Title
How the Mind Exploits Risk-Reward Structures in Decisions under Risk

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 39(0)

Authors
Leuker, Christina
Pleskac, Timothy J.
Pachur, Thorsten
et al.

Publication Date
2017

Peer reviewed

How the Mind Exploits Risk-Reward Structures in Decisions under Risk

Christina Leuker (leuker@mpib-berlin-mpg.de)
Timothy J. Pleskac (pleskac@mpib-berlin.mpg.de)
Thorsten Pachur (pachur@mpib-berlin.mpg.de)
Ralph Hertwig (hertwig@mpib-berlin.mpg.de)
Max Planck Institute for Human Development
Lentzeallee 94, 14195, Berlin, Germany

Abstract

In many natural domains, risks and rewards are inversely related (Pleskac & Hertwig, 2014). We sought to understand how people might use this relationship in choosing among risky gambles. To do so we, manipulated risk-reward structures of monetary gambles to be either negatively or positively correlated, or uncorrelated. After substantial exposure to these environments, participants completed a speeded choice task among non-dominated gambles. Eye-tracking data from this task suggests that participants often shifted their attention to mainly one attribute in the correlated conditions, in which the risk-reward relationship was present. This was an adaptive strategy that resulted in a similar proportion of expected-value maximizing choices, compared to a more compensatory processing strategy.

Keywords: risk-reward relationship; decisions under risk; attention; noncompensatory processing; adaptive cognition

Introduction

How likely is it to win the jackpot in the state lottery? Although many people play this game for a small pay-to-play fee, most of them also know that they are unlikely to win it. In fact, the larger rewards that we desire are usually unlikely to occur. While such a negative relationship between risks and rewards or probabilities and payoffs exists across gambles in many monetary and nonmonetary domains in the environment, this relationship is hardly every present in empirical studies of risky choice (Pleskac & Hertwig, 2014). In this study, we investigated how people's experience with different risk-reward relationships impact how they process explicitly stated payoffs and probabilities in *decisions under risk*. In particular, we studied how an environment in which risks and rewards are correlated would be conducive for the use of noncompensatory processing strategies, that ignore part of the attributes, in a situation where time was limited.

Adaptive Decision Making

According to an adaptive view of cognition, people exploit statistical regularities in the environment (Simon, 1956). As Payne, Bettman and Johnson (1993) found, the extent to which people exploit structures in the environment can largely depend on "the structure of the available alternatives, and [...] the presence of time pressure" (p. 534). For instance, people can decide to rely on a subset of cues in the environment because cues are often interrelated (Brunswik, 1952). Despite using a reduced amount of information, this can lead to good choices (Gigerenzer, Todd, & the ABC Research Group, 1999). Here, we propose that the risk-reward relationship is a key structure that people capitalize on to make fast, adaptive (or value-maximizing) decisions.

Choice in Risk-Reward Environments

When should and do people rely on risk-reward structures to inform their decisions? One case is when information is missing, such as in *decisions under uncertainty*, where the probabilities of obtaining a reward are unknown. In this case, Pleskac and Hertwig (2014) showed that people use a risk-reward heuristic, inferring the probability of a payoff from the magnitude of the payoff itself. In a new set of studies, we have also found that in using the risk-reward heuristic people appear to adapt to different risk-reward structures (Leuker, Pleskac, Pachur, & Hertwig, in prep.). In particular, we exposed participants to different risk-reward environments by asking them to price gambles from different risk-reward environments. Then we asked participants to choose between an uncertain prospect (where the probabilities were unstated) and a certain payoff. Participants' preferences were again consistent with them using a risk-reward heuristic, inferring probabilities from payoff magnitudes. Moreover, their preferences depended on the environment they had been exposed to previously. For example, participants in the negative condition chose the lower payoff, uncertain options more often compared to the positive condition. Based on these results, we sought to examine if and how people adapt their decision-making processes to risk-reward structures in *decisions under risk*, when payoffs and probabilities of the option are known.

The Current Study

Processing strategies. One way to distinguish between processing strategies is to consider the amount of attributes they rely on. *Compensatory* strategies process and trade off of all available and relevant information. *Noncompensatory* strategies "typically reduce processing demands by ignoring potentially relevant information" (Payne, Bettman, & Johnson, 1988). Thus, one important reason to consider noncompensatory processing strategies (despite information being, in principle, available, as in risky choice) is when time or cognitive resources are limited.

