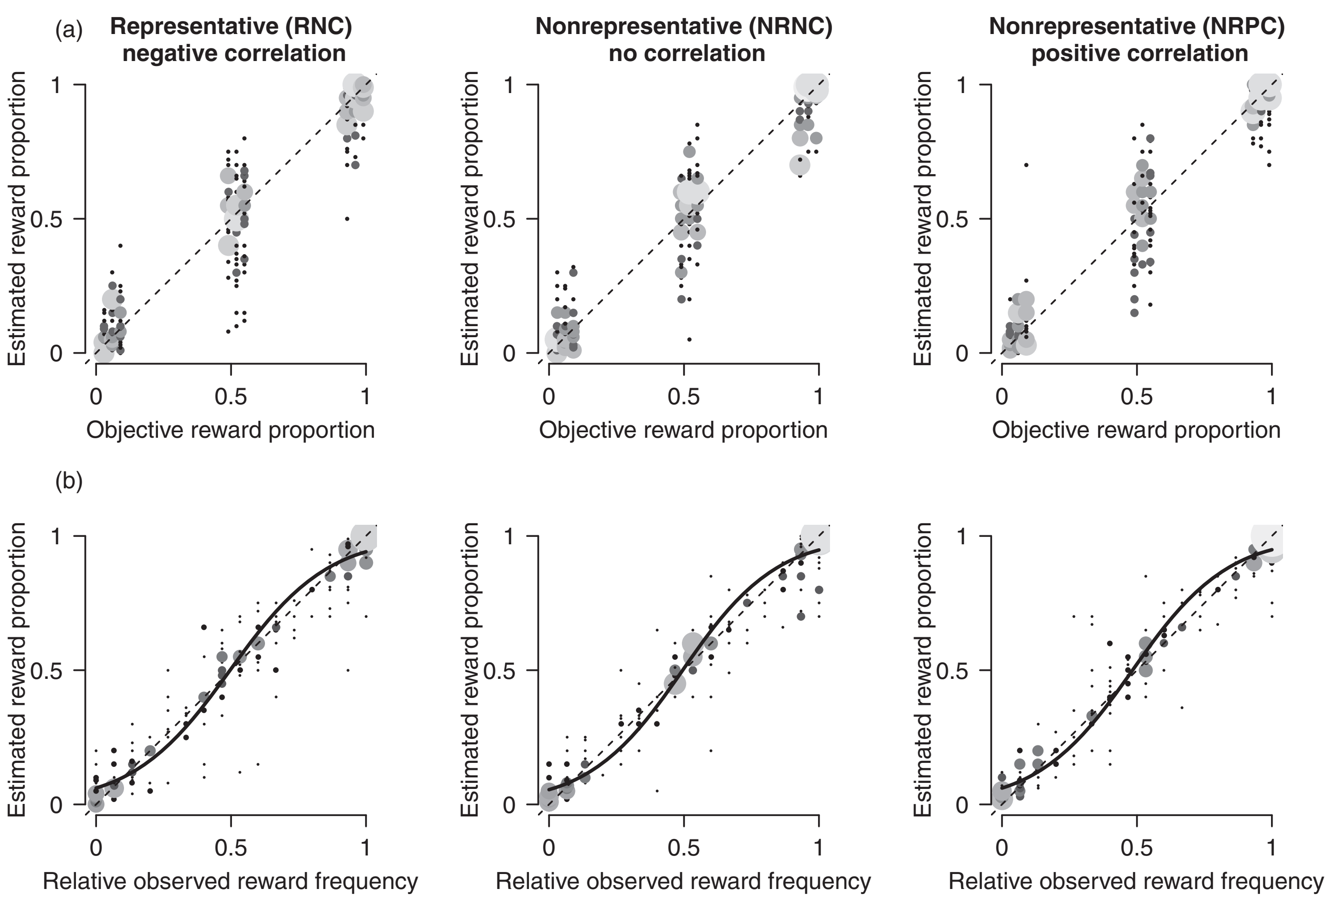

Figure 4. Estimated reward proportions in the learning block as a function of (a) objective reward proportion and (b) relative observed reward frequencies. The sizes of the dots illustrate how often each response was given. The dashed lines represent the 45° identity line. The solid lines represent the regression line based on the median posterior group estimates of the coefficients of the beta regression.

probability estimates increased as a function of increasing objective probabilities and relative observed reward frequencies. This was confirmed by a Bayesian hierarchical beta regression that regressed probability estimates on relative observed reward frequencies. The 95% HDIs of condition-dependent posterior slope and intercept estimates overlapped (see Table 2), indicating no reliable difference in accuracy between conditions. As the regression line in Figure 4b and the intercepts reveal, participants tended to slightly overestimate small probabilities in all conditions.

