Too Good To Be True? Psychological Responses to Uncommon Options

in Risk–Reward Environments

Christina Leuker, Thorsten Pachur, Ralph Hertwig

Max Planck Institute for Human Development

Timothy J. Pleskac

University of Kansas

Max Planck Institute for Human Development

19 October, 2018

Word count: 7273

Author Note

Correspondence concerning this article should be addressed to Christina Leuker (E-mail: leuker@mpib-berlin.mpg.de). Stimuli, data, and analyses can be retrieved via the Open Science Framework (view-only link for review). *Version 2018/10/19. Please do not copy or cite without author's permission.*

Abstract

Risks and rewards, or payoffs and probabilities, are inversely related in many choice environments. We investigated people's psychological responses to uncommon combinations of risk and reward that deviate from learned regularities (e.g., options that offer a high payoff with an unusally high probability) as they evaluated risky options. In two experiments ($N = 183$), participants first priced monetary gambles drawn from environments in which risks and rewards were negatively correlated, positively correlated, or uncorrelated. In later trials, they evaluated gambles with uncommon combinations of risk and reward—that is, options that deviated from the respective environment's risk–reward structure. Pricing, response times, and (in Experiment 2) pupil dilation were recorded. In both experiments, participants took more time when responding to uncommon compared to foreseeable options or when the same options were presented in an uncorrelated risk–reward environment. This result was most pronounced when the uncommon gambles offered higher expected values compared to the other gambles in the set. Moreover, these uncommon, high-value options were associated with an increase in pupil size. These results suggest that people's evaluations of risky options are based not only on the options' payoffs and probabilities, but also on the extent to which they fit the risk–reward structure of the environment.

Too Good To Be True? Psychological Responses to Uncommon Options

in Risk–Reward Environments

On December 25th, 2017, Nicole Coggins and her mother-in-law bought 100 lottery tickets for the South Carolina Holiday Cash Add-A-Play lottery for $1 each. Astonishingly, every ticket was a winner, giving them a total of $18,000 in winnings. Wade Crenshaw, who was working behind the cash register at a convenience store that day, noticed that more and more people were asking to buy Add-A-Play tickets. He later said, "It was weird, everybody winning so much. I didn't know if they were doing some kind of Christmas special" (Fortin, 2017). If you think this sounds too good to be true, you are right. Except for a select few individuals who were able to collect a total windfall of about $1.7 million, when most of the lucky winners went to cash in their prizes, the machine deemed their tickets invalid (Fortin, 2018; Harrington, 2017). As it later emerged, the many winning lottery tickets were not a Christmas special; they were due to a computer glitch.

As the cashier's reaction illustrates, people are usually aware that higher rewards are more unlikely than smaller rewards, a regularity present across many monetary and nonmonetary choice environments (Pleskac & Hertwig, 2014). This link between risks and rewards (or probabilities and payoffs) affords people an ecological structure that they may use when making decisions (Brunswik, 1943), including *decisions under uncertainty*, when they know the possible payoffs of an option but not their probabilities. In decisions under uncertainty, people exploit the risk–reward relationship to infer the probabilities from the magnitude of the payoffs (Leuker et al., 2018; Pleskac & Hertwig, 2014). They may also use the relationship to infer payoffs from probabilities (Skylark & Prabhu-Naik, 2018), or to decide how often they sample in decisions from experience (Hoffart et al., 2018). In contrast to earlier research, our focus in the present article is on *decisions under risk*—that is, decisions in which both payoffs and probabilities are explicitly stated. Specifically, we examine if and how an environmental relationship between risks and rewards results in options being perceived as uncommon in decisions under risk.

Studying people's sensitivity to the ecological structure in decisions under risk is interesting for several reasons. In risky choice, an option is characterized by two attributes, risks and rewards. Detecting uncommon combinations of risk and reward in these environments may be harder than, for instance, detecting uncommon risky options based on only one dimension (e.g., uncommon, higher-than-usual payoff magnitudes). In risk–reward environments, many different combinations of probabilities and payoffs are possible and consistent within a particular risk–reward structure. For instance, $10 with $p = .9$; $50 with $p = .5$; and $90 with $p = .1$ are all in line with a negative risk–reward structure—that is, an environment in which risks and rewards are inversely related. Conversely, $10 with $p = .1$; $50 with $p = .5$; and $90 with $p = .9$ are all in line with a (fairly nonrepresentative) positive risk–reward structure—that is, an environment in which risks and rewards are positively related. An option is uncommon if the payoff–probability combination mismatches a given environment. This also means that an option that is uncommon in one environment can be foreseeable in another environment (e.g., $90 with $p = .9$ is uncommon in a negative risk–reward environment, but foreseeable in a positive one). Furthermore, since risky choice options have an expected value (EV), options that deviate from a risk–reward structure can either offer an uncommon, high-value combination of risks and rewards (e.g., $90 with $p = .9$ in a negative risk–reward environment) or an uncommon, low-value one (e.g., $10 with $p = .1$, in the same environment). Lastly, as these examples show, detecting an uncommon option in risk–reward environments does not require feedback about whether or not an outcome is obtained (as is the case for surprise in, e.g., reward–prediction or risk–prediction errors; see Schultz et al., 1997; Bossaerts, 2010; O'Neill & Schultz, 2013; Preuschoff et al., 2008).

It is currently unclear whether the combination of risks and rewards leaks into how people make decisions under risk in a given choice environment. One reason for this is the way in which risky decisions from description are typically studied: In studies of risky choice, probabilities and payoffs across gambles and gambling problems are typically

factorially combined (risks and rewards are usually only inversely related within a gamble problem Pleskac & Hertwig, 2014). Another, related, reason is that prominent theories of risky choice tend to assume that risks and rewards are treated as independent attributes that determine the subjective value of an option, and ultimately choice (von Neumann & Morgenstern, 1944; Tversky & Kahneman, 1992). Yet there is a growing body of evidence that people are sensitive to the choice ecologies in which they make decisions (Birnbaum, 1992; Ludvig et al., 2014, 2018; Walasek & Stewart, 2015). Relatedly, people's evaluations of payoffs or probabilities can depend on the marginal distributions of payoffs or probabilities they experienced (Stewart et al., 2006). It is unclear whether the joint distribution of payoffs and probabilities (i.e., the risk–reward relationship) affects how subsequent risky options are evaluated—and if so, what the traces of this influence are. Whereas previous research has found that people exploit risk–reward structures in decisions under uncertainty (Leuker et al., 2018), in what ways do such structures leave traces when people make decisions under risk?

## Overview of Experiments and Hypotheses

Here, we present two experiments in which we investigated how participants responded to uncommon combinations of risk and reward. Between participants, the structure of the environment was manipulated to be negatively correlated, positively correlated, or uncorrelated (Figure 1). In both experiments, participants indicated the price at which they would be willing to sell a monetary gamble of the form "p chance of winning x, otherwise nothing."

There were two classes of gambles. *Environment gambles*, denoted as black circles in Figure 1, defined the structure of the environment in a given condition. *Test gambles*, denoted as filled symbols in Figure 1, were common to all conditions and used to test our hypotheses. Depending on the environment in which they were interspersed, the test gambles belonged to one of four groups. A first subset of the test gambles were *uncommon*

gambles (red triangles). These gambles were inconsistent with the risk–reward structure when it was present. A second subset of gambles were *foreseeable gambles* (blue triangles). These gambles were not only consistent with the risk–reward structure when it was present (in negative or positive risk–reward environments), but they could, in principle, even be foreseen or expected. Third, test gambles that were neither uncommon nor foreseeable (in the uncorrelated condition) are referred to as *reference gambles* (gray circles). Fourth, *average gambles* appeared in each environment and were in the midrange of the payoffs and probabilities. Because these gambles fit all risk–reward environments equally well, and had exactly the same characteristics (EV distributions, payoff magnitudes, risk-reward relationship), they were used as control stimuli to examine condition-dependent differences. Participants were not explicitly instructed to pay attention to the underlying risk–reward structure in either experiment. We examined how pricing and response times (RTs) varied as a function of gambles being uncommon, relative to foreseeable gambles in the positively and negatively correlated conditions and to reference gambles in the uncorrelated condition. In Experiment 2, in addition to the behavioral responses, we tracked pupil size in response to uncommon gambles.

All three outcome measures—pricing, RT, and pupil size—may differ depending on whether a gamble is common or foreseeable in an experimental context. First, prices may deviate from gambles' EVs in the direction of the learned risk–reward structure: If high payoffs commonly co-occur with low probabilities, prices given for high-payoff/high-probability options may undershoot the gambles' EVs. Conversely, when high payoffs commonly co-occur with high probabilities, prices given to high-payoff/low-probability options may overshoot the gambles' EVs.