Strategy-environment dependence. Early research on these two classes of strategies demonstrated that their success largely depends on the environment in which they are recruited. Specifically, in environments with nondominated options (e.g., gamble *A* offers a higher payoff x, but gamble *B* offers a higher probability p: $x_A > x_B$ and $p_A < p_B$), people should rely on compensatory strategies (see Table 2, Payne et al., 1988). A decision maker who processes the dimensions

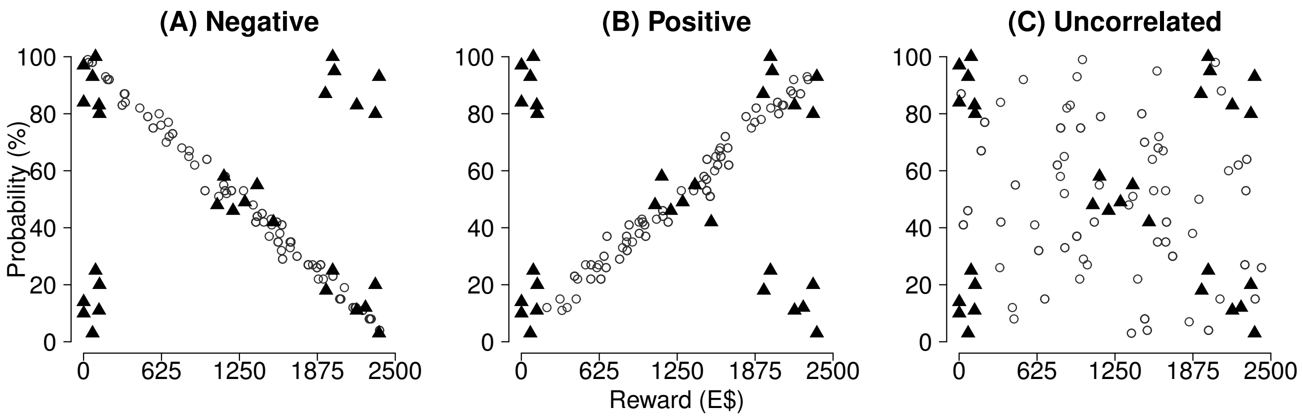

Figure 1: Choice stimuli based on their relationship between probabilities and payoffs. Each point depicts one gamble from the choice phase. Across conditions, probabilities and payoffs were (A) negatively correlated, (B) positively correlated or (C) uncorrelated. Black circles are environment gambles (60 pairs). Triangles are common gambles interspersed in all three conditions (15 pairs). Dominated options not depicted (5 pairs).

in a noncompensatory fashion in these environments—for instance by relying on a simplifying heuristic that attends to outcomes only—will suffer "a substantial loss in accuracy" (Payne et al., 1993, p. 539). In contrast, such noncompensatory processing strategies have been shown to perform well when dominance is possible (that is, one gamble is better on all dimensions: if $x_A > x_B$ and $p_A > p_B$).

Local vs. global environment. By definition, nondominated options create an inverse risk-reward relationship in a given set of alternatives, because the gamble offering a higher payoff will always be associated with a lower probability *relative* to the other gamble ($x_A > x_B$ and $p_A < p_B$). However, this "local" risk-reward relationship (within a pair of gambles) can differ from a "global" risk-reward structure (across a larger reference class of gambles). That is, nondominated alternatives can be drawn from globally structured or unstructured environments. We propose that both the use and performance of either type of strategy is also highly dependent on these global risk-reward structures. Global correlations between risks and rewards make one of the cues redundant (payoffs predict probabilities and vice versa). Therefore, we hypothesized that, when options are drawn from correlated risk-reward environments, noncompensatory strategies can lead to accurate, expected-value maximizing choices even if neither option is dominated. For choices between nondominated options from globally uncorrelated environments, results may resemble those of Payne, Bettman and Johnson (1988).

To test these ideas, we employed a between-subjects design manipulating the global risk-reward relationship between the possible options participants experienced (Figure 1). In a first pricing phase, we showed participants individual gambles and asked them to state their willingness to sell each gamble. We used this phase to expose people to different risk-reward environments. Detailed data from this phase will be reported elsewhere. Our focus in this paper is the second phase, where participants chose between pairs of risky options under moderate time pressure (Figure 2). The gambles in the choice phase were drawn from the same, condition-dependent risk-reward environments, and paired such that neither option was dominated. We tracked participants' eye movements to dissociate between processing strategies across the different risk-reward environments, as choice patterns alone may not be sufficient to do so. As an independent test of whether participants had picked up the different risk-reward relationships, we asked them to estimate probabilities from payoffs at the end of the experiment.