Test block. Before we analyzed the data, we log-transformed the number of outcomes that participants drew to approximate normally distributed data. We then removed trials in which participants sampled more than ±2 *SD* of the condition mean (163 of 4,500 responses). Furthermore, we removed trials in which participants' probability estimates deviated by more than 50 percentage points from observed relative outcome frequencies (52 of 4,337 responses).

Sample size. Across conditions, participants drew on average 21.1 ($SD = 9.5$) outcomes ($M_{\mathrm{RNC}} = 17.3$, $SD_{\mathrm{RNC}} = 6.5$; $M_{\mathrm{NRNC}} = 22.8$, $SD_{\mathrm{NRNC}} = 8.8$; $M_{\mathrm{NRPC}} = 23.2$, $SD_{\mathrm{NRPC}} = 11.3$).

Did participants sample less when rewards and probabilities were (negatively) correlated? We compared search effort between conditions with a Bayesian hierarchical model. We modeled the sample size *y* that a participant *j* drew in a given trial *i* as being drawn from a normal distribution with a specific participant mean μ_j and a participant precision τ_j. The participant means were modeled as draws from overarching condition-dependent normal distributions with mean $\mu_{\mathrm{Condition}}$ and group precision $\tau_{\mathrm{Condition}}$.

The condition-specific parameter $\mu_{\mathrm{Condition}}$ measures the posterior estimates for sample sizes separately for each condition. The posterior densities of these overarching $\mu_{\mathrm{Condition}}$ are shown in Figure 5 alongside the posterior densities of individual mean sample-size estimates (μ_j). The upper bound of the 95% HDI of μ_{RNC} does not overlap with the lower bounds of the HDIs of both other conditions. This indicates that participants reliably sampled less in the representative condition as compared to both nonrepresentative conditions.[1]

Figure 5 shows that the individual parameters μ_j are slightly more dispersed in both nonrepresentative conditions, suggesting larger between-subjects variability of mean sample sizes than in the representative condition. However, the 95% HDI of the posterior group estimates for the precision parameters $\tau_{\mathrm{Condition}}$ of all conditions overlap, which indicates that there were no reliable group differences in precision ($\tau_{\mathrm{RNC}} = 12.5$, $\mathrm{HDI}_{95\%}$: 6.6 to 19.9; $\tau_{\mathrm{NRNC}} = 10.9$, $\mathrm{HDI}_{95\%}$: 5.6 to 17; $\tau_{\mathrm{NRPC}} = 7.3$, $\mathrm{HDI}_{95\%}$: 3.8 to 11.5).

Did sample size decrease over the course of the experiment? Figure 6 displays how sample size evolved over trials in the different conditions. We analyzed whether sample size decreased with a Bayesian hierarchical linear regression with trial number as a predictor. Only for the representative condition did sample size reliably decrease over the course of the experiment, as indicated by the reliably negative slope parameter $\beta_{1,\mathrm{RNC}}$ (*Mdn* $\beta_{1,\mathrm{RNC}} = -0.003$; $\mathrm{HDI}_{95\%}$: −0.006 to −0.001). For both nonrepresentative conditions, the 95% HDIs of the condition-specific slope parameter include zero, suggesting there were no reliable changes of sample sizes over trials (*Mdn* $\beta_{1,\mathrm{NRNC}} = -0.001$, $\mathrm{HDI}_{95\%}$: –.004 to .001; *Mdn* $\beta_{1,\mathrm{NRPC}} = .001$, $\mathrm{HDI}_{95\%}$: –.004 to .001).

Probability judgments. On average, participants' observed relative reward frequencies deviated by 7.3 percentage points (pp) ($SD = 6.2$) from the objective reward probabilities ($M_{\mathrm{RNC}} = 8$ pp, $SD_{\mathrm{RNC}} = 6.9$; $M_{\mathrm{NRNC}} = 6.7$ pp, $SD_{\mathrm{NRNC}} = 5.8$; $M_{\mathrm{NRPC}} = 7.1$ pp, $SD_{\mathrm{NRPC}} = 6$). Figure 7 illustrates participants' probability estimates as a function of objective reward probabilities and relative observed outcome frequencies. Similar to in the learning block, a hierarchical Bayesian beta regression revealed that the 95% HDIs of the posterior slope and intercept coefficients overlapped in all conditions (see Table 2), indicating that accuracy did not differ between conditions.

We further analyzed participants' probability estimates with the Bayesian updating model M_B and the natural-mean heuristic M_{NM} (see Appendix A). We estimated both models with participants' individual data by applying maximum likelihood methods. To compute the likelihood, we assumed participants' probability estimates would follow a truncated normal distribution (0–1) around the model's predicted probability estimate. We estimated the standard deviation of the truncated normal distribution as a free parameter for each participant. With a grid-search approach we identified the set of parameter values that minimized the deviance (negative log likelihood) between predictions and responses. We searched the parameter space for the strength parameter *N* between 2 and 50 in steps of 1 and for the standard deviation of the truncated normal distribution between 0.00001 and 0.3 in steps of 0.01. We compared the models based on the Akaike information criterion (AIC), which accounts for the number of free parameters (P; $\mathrm{AIC} = -2 \times LogL + 2 \times P$; Burnham & Anderson, 2002).