Second, increasing familiarity with an environment's risk–reward structure should accelerate the processing of subsequent foreseeable options and decelerate the processing of subsequent uncommon options. This pattern is consistent with findings in the domain of event sequence learning (i.e., responses to stimuli in different locations that follow a
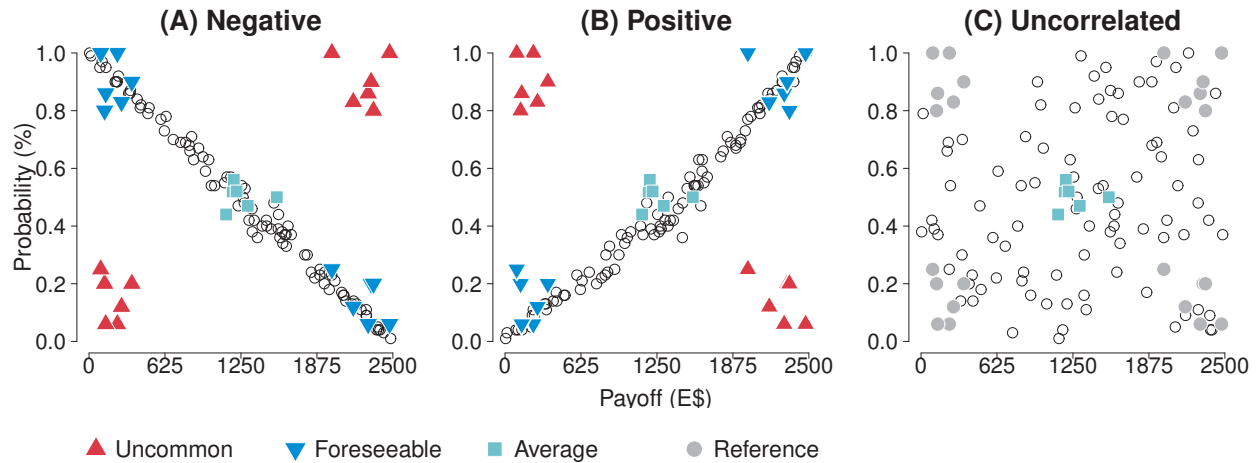
Figure 1. Prototypical gambles used in the two experiments. Environment gambles (black circles) were drawn from one of the three risk–reward environments. A set of test gambles, which was common to all three conditions, was randomly interspersed after two thirds (Experiment 1) or one third (Experiment 2) of the trials. Test gambles are color-coded by gamble type: Test gambles shown in blue downward triangles were consistent with the risk–reward relationships in a condition and could therefore be foreseen ("foreseeable gambles"). Test gambles shown in red upward triangles were inconsistent with risk–reward relationships in a condition and were therefore uncommon ("uncommon gambles"). Test gambles shown in light gray (panel C) served as "reference gambles". Here, participants were unlikely to have any expectations about particular risk–reward relationships. Test gambles shown in green squares (all panels) fit each condition equally well and were therefore used as a control ("average gambles").

sequence, e.g., 4–3–2–1–4–3–2–1). Here, longer RTs for stimuli inconsistent with a learned sequence are taken as direct evidence that a sequential stimulus structure has been learned (Rüsseler & Rösler, 2000). Moreover, prior research found that people may have a mechanism in place that "prevent[s] impulsive responding due to the presence of high value options" (Cavanagh et al., 2014, p. 2; see also Frank et al., 2007). By extension, participants may respond more slowly to uncommon, high-value options than to uncommon, low-value ones. Note that such a mechanism would also lead to RTs increasing as a function of absolute value, rather than as a function of gambles being uncommon (i.e. across conditions).

Third, pupil size may increase in response to uncommon options. Paradigms investigating feedback-based "risk-prediction errors" (Preuschoff et al., 2011) showed that pupil dilation increases in response to surprise in risky choice (an a priori unlikely option

that was eventually obtained, or an a priori likely option that was eventually not obtained). Pupil dilation has also been found to increase in response to high-value options ("win-win trials" in Cavanagh et al., 2014) in general. Here, the EV of an option may increase pupil dilation in addtition to an option being uncommon; or independent of an option being uncommon.

Experiment 1: Behavioral Responses To Uncommon Risk–Reward Combinations

Experiment 1 was a behavioral experiment in which risks and rewards were presented simultaneously. This experiment is a reanalysis of the training data of an earlier experiment (Leuker et al., 2018, Experiment 2). In earlier analyses, our focus was on decisions under uncertainty (test data); here we consider the impact of risk–reward structures on decisions under risk when combinations of risk and reward are uncommon.

Method

Participants.   The sample included 90 participants (53 females, age 24.7 years, $SD = 4.1$ years, proportion students $= .72$) from the participant pool maintained at the Max Planck Institute for Human Development. Each participant completed the experiment (duration 65 min) in exchange for a show-up fee of €10 and a performance-based bonus (€1.99–€7.82).

Stimuli.   Participants priced monetary gambles of the form "p chance of winning x, otherwise nothing." All payoffs were expressed using an experimental currency, $E\$$, with a disclosed conversion rate of $2500E\$ = €1$). We used an experimental currency to minimize the impact of outside norms associated with specific currencies on the experiments. For the condition with a negative correlation between risk and reward, the gambles were constructed as follows: 150 random payoffs were drawn from a uniform distribution with the range $1.01 - 2500$. The probabilities for each payoff were set such that they were inversely related to the payoff x ($p = 1 - x/2500$). We jittered payoffs and probabilities by

adding normally distributed noise with a standard deviation of 0.1 to both the logit transformation of the probabilities and the logit transformation of normalized payoffs. We then transformed those perturbed values back to the scales used in the experiment. For the condition with a positive correlation between risk and reward, we used the same gambles as in the negative condition but reversed the order of the probabilities. For the uncorrelated condition, we randomly linked probabilities and payoffs. This approach controlled for the marginal distribution of payoffs and probabilities across conditions (as marginal distributions of payoffs and probabilities may influence choice; see Birnbaum, 1992; Stewart et al., 2006).

In addition to these 150 condition-dependent environment gambles, participants also priced 22 gambles that were identical across all three conditions, yielding 172 gamble stimuli per condition. Specifically, we included 12 test gambles at the extreme ends of the payoff–probability distribution space. These consisted of different groups: three high payoff/high probability, three high payoff/low probability, three low payoff/high probability, and three low payoff/low probability gambles; see triangles in Figure 1). Payoffs were random draws between $1E\$ - 500E\$$ (low) and $2000E\$ - 2500E\$$ (high). Probabilities were random draws between $0.01 - .2$ (low) and $.8 - 1.0$ (high). Payoffs and probabilities were factorially combined to obtain the four payoff/probability combinations. The test gambles were interspersed after 100 environment-only trials (i.e., in the last third of the trials). We also included 10 average (mid-range payoff/mid-range probability) gambles (payoffs around $1250E\$$, probabilities around .5). These gambles were equally consistent with all conditions because they were not linked to either high or low (i.e., extreme) payoffs (squares in Figure 1). We added them to control for EV differences and to be able to study condition-dependent differences beyond gambles being uncommon, foreseeable, or reference gambles.

Procedure.   Participants indicated their willingness to sell for one gamble at a time (Figure 2), taking self-paced breaks between five blocks. The task was presented in the
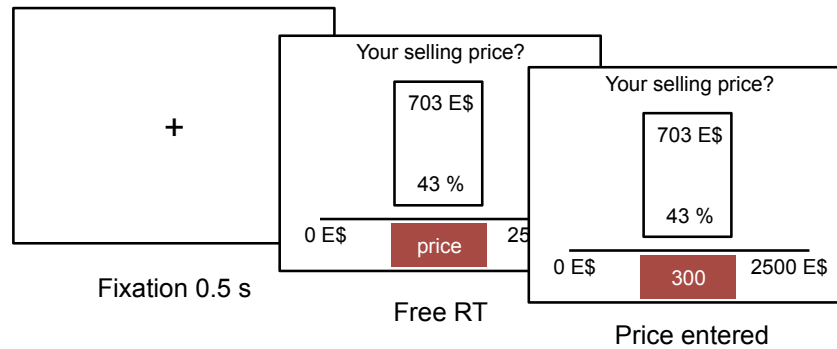
Figure 2. Task. In Experiment 1, participants saw a gamble and priced it in their own time. The price was set by moving an arrow along a rating scale and confirming the price by clicking on it.

form of a game show called "Keep or sell?" ("Behalten oder Verkaufen?"). We used selling rather than buying prices because the former have been shown to deviate less from gambles' EVs than the latter (however, both paradigms elicit responses strongly linked to EVs; see Yechiam et al., 2017); in general, both formats are suitable for exposing participants to different risk–reward structures. Experiments were coded in PsychoPy (Peirce, 2007). To motivate participants to indicate their true valuations of the gambles, we implemented a Becker-DeGroot-Marschak auction (Becker et al., 1964) as follows: Participants entered a price at which they would be willing to sell each gamble by moving the mouse along a rating scale ($0E\$ - 2500E\$$) and confirming the value with a click. To incentivize the task, the experimenter informed participants that 10 gambles would be randomly selected at the end of the experiment. For those 10 gambles, the experimenter then offered a randomly generated buying price between 0 and the absolute payoff in that gamble. If the experimenter's buying price exceeded the participant's selling price, participants sold the gamble and earned the buying price. If the participant's selling price exceeded the experimenter's buying price, the gamble was played out (e.g., 50% chance of $380E\$$). The dominant strategy in this task is to price a gamble based on its subjective value: Setting higher prices can prevent participants from selling unattractive gambles; setting lower prices can lead to them selling attractive gambles under value. In other words, the prices should approximate participants' certainty equivalents for the gambles.