Method

Participants

Ninety-three (55 female) participants (mean age = 25.6 yrs, *SD* = 3.7; *N* = 31 per condition) from the participant pool at the Max Planck Institute for Human Development, Berlin, completed the experiment (duration ∼75 min). All participants were paid a fixed rate of €12 plus a bonus based on their performance in a random subset of trials from the pricing phase and choice task (€3.53-11.67).

Design

The experiment consisted of three phases. In the pricing phase, participants were presented with single gambles and asked to indicate their willingness to sell for each of them. Between subjects, we manipulated the types of gambles people were presented with such that payoffs and probabilities were positively or negatively correlated, or uncorrelated. In the subsequent choice task, these different risk-reward structures were maintained. People were asked to choose between gamble pairs within 3s. All gambles were in the gain domain ("p_1 chance of winning x_1, otherwise nothing"). We used an experimental currency, the E$ (conversion rate 2500E$ = €1, disclosed in the instructions). We collected eye-tracking data during the exposure phase and the choice task. As people are merely *exposed* to different risk-reward structures, participants picking up risk-reward structures despite not being told about the presence of *any* relationship in the data would constitute a form of unsupervised learning. Finally, in the third phase we asked participants to estimate the probabilities they thought were associated with various payoff levels. We did this to test whether participants had picked up the different risk-reward structures from the gambles they were exposed to throughout the study. Participants were not informed about the estimation task beforehand.

Gamble environments. The gambles from the pricing and choice phases were constructed such that across gambles, there was a negative, a positive, or no relationship between risks and rewards. For the negative condition, we drew random payoffs from a uniform distribution (range 1.01 – 2500E\$). The probabilities for each payoff were inversely related to the payoff x such that, $p = 1 - \frac{x}{2500}$. We added normally distributed noise to logit-transformed payoffs and probabilities. For the positive condition, we reversed the order of probabilities such that the highest probabilities were now associated with the highest rewards (and vice versa). For the uncorrelated condition, we re-linked payoffs and probabilities randomly.

Pricing task. The pricing task served to expose participants to different risk-reward environments. Briefly, participants were shown each of the 90 gambles from one of the environments and asked to state a price they would be willing to sell the gamble for. In addition to 90 condition-dependent gambles based on the aforementioned construction rule, participants were also asked to price 30 gambles that were common to each of the three conditions (triangles in Figure 1), yielding 120 gamble stimuli per condition. To motivate participants to report their true valuations of the gambles, we implemented a Becker-DeGroot-Marschak auction (Becker, Degroot, & Marschak, 1964). In particular, ten gambles were selected at the end of the experiment and participants either played out the gamble or received their stated selling price.

Choice task. Gambles were created using the same construction rule as above. An initial set of 100 gambles yields 4950 possible gamble pairs. We randomly drew 60 non-dominated gamble pairs per condition (see circles in Figure 1). By design of the study, expected value differences were largest in the uncorrelated and smallest in the positive condition (uncorrelated: $Md = 173E\$$, .53 - 1374E\$; negative condition: $Md = 134E\$$, .49 - 511E\$; positive condition: $Md = 23E\$$, .43 - 146E\$). In addition, we interspersed 15 gamble pairs that were common to each of the conditions in the second half of the choice task (triangles in Figure 1), and 5 choices with dominated options as catch trials, yielding 80 choices in total. Common gambles allowed us to examine condition-dependent processing differences on precisely the same stimuli. Participants were instructed to choose their preferred gamble within 3s (see Figure 2). Crucially, participants were informed in this task that the gambles were drawn from the same population of gambles they had experienced in the previous pricing task. Five randomly selected choices were played out at the end of the experiment.