Table 2
Beta Regression: Probability Estimates as a Function of Relative Observed Reward Frequencies

Block	Condition	Coefficient	Median	95% HDI
Learning	RNC	Intercept	0.06 (−2.7)	[0.05 (−2.9) to 0.07 (−2.6)]
		Slope	1.06 (0.06)	[1.05 (0.05) to 1.06 (0.06)]
	NRNC	Intercept	0.06 (−2.8)	[0.05 (−3.0) to 0.07 (−2.7)]
		Slope	1.06 (0.06)	[1.06 (0.05) to 1.06 (0.06)]
	NRPC	Intercept	0.06 (−2.7)	[0.06 (−2.9) to 0.07 (−2.6)]
		Slope	1.06 (0.06)	[1.06 (0.05) to 1.06 (0.06)]
Test	RNC	Intercept	0.07 (−2.7)	[0.06 (−2.8) to 0.07 (−2.6)]
		Slope	1.05 (0.05)	[1.05 (0.05) to 1.06 (0.06)]
	NRNC	Intercept	0.06 (−2.8)	[0.06 (−2.8) to 0.07 (−2.7)]
		Slope	1.06 (0.05)	[1.05 (0.05) to 1.06 (0.06)]
	NRPC	Intercept	0.07 (−2.6)	[0.06 (−2.7) to 0.08 (−2.5)]
		Slope	1.05 (0.05)	[1.05 (0.05) to 1.06 (0.06)]

Note. Coefficients of the hierarchical beta regression where we regressed probability estimates on the relative observed reward frequencies. Similar estimates for the median and the highest density interval (HDI) bounds result from rounding. The estimates describe the posterior group estimates of the condition-dependent slope and intercept estimates converted back from the logit scale. A slope coefficient of *x* can be interpreted in terms of the odds of participants' probability estimates: For each percentage point increase in the relative observed reward frequency, the odds of the probability estimates increased by a factor of *x*. The numbers in parentheses describe the median posterior estimates on the logit scale. RNC = representative, negative correlation; NRNC = nonrepresentative, no correlation; NRPC = nonrepresentative, positive correlation.

[1] This finding held (and was even accentuated) when we applied a stricter outlier criterion and removed all trials in which sample sizes deviated by more than 1 *SD* from the condition mean as well as when we did not exclude any data.

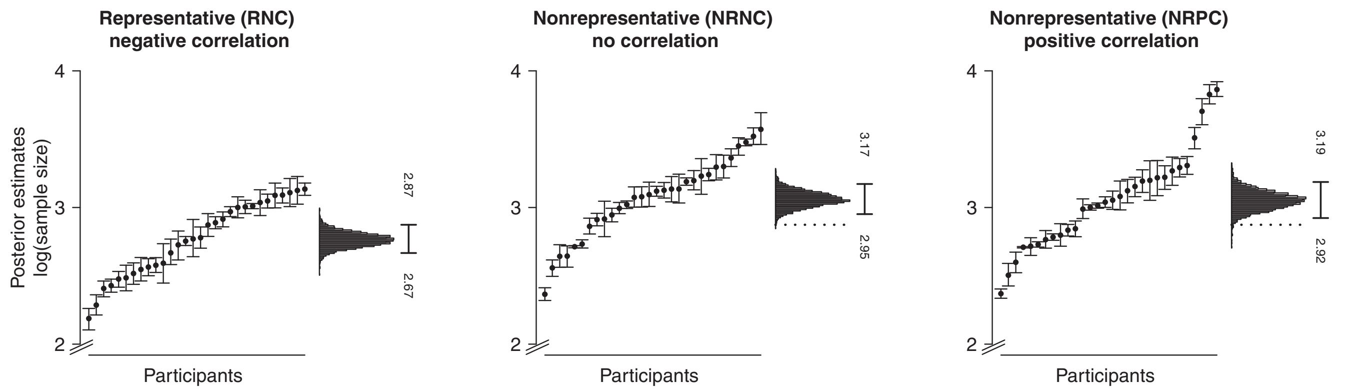

Figure 5. Posterior estimates of the group means (histograms) and posterior individual subject means (dots). The error bars indicate the 95% highest density intervals (HDIs). The dotted lines indicate the upper bound of the 95% HDI of μ_{RNC}.