Statistical Analyses

We used Bayesian estimation techniques in our analyses (Kruschke, 2015). Specifically, we applied Bayesian generalized linear mixed models using Stan in R for regression analyses with the rstanarm package (*RStanArm*, Version 2.9.0–4). We ran three chains using Markov Chain Monte Carlo sampling to draw from posterior distributions of the parameters. Depending on model complexity, we ran $10,000 - 30,000$ samples per chain (to ensure an effective sample size of $>10,000$ for each coefficient) and set a burn-in of 500 samples. We investigated convergence of our posteriors through visual inspection and the Gelman–Rubin statistic (Gelman & Rubin, 1992). In general, we report the mean of the posterior distribution of the parameter or statistic of interest and two-sided 95% equal tail credible intervals (CI) around each value.

In all analyses, we examined how the risk–reward context modulated the behavioral measure of interest (i.e., deviations from prices, deviations from individual median RTs, or pupil dilation). For better interpretability of the RT data, we report and plot parameters and CIs from regression models using untransformed RT data. As the RT data were slightly right-skewed, we reran the key RT analyses using a log transformation. Qualitatively, the conclusions were identical for log-transformed and untransformed data.

Our models were specified as follows. All analyses included a random participant intercept to control for individual variability in responses, and controlled for EV by including it as a predictor. We tested our hypotheses using the following recipe: In the *base effects* models, we modeled all test gambles simultaneously and tested whether responses to uncommon gambles differed from responses to foreseeable and reference gambles as baselines—as a main effect. Thus, the regressions were set up as follows:

```
prices/RTs/pupil ~ (1 | participant) + uncommon + EV, data = test gambles
```

In the *interaction* models, we allowed for an interaction between gamble type (uncommon/foreseeable/reference) and payoff–probability combination (high pay/high

probability; high pay/low probability; low pay/high probability; low pay/low probability).

Low-payoff/low-probability gambles were set as the baseline. In pseudocode,

```
prices/RTs/pupil ~ (1 | participant) + uncommon * pay-prob combn + EV, data =
    test gambles
```

This analysis also revealed the extent to which responses to uncommon combinations

of risk and reward depended on the gambles' EV (e.g., high pay/high probability). The

results of this analysis are depicted as posterior predictive distributions in the figures. In

the *by-type* models, we repeated the same analyses by subsetting the data for each

payoff–probability combination (high pay/high probability; high pay/low probability; low

pay/high probability; low pay/low probability) and modeling differences in responses as

simple effects. In pseudocode,

```
prices/RTs/pupil ~ (1 | participant) + uncommon + EV, data = high pay/high prob
```

In reporting the results, we first examined how participants responded across all trials

including the environment gambles, then using the average gambles that were equally

consistent with all risk–reward environments. Second, we tested for differences between

uncommon and foreseeable or reference gambles using base effects models. Third, we

reported results from the interaction models and by-type, simple models to characterize

how responses varied with different payoff–probability combinations. For model code, see

the Open Science Framework. The complete set of coefficients for all by-type models for

both baseline comparisons (foreseeable and reference gambles) is reported in the

Supplementary Materials.

We excluded trials in which participants indicated prices that exceeded the payoff

offered in the gamble by more than $100E$\$, as this suggests lack of attention to the task

(we did the same in Experiment 2). We used this rather liberal exclusion criterion for two

reasons: First, as a recent meta–analysis shows, selling prices frequently exceed gambles'

EVs (Yechiam et al., 2017, Table 2). Second, the EV of a gamble becomes harder to

compute given four–digit payoffs ($0 - 2500E\$$), compared to smaller payoff ranges. We also removed trials in which RTs deviated by $+/-3SD$ from an individual's median (before log transforms), assuming that these RTs are unlikely to reflect cognitive processing in a specific trial. We tested our hypotheses with and without these excluded trials and note when these exclusions led to qualitatively different results.

Results

Based on our exclusion criteria, we removed 4.6% ($718/15,480$) of the trials across all participants in Experiment 1.

Prices.   We expressed prices as deviations from the respective gamble's EV ($deviation = gamble_{price} - gamble_{EV}$). Thus, deviations could range from $-2500E\$$ to $2500E\$$; with 0 indicating perfect alignment between price and the gamble's EV. Across all trials and conditions, prices were closely aligned to gambles' EVs (all CIs included 0), but in general participants' prices exceeded EVs ($b_{unc} = 30.77E\$$, CI $= [-18.93E\$, 81.10E\$]$, $b_{neg} = 56.45E\$$, CI $= [-64.60E\$, 177.68E\$]$, $b_{pos} = 83.16E\$$, CI $= [-37.43E\$, 204.14E\$]$; deviations across all trials; risk–reward condition as a predictor, controlling for EV). As Figure 3A shows, participants in the positive condition seemed to overprice average gambles by a greater amount compared to participants in the uncorrelated condition ($b_{pos.>unc.} = 71.06E\$$, CI $= [9.32E\$, 133.11E\$]$; using risk–reward condition $\times$ gamble type as predictors, controlling for EV). Since this difference was not reliable across the environment trials, we do not follow up on it further.

Focusing on the test gambles, we tested whether participants provided prices deviating from the gambles' EVs when risk–reward combinations were uncommon. As Figure 4 shows, there were no reliable pricing differences as a function of whether the gamble was uncommon or not ($b_{uncom.>fores.} = 2.52E\$$, CI $= [-26.03E\$, 30.89E\$]$, $b_{uncom.>ref.} = -11.59E\$$, CI $= [-49.15E\$, 26.20E\$]$, base effects models using foreseeable and reference gambles as a baseline, respectively; controlling for EV). Moreover, prices
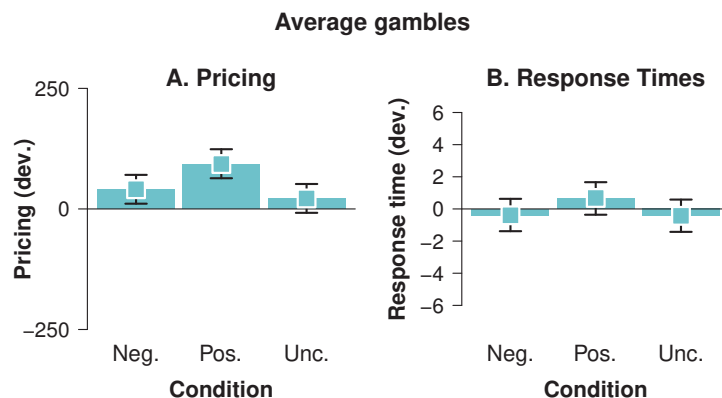
Figure 3. Behavioral results for the average gambles interspersed in Experiment 1. (A) Pricing. Prices deviated more in the positive condition compared to the uncorrelated condition. (B) Response times (RTs). We found little to no difference in RTs across conditions for the average gambles. RTs are expressed as deviations from individual median RTs across trials $(RT_{trial} - RT_{ind.\,Md})$. Symbols and error bars represent the mean and the 95% credible interval of the posterior predictive distribution.

were not drawn towards participants' expectations in a given environment (e.g., prices for uncommon, high-pay/high-probability gambles could have been lower than prices for foreseeable high-pay/high-probability gambles): As a comparison between panels in Figure 4 shows, prices did not vary as a function of different payoff–probability combinations across conditions (all CIs included 0 in the interaction model; the same holds for by-type, simple models; see Supplementary Material).