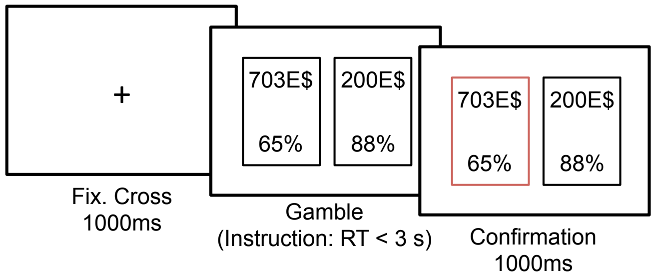

Figure 2: Typical choice trial. For trials that exceeded the time limit, we presented an additional screen informing participants that they had lost payoffs in that particular trial (not shown). Eye-tracking data was recorded throughout, analyses are based on the second screen.

Estimation task. We drew 20 payoffs (range 1.01 – 2500E\$) and asked participants to estimate the probabilities that had been associated with these payoffs in the main experiment.

Eye-tracking

During the pricing and choice tasks, we collected binocular eye position data with an EyeTribe tracker, sampled at 60Hz. The experiment was implemented in PsychoPy 1.83.01 and the eye-tracking interface PyTribe (Dalmaijer, Mathôt, & Van der Stigchel, 2013). Each participant's eye movements were calibrated using the Eyetribe UI with a 9-point grid before each task (< 0.7). Participants were seated approximately 60 cm from the screen using a chinrest affixated to the table, in a room with negligible ambient light. We preprocessed raw samples by parsing eye-tracking data into fixations and saccades using the saccades package in R (*Saccades Version 0.1-1*, 2015), based on a velocity-based algorithm (Engbert & Kliegl, 2003). Eye-tracking analyses in this paper are based on fixation data.

Analysis

The data were analyzed using Bayesian General Linear Models using Stan in R for regression analyses (*RStanArm Version 2.9.0-4*, 2016). We ran 3 chains (2500 samples each, burn-in of 500), and investigated (convergence of) our posteriors visually and with the Gelman-Rubin statistic (Gelman & Rubin, 1992). We report the mean of the posterior distribution of the parameter of interest and two-sided 95% equal tail credible intervals (CI) around each value.

Results

Behavioral

We excluded one participant in the negative condition who chose the dominated option in 4 out of 5 catch trials.

Pricing task. For all participants, prices were strongly related to the expected values of the gambles (credible payoff $\times$ probability interaction, $b = .70$, CI = [.66, .74]).

Choice task. Participants across all three conditions chose the expected-value-maximizing options above chance level ($M = .71$, CI = $[.56, .85]$). As expected, in the positive condition (in which the EV differences between the options were rather small) participants made fewer EV-maximizing choices ($M = .59$, CI = $[.51, .67]$) than in the uncorrelated condition (in which the EV differences were larger; $M = .70$, $b = .11$, CI = $[.00, .22]$), and the negative condition ($M = .74$, $b = .25$, CI = $[.14, .37]$). Controlling for EV differences and individual variation, participants in the negative condition achieved a higher proportion of expected-value maximizing choices ($M = .70$, CI = $[.52, .88]$) compared to the uncorrelated condition ($M = .31$, $b = -.39$, CI = $[-.63, -.16]$), and the positive condition ($M = .39$, $b = -.32$, CI = $[-.54, -.09]$). In both models, the highest accuracy was achieved in the negative condition. In the subset of gambles that were common across all conditions, there were no differences in accuracy between the conditions ($M = .53$, CI = $[.35, .72]$).

Response times were comparable across all conditions and gamble types. In addition, small proportions of timed-out trials (negative: .006, positive: .016, uncorrelated: .013) indicate that participants were well-adjusted to the speed instruction of 3s (*Md* = 1.63s even suggest that people could have taken more time on many trials).

Estimation task. Participants' probability estimates reflected the risk-reward structure they had been exposed to previously. That is, participants in the negative condition provided *lower* probability estimates for gambles with higher payoffs ($b = -.64$, CI = $[-.68, -.60]$, % per 100 E$), and in the positive condition participants provided *higher* probability estimates for gambles with higher payoffs ($b = .16$, CI = $[.10, .14]$). In the uncorrelated condition, participants provided lower probability estimates for gambles with higher payoffs (weaker slope compared to the negative condition: $b = -.32$, CI = $[-.37, -.26]$).