Our model comparison revealed that most participants' data were best fit by the natural-mean heuristic. Only for 4 of the 90 participants did the Bayesian updating model describe the data better (RNC: 1 of 30; NRNC: 0 of 30; NRPC: 3 of 30).[2]

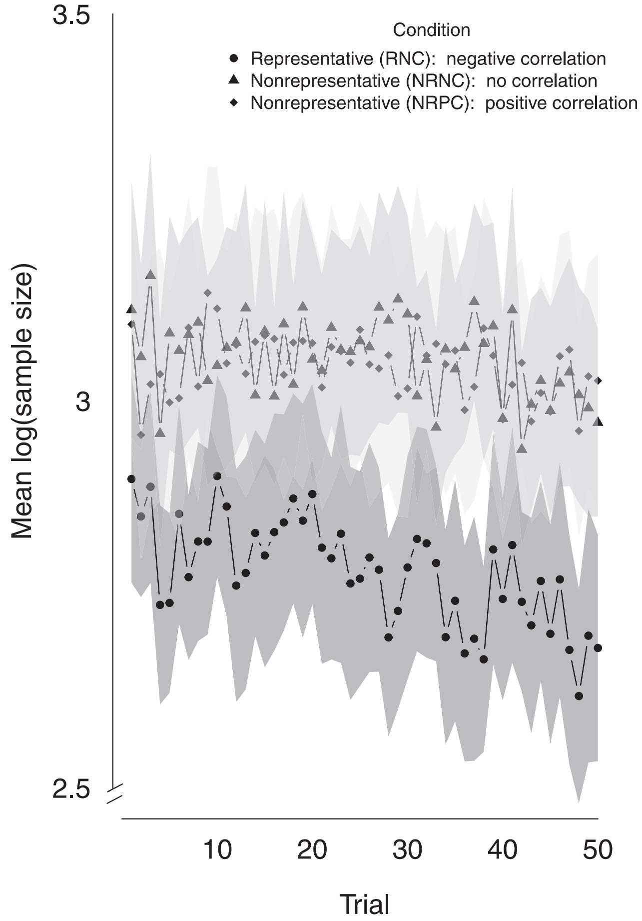

Figure 6. Mean of log sample sizes as a function of trial separately for the three conditions. The shaded area indicates 95% confidence intervals.

Control block. In the control block, participants estimated reward probabilities without sampling. Probability estimates clearly depended on the magnitude of the reward (see Figure 8) in both correlated conditions. This finding was supported by a Bayesian hierarchical beta regression that used reward magnitudes to predict probability estimates. In the RNC condition, probability estimates decreased as a function of reward magnitudes as shown by the median slope parameter $b_{1,\mathrm{Condition}}$ ($b_{1,\mathrm{RNC}}$ = –.023, $\mathrm{HDI}_{95\%}$: –.03 to –.014). In the NRPC condition probability estimates increased as a function of reward amount ($b_{1,\mathrm{NRPC}}$ = .016, $\mathrm{HDI}_{95\%}$: .009 to .022). In the NRNC condition there was no reliable systematic relation between probability estimates and reward magnitudes ($b_{1,\mathrm{NRNC}}$ = –.001, $\mathrm{HDI}_{95\%}$: –.006 to .004). The absolute values of the 95% HDIs' slope parameters of the two correlated conditions overlapped, suggesting that participants in both conditions learned the contingencies equally well.

General Discussion

The present work shows that people generally expect that the size of rewards and the probabilities with which these rewards occur in gambles are negatively correlated. Furthermore, people search for less information under uncertainty when this expectation is met by the stimuli. In Study 1, we asked participants to estimate reward probabilities for two-outcome gambles under complete uncertainty conditions. Their probability estimates were guided by their expectations and varied as a function of reward magnitude: Smaller rewards were assumed to be more likely than larger rewards. With Study 2, we tested whether encountering a task environment that conforms with people's prior beliefs affects their information search and probability judgments. In this study, we asked participants to estimate reward probabilities of two-outcome gambles. Participants sampled outcomes from the gambles. Across participants, we manipulated the relationship between rewards and reward probabilities. In a representative condition, the correlation mimicked the ecological reward and reward

[2] We also tested variants of the Bayesian updating model, where we estimated prior probability beliefs on the individual level from the data or using the regression coefficients of the control block to estimate prior probability beliefs. None of the models outperformed the natural-mean heuristic.