Response times (RTs). The mean RT across all gamble types was 14.17s per trial in the uncorrelated condition (CI $= [12.40s, 15.91s]$) and credibly lower in the negative ($M = 10.3s$, $b_{neg.>unc.} = -3.74s$, CI $= [-6.22s, -1.20s]$), but not in the positive condition ($M = 13.0s$, $b_{pos.>unc.} = -1.30s$, CI $= [-3.78s, 1.17s]$, both models using the uncorrelated condition as a baseline, controlling for EV). Next, to examine changes in RTs for the test gambles and to control for individual differences in baseline RTs, we computed the difference between the observed RT and the median RT as an indicator of change in processing time $(RT_{trial} - RT_{ind.\,Md})$. RTs should be similar for average gambles that fit each risk–reward condition equally well; they should also deviate very little from median RTs. As Figure 3B shows, this was indeed the case: Participants' RTs for the average
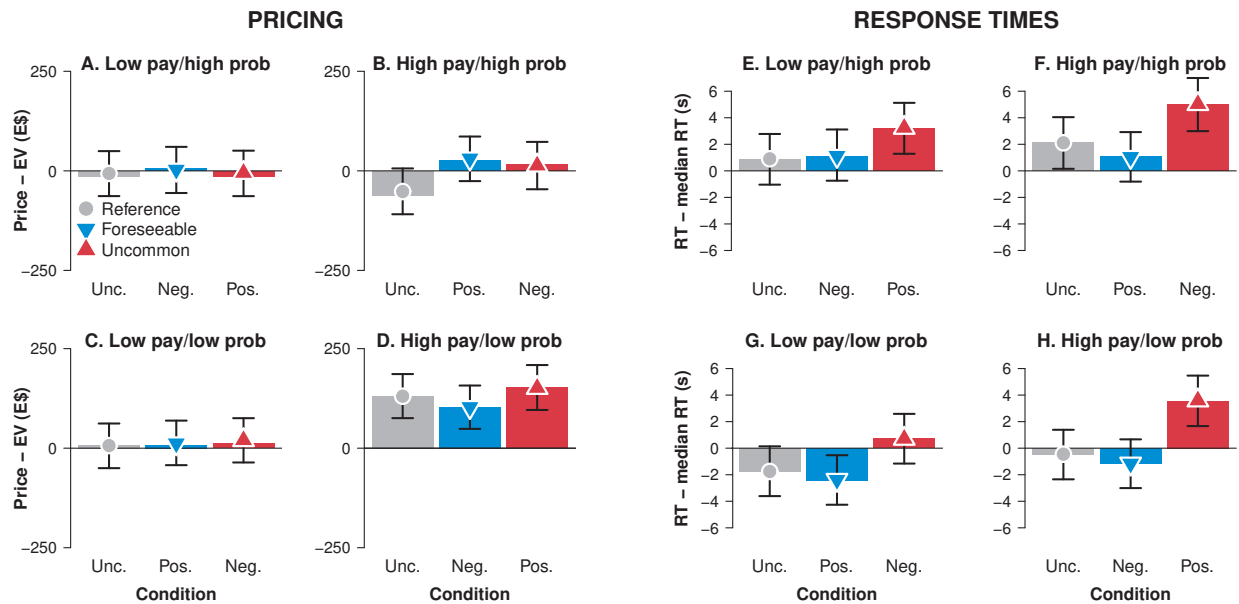
Figure 4. Behavioral results for the test gambles interspersed in Experiment 1. (A–D) Pricing. We found no difference in how well prices were adjusted to the expected values when gambles were uncommon. (E–H) Response times (RTs). Participants slowed down when gambles were uncommon. RTs are expressed as deviations from individual median RTs across trials $(RT_{trial} - RT_{ind.\,Md})$. Symbols and error bars represent the mean and the 95% credible interval of the posterior predictive distribution.

gambles were similar across conditions and all CIs included 0.

How did RTs change when gambles were uncommon? As Panels E–H in Figure 4 show, participants slowed down in response to uncommon gambles (red bars). On average, they spent around 3.5s longer on uncommon gambles than on foreseeable gambles $(b_{uncom.>fores.} = 3.40s$, CI $= [1.49s, 5.32s]$, controlled for EV and individual variation). The RTs to the reference gambles from the uncorrelated environment were statistically indistinguishable from those to foreseeable gambles. Consequently, the effects involving uncommon gambles were qualitatively the same as when using these gambles as a baseline (see Supplementary Materials 1.3). The results were comparable when all trials are maintained in the analyses $(b_{uncom.>fores.} = 3.63s$, CI $= [2.56s, 4.98s]$, no exclusions, controlled for EV and individual variation), and when using normalized, log-transformed RTs $(b_{uncom.>fores.} = 0.32$, CI $= [0.25, 0.39]$, same model but log-transformed RTs).

Lastly, we examined the extent to which RTs—either generally or for uncommon

gambles—depended on particular payoff–probability combinations. In general, as the mostly negative values in Figure 4G and the mostly positive values in Figure 4F show, participants responded faster to low-value gambles compared to high-value gambles, relative to their median RT per gamble. This result is captured by a main effect of EV on RTs ($b_{EV} = 0.004s$, CI $= [0.000s, 0.008]$; interaction model gamble type $\times$ uncommon, main effect of EV).

Responses to uncommon gambles (compared to foreseeable gambles) were most pronounced when they involved high payoffs (Figure 4, F & H). However, there were no reliable interactions between these payoff–probability combinations and gambles being uncommon for either high-pay/high-probability ($b_{uncom.>fores.} = .37s$, CI $= [-2.33s, 3.07s]$) or high-pay/low-probability gambles ($b_{uncom.>fores.} = 1.57s$, CI $= [-1.19s, 4.35s]$; interaction uncommon $\times$ pay–prob comb uncommon; baseline low pay/low probability; controlling for EV). All results were qualitatively the same in a model using the reference gambles as a baseline (see Supplementary Material).

Summary of Experiment 1

Experiment 1 suggests that people form expectations about risk–reward structures that reflect the experienced environments. This was evident in participants' RTs for options that did not match these expectations: They slowed down in response to uncommon options for all payoff–probability combinations. Participants slowed down most for uncommon high-payoff gambles—where indicating an accurate price seems more important. We note that the differences between payoff–probability combinations were not reliable and the value-dependent variation in RTs is best captured as a main effect. The finding that participants took longer in responding to uncommon, high-pay/high-probability gambles is consistent with earlier research showing that participants are sensitive to the range of possible rewards in a given experiment (for a related result see Ludvig et al., 2014, 2018): Such gambles were uncommon in the negative

condition, which was characterized by otherwise similar and smaller EVs (all below 700E$). At the same time, prices were equally adjusted to the gambles' EVs, regardless of whether gambles were uncommon or not. Notably, participants were sensitive to uncommon combinations of risk and reward irrespective of any feedback on whether or not an outcome was obtained and even though they had experienced the entire range of probabilities and payoffs before being exposed to uncommon combinations of risk and reward.
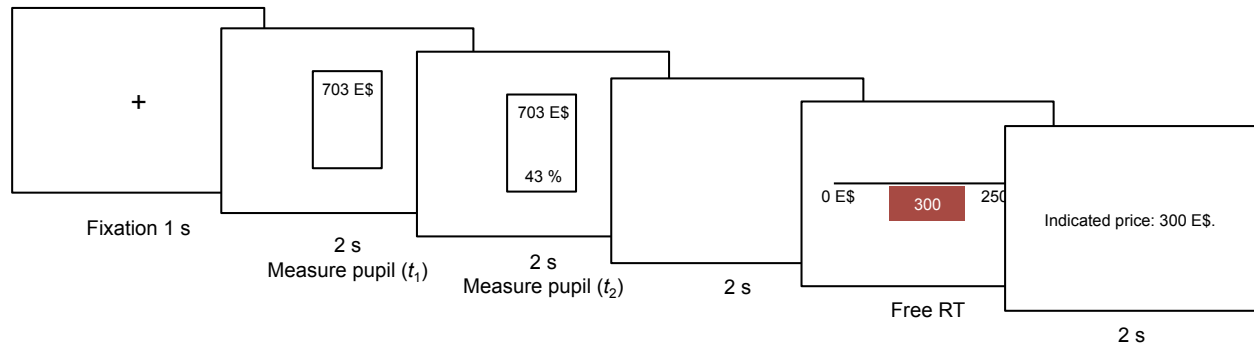
Experiment 2: Pupillary and Behavioral Responses to Uncommon Risk–Reward Combinations

In Experiment 2, we adapted the methodology from Experiment 1 to include an eye-tracking component that allowed us to measure pupillary responses to uncommon options. We outline key differences below. After the pricing task, participants completed a two-alternative forced choice task reported in a proceedings paper (Leuker et al., 2017).

Methods

Participants. Ninety-three (55 female) participants (age $M = 25.6$ years, $SD = 3.7$ years) 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 (€3.53–€11.67).

Stimuli. We reduced the number of stimuli to 90 condition-dependent gambles and increased the number of gambles common to each of the three conditions. We reduced the number of trials to ensure that all participants could complete the experiment in a reasonable time (presenting payoffs and probabilities in a fixed sequence takes more time than presenting gambles in full). At the same time, we aimed at a larger set of test gambles per participant. To that end, we created six gambles for each payoff–probability combination. Overall, these procedures resulted in 120 gambles. The test gambles were interspersed after 40 environment-only trials. As in Experiment 1, we also added six average-payoff/average-probability gambles.