Eye-tracking

We defined four areas-of-interest (AOIs), one for each payoff and probability. We visually inspected the quality of every participant's eye-tracking data by plotting their fixations over time. Seven participants whose fixations did not map onto the screen correctly were excluded, a possible result of the eye-tracker being moved during the experiment. We excluded one further participant who was blind in one eye, leaving $N = 84$ for the eye-tracking analyses (27, 29, and 28 in the negative,

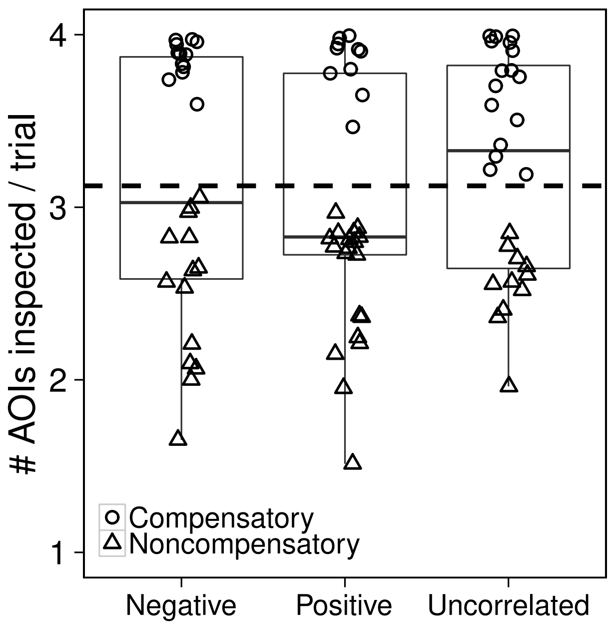

Figure 3: Number of AOIs inspected per condition. Each dot represents one participant's average number of AOIs visited per trial. Dashed line = mean number of AOIs visited across participants. Differences between conditions are driven by the composition of compensatory/noncompensatory strategies (percent compensatory in negative: 44%; positive: 34%, uncorrelated: 61%).

positive, and uncorrelated conditions, respectively).

Number of AOIs viewed. To test whether the presence of a risk-reward relationship led to more noncompensatory processing, we averaged the number of AOIs each participant viewed (max. 4). Participants in the uncorrelated condition inspected the largest number of AOIs ($M = 3.46$, CI = [3.45, 3.48]). Participants in the positive condition inspected a credibly lower number of AOIs ($M = 3.20$, $b = -.20$, CI = $[-.23, -.18]$). The negative condition also inspected credibly fewer AOIs, but the difference was smaller ($M = 3.40$, $b = -.06$, CI = $[-.08, -.03]$, model on the trial level). Using the average number of fixations as an alternative indicator resulted in the same pattern of results, and only a marginally higher count (uncorrelated condition $M = 3.71$, negative condition $M = 3.51$, positive condition $M = 3.39$), likely because the time limit imposed in the experiment did not allow for many re-acquisitions (i.e., fixations back to a previously acquired AOI). Note that the mean number of fixations is rather low (i.e. < 4). We ran the same AOI model using only common gamble data (see triangles in Figure 1). Again, the uncorrelated condition inspected most AOIs ($M = 3.21$, CI = [3.09, 3.34]). This number was lower in the positive condition ($M = 2.83$, $b = -.38$, CI = $[-.57, -.20]$, difference credible), and in the negative condition ($M = 3.07$,

$b = -.14$, CI = $[-.34, .04]$, difference however not credible). Because these gambles were identical across conditions, this suggests condition-dependent processing strategies are not merely a by-product of specific risk-reward environments that vary on crucial dimensions such as EV differences between gambles. Figure 3 suggests substantial individual differences among participants in the conditions (indeed, differences in numbers of AOIs inspected can be accounted for by including participant as a grouping factor). More importantly, however, visual inspection of the data suggests two subgroups that can roughly be split by the mean number of AOIs inspected across participants ($M = 3.11$, dashed line in Figure 3): Participants who tend to inspect all four AOIs ("compensatory") and participants who ignore some of the AOIs ("non-compensatory"). Thus, differences between conditions may be driven by the composition of compensatory/noncompensatory strategies (proportion compensatory in uncorrelated: .61, positive: .34, negative: .44). That is, participants in the uncorrelated condition were 2.85 times more likely to rely on a compensatory strategy than participants in the positive condition ($b = 1.05$, CI = $[.07, 2.04]$). The difference between the negative and positive risk-reward environments was not credible ($b = .68$, CI = $[-.29, 1.73]$, OR = 1.97). A majority of participants in the correlated environments thus seemed to rely on a noncompensatory strategy (note that here such a strategy could also mean attending to three out of four AOIs per trial, see Figure 3).