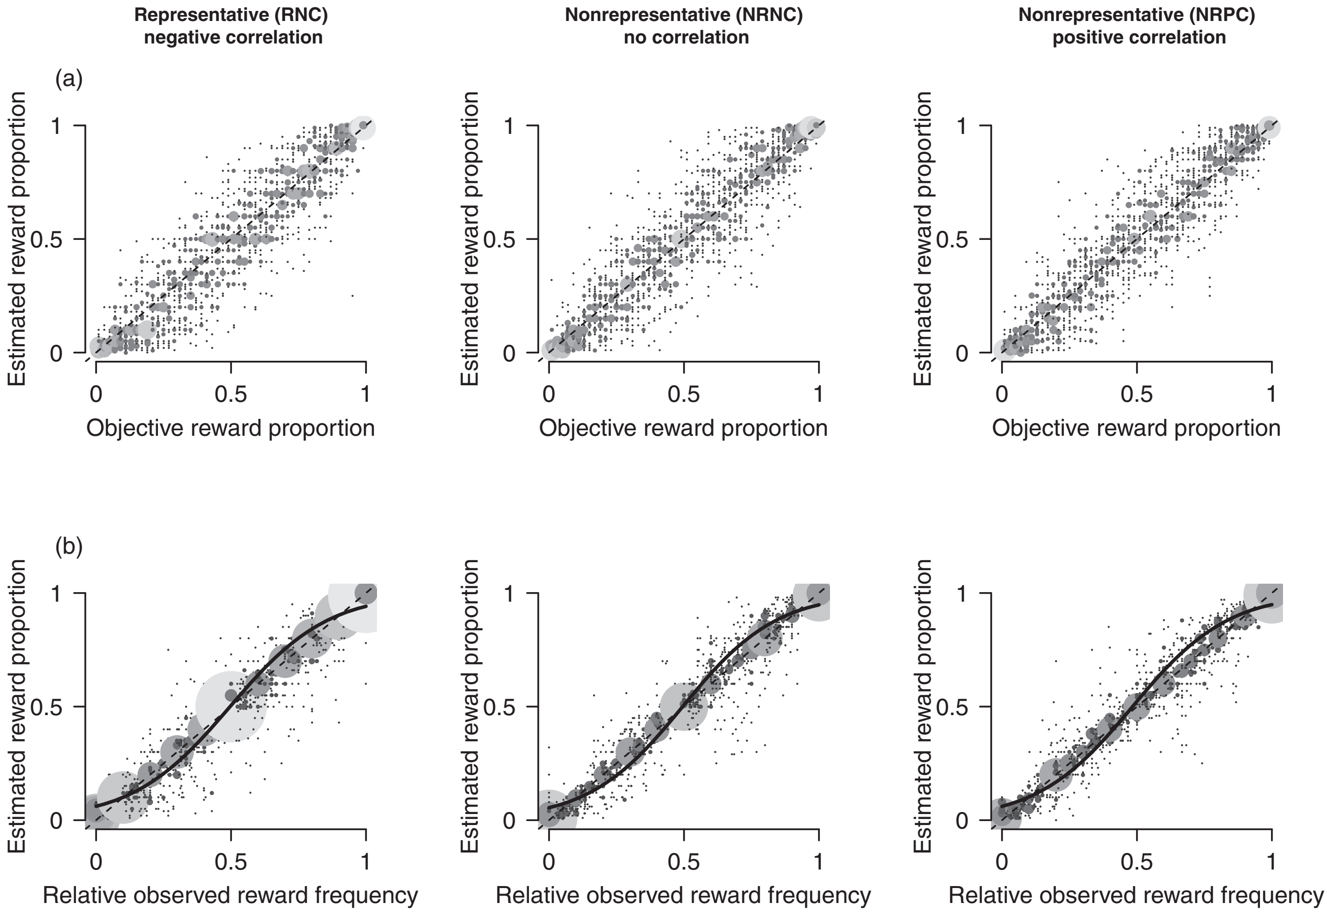

Figure 7. Estimated reward proportions in the test block as a function of (a) objective reward proportion and (b) relative observed reward frequencies. The sizes of the dots illustrate how often each response was given. The dashed lines indicate the 45° identity line. The solid lines indicate the regression line based on the median posterior group estimates of the coefficients of the beta regression.

probability relationship such that larger rewards were less likely than smaller rewards. In two nonrepresentative conditions, rewards and reward probabilities were either not correlated or positively correlated. We hypothesized that a correlation between reward and reward probabilities could influence people's search effort and how they integrate the acquired information for their final probability judgment.

In summary, participants searched for less information when rewards and reward probabilities were negatively correlated as compared to situations in which they were positively correlated or not correlated. But the way participants integrated the acquired information, that is, their judgment strategy, did not depend on different reward and reward probability relationships: They treated the outcomes that they drew as if they were representative of the true outcome distribution.

These findings contradict a full Bayesian approach, which would suggest that people integrate knowledge about how rewards and probabilities are correlated to determine how much to sample and when to estimate probabilities. However, our data suggest that participants heuristically applied a two-stage strategy: First, they sample *X* outcomes. Second, they estimate the relative reward frequency of the outcomes that they observed. But what is the mechanism that influenced the sample size *X* (i.e., the number of outcomes participants drew) and led to sample sizes lower in the representative condition than in both nonrepresentative conditions?

Does Awareness of Correlation Influence Sample Size?