Figure 5. In Experiment 2, participants priced gambles after seeing the payoff and probability of each gamble in a fixed sequence. After 4s in total, the gamble disappeared from the screen. We added a blank screen to achieve sufficient spacing between the critical stimuli (i.e., a gamble's risk–reward relationship) and the participant's manual response. Participants then had as much time as they wanted to indicate a price for the gamble on a rating scale. We measured pupil dilation as a response to a gamble's payoff alone (t_1) and as a response to its payoff and probability combined (t_2).

Procedure

Experiment 2 differed from Experiment 1 in that the payoff and probability information appeared sequentially: After a fixation cross, the payoff appeared for 2s, followed by the probability for another 2s. After a blank screen (2s), the screen automatically switched to a rating scale. Participants entered the prices they were willing to sell the gamble for by moving the mouse along the rating scale ($0E\$ - 2500E\$$) and clicking on the value to confirm (free RT). This sequential presentation of only payoffs and then payoffs and probabilities combined allowed us to analyze pupillary responses to payoffs only and pupillary responses to the joint presentation of payoffs and probabilities (Figure 5). To control for the pupillary light reflex, we matched the gambles' luminance to the background of the screen by defining stimuli colors in the Derrington, Krauskopf, and Lennie color space (see Derrington et al., 1984; MacLeod & Boynton, 1979). Gambles were presented in orange on a gray background with the same luminance. For clarity, these colors are not shown in the figure.

Eye Tracking

We collected binocular eye tracking data with an EyeTribe tracker, sampled at $60Hz$. The experiment was implemented in PsychoPy 1.83.01 (Peirce, 2007) and the eye-tracking interface PyTribe (Dalmaijer et al., 2014). Before the task, each participant's eye movements were calibrated using the EyeTribe UI with a 9-point grid (< 0.7 degrees of visual angle). Participants were seated approximately 60 cm from the screen with their chin on a chinrest affixed to a table, in a room with negligible ambient light. We obtained pupil size from the left and the right eye (arbitrary units, measured at every Hz).

Eye-Tracking Analyses

Pupillary data were preprocessed as follows. We used EyeTribe's default settings to detect fixations and removed saccade data, because pupillary responses during (and even before) saccades differ systematically from those during fixations (Mathôt et al., 2016). Trials were discarded if pupil size deviated more than $3SDs$ from a participant's median pupil size and if it was outside plausible values (range [10,40] in arbitrary units given by Eyetribe). This procedure removed blinks (rows with values [0,0]) and measurement error. We smoothed the data using a lowess filter, and we averaged the pupil size of the left and the right eye. We removed trials with fewer than 15 samples, which indicates extremely poor eye tracking (one would expect 240 samples, minus a few blinks, per 2s period). This applied to a small proportion of all trials: .07 of all trials at t_1 (when reward was shown) and .05 of all trials at t_2 (when reward and probability were shown).

To facilitate comparisons across participants and control for differences in pupil size that were nonspecific to our conditions, (Preuschoff et al., 2008; Cavanagh et al., 2014; Fiedler & Glöckner, 2012), we analyzed pupillary signals aligned to a baseline pupil size. This is consistent with the literature. As our baseline pupil size, we used the offset of stimulus presentation. Specifically, we used the median value in the first 0.1s at t_1 (after the reward was shown) and at t_2 (after reward and probability were shown). We used this
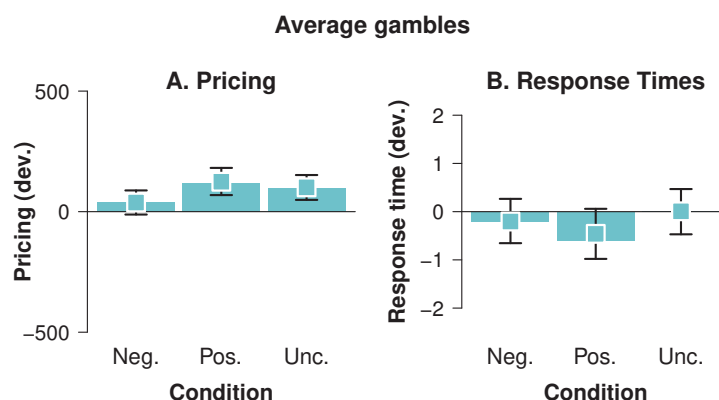
baseline to obtain a similar measure of pupil changes for both t_1 and t_2. Often the fixation cross time is used as a baseline; in our paradigm, however, the fixation cross only preceded stimulus appearance at t_1 (Figure 5). We obtained the normalized pupil size by subtracting the signal at each time point from the baseline signal and then dividing by the baseline signal, resulting in a percentage change relative to the stimulus onset.

Pupillary responses to psychologically relevant stimuli are assumed to occur after approximately 1s, and are therefore conceptually different from the pupillary light reflex that occurs after milliseconds (Gagl et al., 2011; van Steenbergen & Band, 2013). As in previous research using pupillary responses in the context of choice (Cavanagh et al., 2014), we therefore set an a priori region of interest from 1 to 2s poststimulus at t_1 (payoff visible on screen) and t_2 (payoff and probability visible together on screen) for our statistical analyses (Figure 1). We obtained the median percentage change in pupil dilation within this a priori region of interest. We compared the results of this analysis with results using the mean dilation (e.g., as in Mathôt et al., 2016) and peak dilation (e.g., as in Fiedler & Glöckner, 2012), which were qualitatively very similar (see Supplementary Materials for results using the other two indicators). We would like to stress that our focus was not on the time course of pupil dilation because we introduced a fixed lag between payoffs, probabilities, and participants' ability to respond (however, for completeness, we plot the time course of pupillary responses in the Supplementary Material). While some research has studied pupil dilation shortly before a decision is made (Fiedler & Glöckner, 2012, finding that pupil dilation increases while the participant is deciding), in our setup we could not determine a unique time point at which participants made their choice, as they could have reached a decision prior to being able to enter it on the rating scale.

Results

Based on our exclusion criteria, we removed 9.2% $(1,027/11,160)$ of the trials across all participants in Experiment 2.

Pricing. As in Experiment 1, we expressed prices as deviations from a gamble's EV ($gamble_{price} - gamble_{EV}$), with 0 indicating perfect alignment between price and the EV. Across all trials and conditions, prices were closely aligned to gambles' EVs, but in general participants' prices exceeded EVs ($b_{unc} = 94.96E\$$, CI $= [-27.74E\$, 217.75E\$]$, $b_{neg.} = 24.17E\$$, CI $= [-110.93E\$, 256.93E\$]$, $b_{pos.} = 121.44E\$$, CI $= [-110.93E\$, 354.35E\$]$, main effects of condition; controlling for EV). Figure 6A plots the extent to which participants' prices deviated from EVs for the average gambles, and shows that there were no reliable differences between conditions for these gambles ($b_{pos.>unc.} = 26.69E\$$, CI $= [-79.31E\$, 134.68E\$]$, $b_{neg.>unc.} = -69.23E\$$, CI $= [-175.88E\$, 37.79E\$]$; using risk–reward condition as a predictor, controlling for EV). The result that participants tended to overprice gambles in the willingness-to-sell paradigm is consistent with Experiment 1.

Figure 6. Behavioral results for the average gambles in Experiment 2. (A) Pricing. Prices exceeded gambles' expected values in all conditions. (B) Response times (RTs). We found little to no difference in RTs across conditions. RTs are expressed as deviations from individual median RTs across trials ($RT_{trial} - RT_{ind. Md}$). Symbols and error bars represent the mean and the 95% credible interval of the posterior predictive distribution.

As in Experiment 1, we analyzed the extent to which prices differed when gambles were uncommon versus foreseeable or reference gambles. Across all gamble types, uncommon gambles were priced slightly—but not reliably—higher than foreseeable gambles ($b_{uncom.>fores.} = 28.16E\$$, CI $= [-2.74; 58.52]$) and reference gambles

($b_{uncom.>ref.} = 12.61E\$$, CI $= [-62.28; 87.53]$, deviations across all trials; risk–reward condition as a predictor, controlling for EV). As panels B and D in Figure 7 show, participants' evaluation of uncommon gambles depended on payoff–probability combinations as well as payoff magnitudes. Specifically, prices were lower for uncommon high-payoff/high-probability gambles ($b_{uncom.>fores.} = -101.84E\$$, CI $=$ $[-234.88E\$, 29.88\$]$, by-type model, controlling for EV). Conversely, prices were higher for uncommon high-payoff/low-probability gambles ($b_{uncom.>fores.} = 208.20E\$$, CI $=$ $[-1.36E\$, 418.71E\$]$, by-type model, controlling for EV). This pattern of results was similar when using reference gambles as a baseline (see gray bars in Figure 7, and Supplementary Material), but note that all CIs included 0.