Attention to attributes. Which attributes did participants attend to, especially when choosing to ignore some of the information? All participants fixated most on payoff information (.57 of fixations, CI = [.53, .61]). This proportion decreased for participants who inspected more AOIs ($b = -.98$, CI = $[-1.56, -.38]$; no credible effect of condition). At the extreme end, participants who, on average, inspected roughly two AOIs fixated on the payoff 80% of the time. An alternative viable noncompensatory strategy would have been to focus more on the information presented at the top or the bottom of the screen. We counterbalanced the location of attributes (between-participants). However, top/bottom fixation proportions were unrelated to the number of AOIs inspected ($b = .03$, CI = $[-.60, .66]$), suggesting that participants considered payoff information as more relevant when using noncompensatory strategies.

EV choices by strategy. Do compensatory or noncompensatory strategies differ in performance within the three environments? Figure 4 shows that users of a noncompensatory strategy (triangles) achieved similar levels of EV-maximizing choices compared to users of a compensatory strategy (circles), overall ($M = -.10$, CI = $[-.39, .18]$). Unexpectedly, this held irrespective of condition (no credible strategy $\times$ condition interaction). This result also held when controlling for differences between the condition in EV difference between the options. We expected that in the uncorrelated condition, decision performance would be compromised for users of a noncompensatory strategy—who should lack critical information to determine the EV maximizing option.

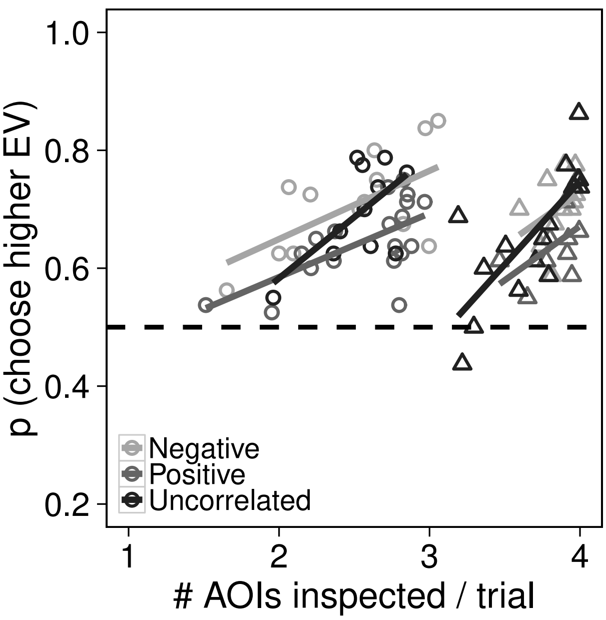

Figure 4: Proportion of higher EV choices by a participant's average number of AOIs inspected per trial, condition, and decision strategy. Circles = noncompensatory strategy users, triangles = compensatory strategy users. The gain in EV-maximizing choices associated with inspecting more AOIs is more pronounced for participants in the uncorrelated condition.

At the same time, across conditions, the proportion of EV-maximizing choices was higher for participants who inspected more AOIs (main effect of AOIs inspected on EV choice irrespective of processing strategy, $b = .044$, CI = [.001, .089]). Within each subgroup, the increase in EV choices with increasing numbers of AOIs inspected is more pronounced for participants in the uncorrelated condition (see black regression line, Figure 4). Yet, this interaction effect is not credible, potentially due to the small number of participants in each subgroup (AOI $\times$ condition interaction with the positive condition as a reference; compensatory: $b = .02$, CI = [-.04, .09]), noncompensatory: $b = .05$, CI = [-.01, .11]). In general, one would indeed expect that the increase in the proportion of EV choices with higher number of AOIs inspected is more pronounced for the uncorrelated condition because this condition allows for least simplification.

Discussion

Risk-reward relationships allow people to make fast, value-maximizing decisions. A majority of people exposed to correlated risk-reward structures used noncompensatory processing strategies, likely as a result of time pressure. With fewer AOIs inspected, participants focused more on payoff information. In turn, most people who experienced an uncorrelated risk-reward environment attempted to take into account all attributes, speaking in favor of a more compensatory processing strategy. This strategy use is adaptive given the affordances of the different environments. While correlated risk-reward environments made one of the attributes redundant, such a relationship did not exist for gamble problems in uncorrelated risk-reward environment. Condition-dependent processing differences (i.e., numbers of AOIs visited) persisted when restricting the analysis to a common set of gambles interspersed in each of the conditions. Surprisingly, these differences only had a minor impact on EV choice.