Generally, people's search efforts could be influenced by any correlation between rewards and probabilities. Lejarraga et al. (2012) argued that in decisions based on experience, people learn how gambles are structured. Consequently, people's search efforts are directly influenced by their knowledge. In our experiment, participants in both correlated conditions learned the relationship between rewards and reward probabilities. This becomes evident from the results of the control block, where participants estimated reward probabilities without sampling. We found that in all conditions, their responses resembled the contingencies between rewards and reward probabilities that they had observed in the previous blocks (see Figure 8). We conclude that a general awareness that a correlation exists does not influence sample size. We argue instead that familiarity or belief in the plausibility of a correlation between rewards and probabilities influences sample size.

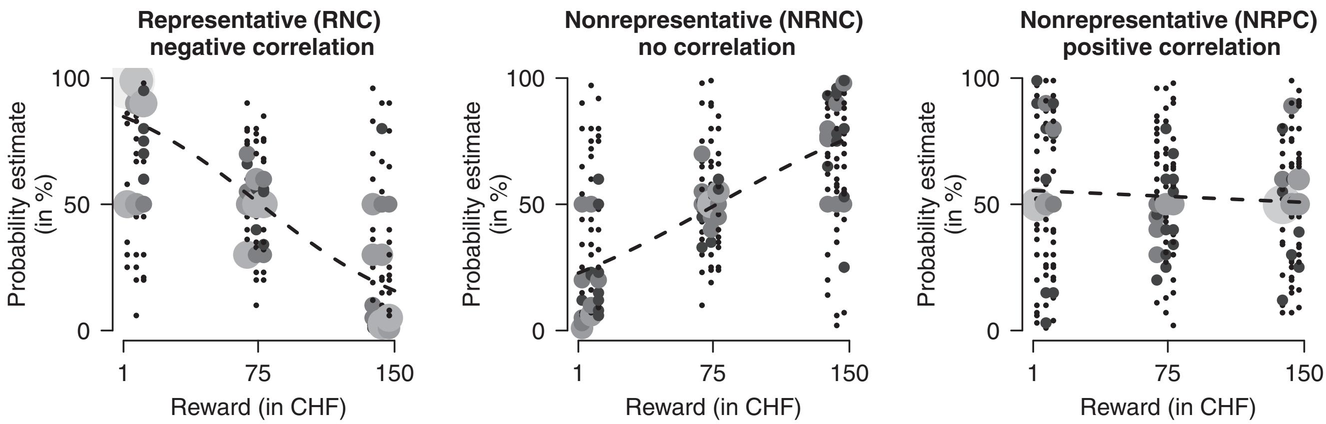

Figure 8. Probability estimates in the control block as a function of reward amount. Sizes of the dots indicate how often individual responses were given. The dashed lines represent the regression line based on the median estimates of the posterior intercept and slope distributions. CHF = Swiss francs.

Belief in the Plausibility of Correlation

Arguably, participants learned the correlation in the nonrepresentative condition as well as in the ecologically representative condition but apparently did not trust this correlation as much. The notion that larger rewards are more likely than smaller rewards could appear so counterintuitive that in each trial, participants felt that they needed to sample more outcomes to verify that the correlation still existed. This hypothesis receives support from the result that only in the representative condition did sample size decrease over the course of trials. Future research should test this hypothesis, for instance, by assessing confidence in judgments.

Conclusion and Outlook

Our study provides evidence that people exploit representative reward–probability regularities of the environment when they search for information in experience-based judgment tasks. This finding is crucial for the well-studied domain of decisions from experience (Hertwig et al., 2004). There, the expected values of the two competing gambles that participants choose between are often matched. This matching of expected values creates an environment where rewards and reward probabilities of competing gambles are negatively correlated (see Figure 1). People potentially exploit this correlation by inferring properties of one gamble from knowledge about the other gamble. Suppose a participant draws a few outcomes from Gamble A, showing small rewards, and a few outcomes from Gamble B, showing a number of zero outcomes. The person expects that rewards and probabilities are negatively correlated and correctly guesses from her observations that the nonobserved reward value in Gamble B must be relatively high. Hence, the subjective representation of the gambles may be more accurate than would be assumed based on the observed outcome frequencies which potentially influences a decision-maker's search effort and/or decisions.

More generally, our study shows an aspect of human cognition that is often overlooked: People behave differently depending on whether the research environment (e.g., stimuli) corresponds to the structure of everyday situations outside the laboratory. When one ignores people's prior beliefs about the structure of the world, one might fail to observe the ecological rationality of human cognition.

References

Alloy, L. B., & Tabachnik, N. (1984). Assessment of covariation by humans and animals: The joint influence of prior expectations and current situational information. *Psychological review, 91,* 112–149.

Brunswik, E. (1947). *Systematic and representative design of psychological experiments. With results in physical and social perception*. Berkeley, CA: University of California Press.

Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review, 62,* 193–217.

Burnham, K. P., & Anderson, D. R. (2002). *Model selection and multimodel inference: A practical information-theoretic approach*. New York, NY: Springer.