Taken together, these pricing patterns are consistent with the payoff–probability combinations that participants could expect in a given environment: The negative condition prompts the expectation that a high payoff will be accompanied by a low probability; the positive condition prompts the expectation that a high payoff will be accompanied by a high probability. However, participants adjusted their prices to the environment-dependent expectations for high-payoff, but not low-payoff, gambles. Note that these results differ from those in Experiment 1, where prices did not vary depending on whether a gamble was uncommon or foreseeable. One reason for this discrepancy could be the limited presentation duration of the gambles in Experiment 2. Having less time to process a gamble while it was on screen (4s in total) may have led participants to only partially price the gambles based on experiences in a given environment.

Response times (RTs). Participants took 4.33s on average per trial from seeing the empty screen (2s enforced) to entering a response on the rating scale. The mean RT in the uncorrelated condition was 4.02s (CI $= [3.44s, 4.62s]$) and not reliably different in the negative ($M = 4.39s$, $b_{neg.>unc.} = 0.37s$, CI $= [-0.46s, 1.19s]$), or the positive condition ($M = 3.99s$, $b_{pos.>unc.} = -0.03s$, CI $= [-0.88s, 0.80s]$, both models using the uncorrelated condition as a baseline, controlling for EV). The range of RTs was smaller in Experiment 2
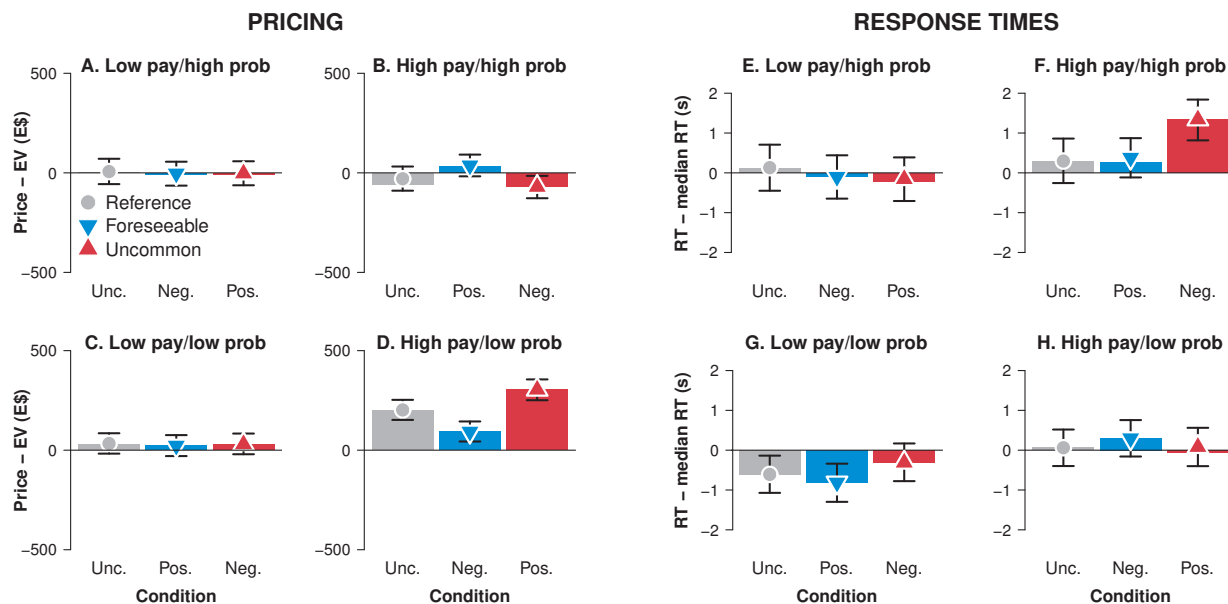
Figure 7. Behavioral results for the test gambles interspersed in Experiment 2. (A–D) Pricing. For low payoffs, there was little to no difference in how well prices of uncommon gambles were aligned with the expected values of the gambles (A, C). For high payoffs, prices of uncommon options differed from prices of foreseeable options (B, D). (E–H) Response times (RTs). Participants slowed down when options were uncommon. This effect was driven by the uncommon, high-value options (panel F). RTs are expressed as deviations from individual median RTs across trials ($RT_{trial} - RT_{ind.\,Md}$). Symbols and error bars represent the mean and the 95% credible interval of the posterior predictive distribution.

than in Experiment 1 due to the experimental design, in which a price could not be entered before seeing payoff and probability information for 2s each.

As in Experiment 1, to examine changes in RTs for the test gambles and to control for individual differences in baseline RTs, we computed the difference between the observed RT and the median RT as an indicator of change in processing time ($RT_{trial} - RT_{ind.\,Md}$). RTs should be similar for average gambles that fit each risk–reward condition equally well; they should also deviate very little from median RTs. As Figure 6B shows, this was indeed the case: Participants' RTs for average gambles were similar across conditions and all CIs included 0.

Like Experiment 1, our focus was on changes in RTs for uncommon gambles. Again, participants spent more time evaluating uncommon gambles than foreseeable gambles ($b_{uncom.>fores.} = 0.28s$; CI $= [0.03s, 0.53s]$, control for EV). They also spent more time

evaluating uncommon versus reference gambles, but the CI included 0 ($b_{uncom.>ref.} = 0.26s$; CI $= [-0.29s, 0.80s]$; control for EV). Overall, the effect sizes were smaller and the differences in RTs less pronounced—likely a result of the fixed presentation durations of payoffs and probabilities. The results were comparable when all trials are maintained in the analyses ($b_{uncom.>fores.} = 0.55s$, CI $= [0.27s, 0.84s]$, no exclusions, controlled for EV and individual variation), and when using normalized, log-transformed RTs ($b_{uncom.>fores.} = 0.28$, CI $= [0.03, 0.53]$, same model but log-transformed RTs).

We again examined the extent to which RTs varied by different payoff–probability combinations: In general, as the mostly negative values in Figure 4G and the mostly positive values in Figure 4F show, participants responded slightly faster to low-value gambles compared to high-value gambles relative to their median RT per gamble. However, most CIs in panels G and F include 0, and there was no reliable effect of EV on RTs ($b_{EV} = 0.0005s$, CI $= [-0.0002s, 0.00012s]$; interaction model gamble type $\times$ uncommon, main effect of EV).

As Figure 7F shows, RTs between uncommon and foreseeable gambles only differed for high-payoff/high-probability gambles ($b_{uncom.>fores.} = 0.76s$; CI $= [0.08s, 1.45s]$, interaction pay–prob combn $\times$ uncommon, controlling for EV, baseline low payoff/low probability). In other words, the effect for the uncommon gambles was driven by high-payoff/high-probability gambles. These results are reliable for both baseline conditions ($b_{uncom.>fores.} = 0.94s$; CI $= [0.03s, 1.84s]$; $b_{uncom.>ref.} = 1.06s$; CI $= [0.15s, 1.97s]$), and more pronounced when no trials were excluded ($b_{uncom.>fores.} = 1.15s$; CI $= [0.01s, 2.29s]$; $b_{uncom.>ref.} = 1.52s$; CI $= [0.37s, 2.67s]$, by-type models, controlling for EV). As Figure 7, panels E, G, & H show, there were no reliable differences in responses to uncommon versus foreseeable or expected gambles for the other three payoff–probability combinations (all CIs include 0; see Supplementary Material). Note that this result deviates from Experiment 1, where all uncommon combinations of risk and reward were evaluated more conservatively. Across both experiments, the most pronounced change in RTs was observed when gambles

were both uncommon and of high value. This finding is consistent with uncommon, high-value options redefining the range of possible EVs: High-payoff/high-probability gambles were uncommon in the negative condition (all other EVs below $600E\$$), where such high-value gambles had not been experienced in the environment gambles.

Pupil dilation.   We analyzed pupil responses at two time points. Pupillary responses to payoffs only (t_1) were used to test the influence of payoff magnitude: We expected that pupil size would not vary between high and low payoffs at t_1. This was indeed the case across all conditions ($b_{unc} = -0.07$, CI $= [-0.91, 0.76]$, using the uncorrelated condition as baseline; $b_{neg} = -0.40$, CI $= [-1.59, 0.71]$, $b_{pos} = -0.33$, CI $= [-1.52, 0.88]$, all comparisons high payoff > low payoff. Results were qualitatively the same with and without excluded trials; see Supplementary Materials). We expected pupil dilation to change as a function of whether or not a gamble had an uncommon payoff–probability combination. Figure 8 shows mean changes in pupil dilation after both payoff and probability information were presented (t_2). Panels A–D suggest that pupil size was not associated with gambles being uncommon ($b_{uncom.>fores.} = 0.06$, CI $= [-0.60, 0.73]$, $b_{uncom.>ref.} = 0.05$, CI $= [-0.98, 1.09]$, base effects model across gamble types; controlled for EV).