Earlier research suggested that noncompensatory strategies fare well when dominance is possible, but not when neither option is clearly dominated (Payne et al., 1988). We identify one qualification of this prediction, showing that noncompensatory processing strategies can also perform well for non-dominated option pairs; namely when a risk-reward relationship is present in the global set that gamble pairs are drawn from. Researchers have studied the influence such contextual factors before. Birnbaum (1992) found that participants' certainty equivalents for gambles were larger when a set of certainty equivalents to choose from was positively skewed (vs. negatively skewed). In addition, the *marginal* distributions of payoffs, probabilities and delays can account for psycho-economic functions that are often described in the literature (Stewart, Chater, & Brown, 2006). Here, we extend such considerations by manipulating the *joint* distribution of payoffs and probabilities.

Several limitations of the current study should be mentioned. First, it is currently unclear what underlies the strong individual differences in noncompensatory/compensatory strategy use in each condition. Potentially, some participants did not perceive a time limit of 3s as pressing enough to opt for noncompensatory strategies, or turned to different simplification strategies. Overall, the dichotomous distinction between compensatory and noncompensatory processors may be too simplistic: For instance, some individuals attended to three attributes on average (i.e., more than one class of attributes such as two payoffs, one probability). Another possibility is that users of a noncompensatory strategy fixated on some but glanced the other attributes (covert attention), or changed strategies across trials. Lastly, more research is needed to study the process by which people learn about different risk-reward structures (Klayman, 1988).

Conclusion

People's choices and processing strategies are impacted by the risk-reward structure in a given environment. Specifically, correlated risk-reward environments allow decision makers to use noncompensatory strategies when they need to reduce processing demands. This strategy use is adaptive, given that it does not need to compromise accuracy if it matches the environment. Many natural environments exhibit an inverse relationship between payoffs and probabilities that can thus be exploited in a similar way, when time or cognitive resources are limited. These findings challenge theories of decision making under risk, that often treat payoffs and probabilities as independent attributes determining the value of an option. In comparison, an adaptive decision maker may often have good ecological reasons to process payoffs and probabilities dependently.

Acknowledgments

This research was supported by the MaxNetAgeing Research School (CL). We would like to thank Chantal Wysocki and Lisa-Marie Jirak for assistance with data collection.

References

Becker, G. M., Degroot, M. H., & Marschak, J. (1964). Measuring utility by a single response sequential method. *Behavioral Science*, *9*(3).

Birnbaum, M. H. (1992). Violations of Monotonicity and Contextual Effects in Choice-Based Certainty Equivalents. *Psychological Science*, *3*(5), 310–314.

Brunswik, E. (1952). The conceptual framework of psychology. *Journal of Consulting Psychology*, *16*(6), 475–475.

Dalmaijer, E. S., Mathôt, S., & Van der Stigchel, S. (2013). PyGaze: An open-source, cross-platform toolbox for minimal-effort programming of eyetracking experiments. *Behavior Research Methods*, 1–16.

Engbert, R., & Kliegl, R. (2003). Microsaccades uncover the orientation of covert attention. *Vision Research*(9), 1035–1045.

Gelman, A., & Rubin, D. B. (1992). Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*, *7*(4), 457–511.

Gigerenzer, G., Todd, P., & the ABC Research Group. (1999). *Simple heuristics that make us smart*.

Klayman, J. (1988). On the how and why (not) of learning from outcomes. *Advances in Psychology*, *54*, 115 - 162.

Leuker, C., Pleskac, T., Pachur, T., & Hertwig, R. (in prep.). *Risk and reward: Exploiting the environments risk-reward structures in decisions under uncertainty.*

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*(3), 534–552.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge University Press.

Pleskac, T. J., & Hertwig, R. (2014). Ecologically rational choice and the structure of the environment. *Journal of Experimental Psychology: General*, *143*(5), 2000–2019.

RStanArm version 2.9.0-4. (2016). `cran.r-project.org/web/packages/rstanarm/`.

Saccades version 0.1-1. (2015). `cran.r-project.org/web/packages/saccades/`.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63*(2), 129–138.

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, *53*(1), 1–26.