Busemeyer, J. R., Byun, E., Delosh, E. L., & McDaniel, M. A. (1997). Learning functional relations based on experience with input–output pairs by humans and artificial neural networks. In K. Lamberts & D. Shanks (Eds.), *Knowledge concepts and categories* (pp. 405–437). New York, NY: Psychology Press.

Busemeyer, J. R., & Rapoport, A. (1988). Psychological models of deferred decision making. *Journal of Mathematical Psychology, 32,* 91–134.

Camilleri, A. R., & Newell, B. R. (2013). Mind the gap? Description, experience, and the continuum of uncertainty in risky choice. In N. Srinivasan & C. Pammi (Eds.), *Progress in brain research: Vol. 202. Decision making: Neural and behavioural approaches* (pp. 55–71). Oxford, UK: Elsevier.

Dieckmann, A., & Rieskamp, J. (2007). The influence of information redundancy on probabilistic inferences. *Memory & Cognition, 35,* 1801–1813.

Fiedler, K. (2000). Beware of samples! A cognitive–ecological sampling approach to judgment biases. *Psychological Review, 107,* 659–676.

Gigerenzer, G., Hoffrage, U., & Kleinbolting, H. (1991). Probabilistic mental models—A Brunswikian theory of confidence. *Psychological Review, 98,* 506–528.

Gigerenzer, G., & Hug, K. (1992). Domain-specific reasoning: Social contracts, cheating, and perspective change. *Cognition, 43,* 127–171.

Gigerenzer, G., Todd, P. M., & the ABC Research Group. (1999). *Simple heuristics that make us smart*. New York, NY: Oxford University Press.

Griffin, D., & Tversky, A. (1992). The weighing of evidence and the determinants of confidence. *Cognitive Psychology, 24,* 411–435.

Hau, R., Pleskac, T., & Hertwig, R. (2010). Decisions from experience and statistical probabilities: Why they trigger different choices than a priori probabilities. *Journal of Behavioral Decision Making, 68,* 48–68.

Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description–experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making, 21,* 493–518.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science, 15,* 534–539.

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences, 13,* 517–523.

Hertwig, R., & Pleskac, T. J. (2008). The game of life: How small samples render choice simpler. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind: Prospects for rational models of cognition* (pp. 209–236). Oxford, UK: Oxford University Press.

Knight, F. H. (1921). *Risk, uncertainty, and profit*. New York, NY: Sentry Press.

Lejarraga, T., Hertwig, R., & Gonzalez, C. (2012). How choice ecology influences search in decisions from experience. *Cognition, 124,* 334–342.

Markowitz, H. (1952). Portfolio selection. *Journal of Finance, 7,* 77–91.

Mata, R., Schooler, L. J., & Rieskamp, J. (2007). The aging decision maker: Cognitive aging and the adaptive selection of decision strategies. *Psychology and Aging, 22,* 796–810.

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,* 328–339.

Pleskac, T. J., & Hertwig, R. (2014). Ecologically rational choice and the structure of the environment. *Journal of Experimental Psychology: General, 143,* 2000–2019.

Plonsky, O., Teodorescu, K., & Erev, I. (2015). Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychological Review, 122,* 621–647.

Pohl, R. F. (2006). Empirical tests of the recognition heuristic. *Journal of Behavioral Decision Making, 19,* 251–271.

Rakow, T., Demes, K. A., & Newell, B. R. (2008). Biased samples not mode of presentation: Re-examining the apparent underweighting of rare events in experience-based choice. *Organizational Behavior and Human Decision Processes, 106,* 168–179.

R Core Team. (2014). *R: A language and environment for statistical computing* [computer software manual]. Vienna, Austria: R Foundation. Retrieved from http://www.R-project.org/

Rieskamp, J. (2008). The probabilistic nature of preferential choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34,* 1446–1465.

Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General, 135,* 207–236.

Roth, Y., Wänke, M., & Erev, I. (2016). Click or skip: The role of experience in easy-click checking decisions. *Journal of Consumer Research, 43,* 583–597.

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance, 19,* 425–442.

Todd, P. M., & Gigerenzer, G. (2012). *Ecological rationality: Intelligence in the world*. New York, NY: Oxford University Press.

Tversky, A., & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin, 76,* 105–110.

Appendix A

Bayesian Updating Model

The Bayesian updating model describes how people can make relatively accurate probability judgments based on a few observed outcomes. Basically, the model suggests that people learn the underlying patterns of correlation between rewards and probabilities and integrate this information into their judgments. The model makes two assumptions: First, in every trial (t), people start sampling with a prior belief about the reward probability of this trial $p_{t=0}$. Second, people update this belief as they sample.