Rather, as for the RTs, there was an interaction between gambles with an uncommon payoff–probability combination and the gambles' value ($b_{uncom.>fores.} = 3.24$; CI $= [1.44, 5.07]$, $b_{uncom.>ref.} = 1.63$; CI $= [-0.22, 3.48]$, interaction pay–prob combn $\times$ uncommon, controlling for EV). Note that in Figure 8B the pupil size was reliably different from 0 but not from gambles being foreseeable $b_{uncom.>fores.} = 1.16$; CI $= [-0.38, 2.17]$, by-type model for the high-pay/high-probability gamble, controlling for EV). Panels A, C, and D show that the change in pupil dilation for the other gamble types was statistically indistinguishable from a 0% change, regardless of whether these gambles were uncommon or not.

In sum, both RTs and pupil dilation differed as a function of gambles being

uncommon and of high value. In an exploratory analysis, we found that these two indicators were also correlated with each other for the high value, but not low-value gambles ($b_{rt} = 0.51$; CI $= [0.06, 0.96]$, interaction RT $\times$ payoff–probability combn using low-pay/low-probability gambles as a baseline). Across all gambles, RTs and pupil dilation were even anticorrelated ($b_{rt} = -0.08$; CI $= [-0.13, -0.33]$, predicting pupil dilation from RTs as a main effect, controlling for participant variation and EV).
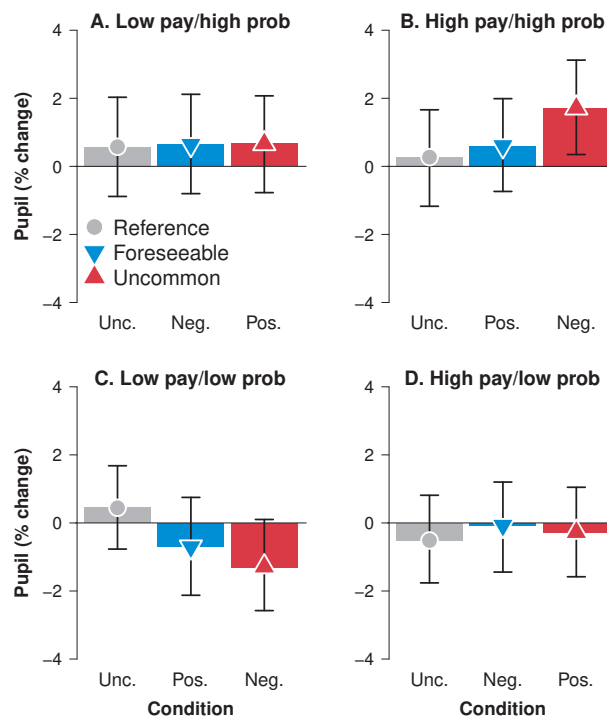
Figure 8. Pupil dilation after gambles were presented (with both payoff and probability information on the screen). Pupil dilation was computed as the median percentage change from 1000–2000ms after the stimulus. Black dots and error bars represent the mean and the 95% posterior predictive distribution.

Summary of Experiment 2

Experiment 2 corroborates the finding that people develop expectations about options appearing in a particular risk–reward environment. As in Experiment 1, this was evident in longer RTs for uncommon gambles. This effect was driven by high-payoff/high-probability gambles that were of higher value than the other gambles in the set. Moreover, there were some differences in pricing, such that prices partially

reflected participants' expectations in a given environment; this result, however, was not robust across payoff magnitudes or across baseline comparisons, so we do not elaborate on it further. Lastly, like RTs, uncommon high-value options were associated with a reliable increase in pupil dilation. However, pupil dilation was not sensitive per se to whether a risk–reward combination was foreseeable or uncommon.

General Discussion

Decision makers often face environments in which risks and rewards are systematically related (Pleskac & Hertwig, 2014). We found evidence that such regularities leak into people's evaluations of risky options—especially if these options represent an uncommon combination of risk and reward. We investigated the extent to which uncommon options were linked to three indicators: prices, RTs, and pupil dilation and found that some, but not all, of these indicators were consistently linked to whether a gamble matched or mismatched the learned environmental structure. Next, we discuss these three indicators in detail and consider their broader implications for preferential choice.

Do people evaluate uncommon options differently from foreseeable options? Across experiments and gamble types, we did not find robust evidence for this as indicated by the selling prices participants provided. This suggests that risk–reward structures may play only a tangential role (if any) in shaping preferences in decisions under risk, when both payoffs and probabilities are known. Earlier research suggested that risk–reward structures do shape preferences in decisions under uncertainty—decisions in which probabilities are unknown or difficult to ascertain (Pleskac & Hertwig, 2014). It has been shown that people infer the unknown probability directly from the payoff in decisions under uncertainty and adapt their inferences to previously learned risk–reward structures. As the inferred probability, together with the payoff, determines the subjective value of the option, different choice ecologies can lead to different choices (Leuker et al., 2018). In decisions under risk, by contrast, such inferences are not necessary and apparently are not made.

How do people process uncommon options? Our data show that people take more time when evaluating an option when it is uncommon rather than foreseeable. How much longer people take to evaluate an uncommon option also depends on whether the option is of high value (high payoff/high probability) or of low value (low payoff/low probability). Specifically, high-payoff/high-probability options were uncommon in the negative condition, which was characterized by otherwise similar and smaller EVs (all below 700E$), suggesting that participants were sensitive to the range of possible rewards in the given experiment (see Ludvig et al., 2014, 2018, for a similar finding in decisions from experience). These results are consistent with range–frequency theory (Parducci, 1965). This uncommon high-value effect was also present in pupil dilation. The differences between high-value and low-value gambles are consistent with a "hold your horses" mechanism by which people behave less impulsively in the presence of high-value options (Cavanagh et al., 2014; Frank et al., 2007). That is, more scrutiny is required when high payoffs are at stake.

We note that processing differences in response to these uncommon combinations of risk and reward (that is, "risk–reward novelty") can be dissociated from pure novelty. While participants' RTs for novel gambles (first five trials) were substantially longer than RTs for any of the test gambles, differences between uncommon and foreseeable/reference gambles persisted (see Supplementary Materials S1.4 and 2.5 for these analyses).

More generally, our data bring a new perspective to the growing body of research on how the environmental distribution of monetary payoffs and probabilities influences how otherwise identical options are evaluated (Birnbaum, 1992; Ludvig et al., 2014, 2018; Stewart et al., 2006, 2015; Walasek & Stewart, 2015). Here we show that people go beyond evaluating options from the given information. Instead, our results point to a mixture of evaluations from given information and evaluations from the environment. For instance, prices for uncommon options approximated the options' EVs but sometimes shifted in the direction of environmental expectations. What is more, we observed systematic shifts in the way that uncommon options were processed. Such environment-based evaluations are

not anticipated by prominent theories of choice, which conceptualize risks and rewards as independent attributes that determine the expected utility of an option (von Neumann & Morgenstern, 1944) or its subjective value (Tversky & Kahneman, 1992).

In addition, our results bring a novel dimension to studies of surprise. We show that feedback is not a prerequisite for the mind to detect uncommon options in risk–reward environments, as is the case for "risk-prediction errors" (Bossaerts, 2010; O'Neill & Schultz, 2013; Preuschoff et al., 2008) or "reward-prediction errors" (Schultz, 2002; Schultz et al., 1997). Instead, it is necessary for the mind to have a representation of the risk–reward structure of the environment. Similar to feedback–based paradigms, in our experiments it also mattered whether the direction of the uncommon stimulus was high or low (i.e., whether an option was of high or low value).

What is the function of detecting uncommon options in risk–reward environments? In value-based choice, reward–prediction errors are thought to aid reward-based learning—e.g. after obtaining an unexpected reward, a person may expect similar rewards to occur again. A similar mechanism would be plausible in risky choice if participants anticipated the uncommon high-value option being added to their bonus at the end of the study (without feedback after each choice). Paying attention to risk–reward combinations is crucial to maintaining an accurate model of the world, as risk–reward relationships may vary across environments. For instance, the relationship is stronger in monetary domains than in nonmonetary domains, see (Pleskac & Hertwig, 2014). However, the amount of evidence needed before a model is adjusted remains an open question. Our previous experiments suggested that the risk–reward regularity people extract largely reflects the structure of the environment trials (Leuker et al., 2018). Maintaining a representation of the risk–reward relationship even after seeing uncommon risk–reward combinations would be consistent with a "rule-plus-exception model," according to which people may learn exceptions to a rule instead of updating that rule if they are unable to identify the rule that would account for the exceptions (Nosofsky et al., 1994). For instance, a person

"might regard a single-dimension rule as tentatively acceptable as long as it correctly classifies 60% of the incoming exemplars" (p. 56).

In value-based decision making, extracting a rule from the overall environment and forming expectations may serve a very particular function: to help people identify and correctly respond to better-than-average options. Forming expectations in this way may translate to nonlaboratory environments in which risks and rewards are inversely related, for instance when playing the lottery, betting at the horsetrack (Pleskac & Hertwig, 2014), or investing in the stock market. In these environments, people know that there is usually "no free lunch," in that the larger rewards they desire occur only rarely—and if they are lucky enough to get a "free lunch" once, it does not mean that their model of risk–reward environments will change.