We assumed that people's prior beliefs depend on the objective correlation between rewards and probabilities and can be described as $p_{t=0} \sim \text{Beta}(a, b)$. To estimate a and b, we assumed that in both correlated conditions the mode of people's prior belief distribution is approximated by the objective reward probability. In the uncorrelated condition, we used 0.5 as an approximation of the mode of people's prior belief distribution. With these assumptions, parameter $b = N \times -p + N + 2 \times p - 1$. Here N is a free parameter that describes the strength of the prior belief. The larger that N is, the more influence the prior belief has relative to the sampled information. N has to be greater than 2. Parameter $a = N - b$. We assumed that people integrate new information as they sample. Mathematically, this can be described by adding 1 to parameter a of the beta-distributed probability belief p_t every time a person samples a reward and by adding 1 to parameter b of the beta-distributed probability belief p_t every time a person samples a zero. In the third step, people deduct an estimate of the reward probability ($P_{MB,reward,t}$). After t outcomes, this estimate equals the mean of the updated beta distribution $p_{M_B,reward,t} = \frac{a_t}{a_t+b_t}$

Natural-Mean Heuristic

We compared the Bayesian model described above to a variant of the natural-mean heuristic $M_{NM.}$ There, people's probability estimates are derived from the relative frequency of reward and zero outcomes. This M_{NM} model implies that people treat the observed outcomes as if they describe the underlying outcome probabilities comprehensively.

Mathematically, the probability judgment about the reward probability $p_{MNM,reward}$ after t observations is defined as $p_{M_{NM},reward,t} = \frac{1}{t} \times \sum_{t=1}^{t} f(t)$, where $f(t)$ is a sign function that equals 1 when the observed outcome was a reward and zero otherwise.

(*Appendices continue*)

Appendix B

Description of Bayesian Models

The following code describes the hierarchical beta regression that we used in Studies 1 and 2.

```
model {
for(i in 1:Ndata) {
y[i] ~dbeta(alpha[i], beta[i])
alpha[i] <- mu[i] * tau[subj[i]]
beta[i] <- (1-mu[i]) * tau[subj[i]]
mu[i] <- 1/(1+ exp(-1* (b0[subj[i]]+inprod(b1[subj[i]],x[i]))))
}
for(vp in 1:Nsubj){
tau[vp] ~dgamma(.001,.001)
b0[vp] ~dnorm(mub0[Csubj[vp]], taub0[Csubj[vp]])
b1[vp] ~dnorm(mub1[Csubj[vp]], taub1[Csubj[vp]])
}
for(c in 1:cond){
mub1[c] ~dnorm(0,.0001)
taub1[c] ~dgamma(.001,.001)
mub0[c] ~dnorm(.5,.0001)
taub0[c] ~dgamma(.001,.001)}
}
•Ndata = Total number of data points;
•y = Probability estimate;
•subj = Identifies to which subject each individual data point belongs;
•x = Predictor;
•Nsubj = Number of participants;
•Csubj = Identifies in which condition a subject is;
•cond = Number of between-subjects conditions.
```

The following code describes the hierarchical model of the sample-size analysis.

```
model {
for(i in 1: Ndata) {
y[i] ~dnorm(mu[subj[i]], tau[subj[i]]) T(0,)
}
for (j in 1: Nsubj) {
mu[j] ~dnorm(muG[Csubj[j]], tauG[Csubj[j]]) T(0,)
tau[j] ~dgamma(.001,.001)
}
for(c in 1:cond){
muG[c] ~dnorm(2.95,.0001)T(0,)
tauG[c] ~dgamma(.001,.001)
}
dif1 <- muG[2] - muG[1]
dif2 <- muG[3] - muG[1]
}
•Ndata = Total number of data points;
•y = Data point (sample size);
•subj = Identifies to which subject each individual data point belongs;
•Nsubj = Number of participants;
•Csubj = Identifies in which condition a subject is;
•cond = Number of between-subjects conditions.
```

(*Appendices continue*)

The following code describes the regression of sample size on trial number.

```
model {
for(i in 1:Ndata){
y[i] ~dnorm(mu[i], tau[subj[i]])
mu[i] <- b0[subj[i]]+b1[subj[i]]*trial[i]
}
for(vp in 1:Nsubj){
tau[vp] ~dgamma(.001,.001)
b0[vp] ~dnorm(mub0[Csubj[vp]], taub0[Csubj[vp]])
b1[vp] ~dnorm(mub1[Csubj[vp]], taub1[Csubj[vp]])
}
for(c in 1:cond){
mub1[c] ~dnorm(0,.0001)
taub1[c] ~dgamma(.01,.01)
mub0[c] ~dnorm(0,.0001) T(0,)
taub0[c] ~dgamma(.01,.01)}
}
•Ndata = Total number of data points;
•y = Data point (sample size);
•trial = Predictor (trial number);
•subj = Identifies to which subject each individual data point belongs;
•Nsubj = Number of participants;
•Csubj = Identifies in which condition a subject is;
•cond = Number of between-subjects conditions.
```

Received June 9, 2017
Revision received January 11, 2018
Accepted January 18, 2018 ■