Conclusion

When making decisions under risk, people take into account the degree to which an option matches or deviates from the general risk–reward regularities in an environment. Although this does not impact their evaluations of the option's value, they inspect uncommon options more carefully; this holds primarily when uncommon options offer a high EV. Forming expectations about risk–reward relationships may help people to identify above-average options and channel limited attentional resources there. Doing so can be adaptive in many environments in which risks and rewards are inversely related.

References

Becker, G. M., Degroot, M. H., & Marschak, J. (1964). Measuring utility by a
single-response sequential method. *Behavioral Science*, *9*(3), 226–232. doi:
10.1002/bs.3830090304

Birnbaum, M. H. (1992). Violations of monotonicity and contextual effects in choice-based
certainty equivalents. *Psychological Science*, *3*(5), 310–314. doi:
10.1111/j.1467-9280.1992.tb00679.x

Bossaerts, P. (2010). Risk and risk prediction error signals in anterior insula. *Brain
Structure and Function*, *214*(5), 645–653. doi: 10.1007/s00429-010-0253-1

Brunswik, E. (1943). Organismic achievement and environmental probability.
Psychological Review, *50*(3), 255–272. doi: 10.1037/h0060889

Cavanagh, J. F., Wiecki, T. V., Kochar, A., & Frank, M. J. (2014). Eye tracking and
pupillometry are indicators of dissociable latent decision processes. *Journal of
Experimental Psychology: General*, *143*(4), 1476–1488. doi: 10.1037/a0035813

Dalmaijer, E. S., Mathôt, S., & Van der Stigchel, S. (2014). PyGaze: An open-source,
cross-platform toolbox for minimal-effort programming of eyetracking experiments.
Behavior Research Methods, *46*, 913–921. doi: 10.3758/s13428-013-0422-2

Derrington, A., Krauskopf, J., & Lennie, P. (1984). Chromatic mechanisms in lateral
geniculate nucleus of macaque. *Journal of Physiology*, *357*, 241–265. doi:
10.1113/jphysiol.1984.sp015499

Fiedler, S., & Glöckner, A. (2012). The dynamics of decision making in risky choice: An
eye-tracking analysis. *Frontiers in Psychology*, *3*, 335. doi: 10.3389/fpsyg.2012.00335

Fortin, J. (2017, December 30). Glitch in South Carolina lottery could mean $19.6 million
in winnings. *New York Times*. Retrieved from
https://www.nytimes.com/2017/12/30/us/lottery-south-carolina.html?smid=pl-share

Fortin, J. (2018, May 31). Sorry, lottery winners. South Carolina won't pay, but here's
your $1. *New York Times*. Retrieved from https://nyti.ms/2xDTp8L

Frank, M. J., Samanta, J., Moustafa, A. A., & Sherman, S. J. (2007). Hold your horses:
Impulsivity, deep brain stimulation, and medication in parkinsonism. *Science*,
318(5854), 1309–1312. doi: 10.1126/science.1146157

Gagl, B., Hawelka, S., & Hutzler, F. (2011). Systematic influence of gaze position on pupil
size measurement: Analysis and correction. *Behavior Research Methods*, *43*(4),
1171–1181. doi: 10.3758/s13428-011-0109-5

Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple
sequences. *Statistical Science*, *7*(4), 457–511.

Harrington, R. (2017, December 31). A computer glitch on Christmas could give mistaken
winners in the South Carolina lottery almost $20 million. *Business Insider*. Retrieved
from http://www.businessinsider.de/
computer-glitch-20-million-dollars-winnings-south-carolina-lottery-2017-12

Hoffart, J. C., Rieskamp, J., & Dutilh, G. (2018). How environmental regularities affect
people's information search in probability judgments from experience. *Journal of
Experimental Psychology: Learning, Memory, and Cognition*, No Pagination
Specified–No Pagination Specified. doi: 10.1037/xlm0000572

Kruschke, J. (2015). *Doing bayesian data analysis: A tutorial introduction with R*.
London, United Kingdom: Academic Press.

Leuker, C., Pachur, T., Hertwig, R., & Pleskac, T. J. (2018). Exploiting risk–reward structures in decision making under uncertainty. *Cognition*, 186–200. doi: doi:10.1016/j.cognition.2018.02.019

Leuker, C., Pleskac, T. J., Pachur, T., & Hertwig, R. (2017). How the mind exploits risk–reward structures in decisions under risk. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. Davelaar (Eds.), *Proceedings of the 39th Annual Conference Cognitive Science Society* (pp. 2543–2548). London, UK.

Ludvig, E. A., Madan, C. R., & Spetch, M. L. (2014). Extreme Outcomes Sway Risky Decisions from Experience. *Journal of Behavioral Decision Making*, *27*(2), 146–156. doi: 10.1002/bdm.1792

Ludvig, E. A., McMillan, N., Madam, C. R., Xu, H., & Spetch, M. L. (2018). Living near the edge: how extreme outcome and their neighbours drive risky choice. *Journal of Experimental Psychology: General*.

MacLeod, D. I. A., & Boynton, R. M. (1979). Chromaticity diagram showing cone excitation by stimuli of equal luminance. *Journal of the Optical Society of America*, *69*(8), 1183–1186. doi: 10.1364/JOSA.69.001183

Mathôt, S., van der Linden, L., Grainger, J., & Vitu, F. (2016). The pupillary light response reflects eye-movement preparation. *Journal of Experimental Psychology: Human Perception and Performance*, *41*(1), 28–35. doi: dx.doi.org/10.1037/a0038653

Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. , *101*(1), 53–79. doi: 10.1037/0033-295X.101.1.53

O'Neill, M., & Schultz, W. (2013). Risk prediction error coding in orbitofrontal neurons. *The Journal of Neuroscience*, *33*(40), 15810–15814.

Parducci, A. (1965). Category judgment: A range-frequency model. *Psychological Review*, *72*(6), 407–418. doi: 10.1037/h0022602

Peirce, J. W. (2007). PsychoPy - Psychophysics software in python. *Journal of Neuroscience Methods*, *162*(1–2), 8–13. doi: 10.1016/j.jneumeth.2006.11.017

Pleskac, T. J., & Hertwig, R. (2014). Ecologically rational choice and the structure of the environment. *Journal of Experimental Psychology: General*, *143*(5), 2000–2019. doi: 10.1037/xge0000013

Preuschoff, K., Quartz, S. R., & Bossaerts, P. (2008). Human insula activation reflects risk prediction errors as well as risk. *Journal of Neuroscience*, *28*(11), 2745–2752. doi: 10.1523/JNEUROSCI.4286-07.2008

Preuschoff, K., 't Hart, B. M., & Einhäuser, W. (2011). Pupil dilation signals surprise: Evidence for noradrenaline's role in decision making. *Frontiers in Neuroscience*, *5*, 115. doi: 10.3389/fnins.2011.00115

RStanArm. (Version 2.9.0–4). Retrieved from cran.r-project.org/web/packages/rstanarm/

Rüsseler, J., & Rösler, F. (2000). Implicit and explicit learning of event sequences: Evidence for distinct coding of perceptual and motor representations. *Acta Psychologica*, *104*(1), 45–67. doi: 10.1016/S0001-6918(99)00053-0

Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron*, *36*(2), 241–263. doi: https://doi.org/10.1016/S0896-6273(02)00967-4

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*(5306), 1593–1599. doi: 10.1126/science.275.5306.1593

Skylark, W. J., & Prabhu-Naik, S. (2018). A new test of the risk-reward heuristic. *Judgment and Decision Making*, *13*(1), 73.

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, *53*(1), 1–26. doi: 10.1016/J.Cogpsych.2005.10.003

Stewart, N., Reimers, S., & Harris, A. J. (2015). On the origin of utility, weighting, and discounting functions: How they get their shapes and how to change their shapes. *Management Science*, *61*(3), 687–705. doi: 10.1287/mnsc.2013.1853

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*(4), 297–323. doi: 10.1007/BF00122574

van Steenbergen, H., & Band, G. (2013). Pupil dilation in the Simon task as a marker of conflict processing. *Frontiers in Human Neuroscience*, *7*, 215. doi: 10.3389/fnhum.2013.00215

von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior.* Princeton, NJ: Princeton University Press.

Walasek, L., & Stewart, N. (2015). How to make loss aversion disappear and reverse: Tests of the decision by sampling origin of loss aversion. *Journal of Experimental Psychology: General*, *144*(1), 7–11. doi: 10.1037/xge0000039

Yechiam, E., Ashby, N. J., & Pachur, T. (2017). Who's biased? A meta-analysis of buyer-seller differences in the pricing of lotteries. *Psychological Bulletin*, *143*(5), 543–563. doi: 10.1037/bul0000095