Supplementary Material for "Do People Exploit Risk–Reward Structures To Simplify Information Processing in Risky Choice?"

Christina Leuker, Thorsten Pachur, Ralph Hertwig and Timothy J. Pleskac

Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany

May 6, 2019

Contents

1 Learning phase

1.1 Setup

Participants indicated their willingness to sell (WTS) for one gamble at a time, with self-paced breaks between five blocks. To motivate participants to indicate their true valuations of the gambles, we implemented a Becker-DeGroot-Marschak auction (Becker et al., 1964): Participants entered a price at which they would be willing to sell each gamble by moving the mouse along a rating scale ($0 - 100E\$$) and confirming the value with a click. 10 gambles were played out 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 experimenters buying price exceeded the participants selling price, participants sold the gamble and earned the buying price. Otherwise, the gamble was played out (e.g., 50% chance of $38E\$$). The dominant strategy in this task is to price a gamble based on its subjective value: Setting higher prices can prevent selling unattractive gambles; lower prices can lead to selling attractive gambles under value.

1.2 Instructions

Welcome to this study! Today you are a candidate in the game show Keep or Sell? that is all about games of chance.

–

In addition to your base payment you can earn a bonus that depends on your choices.

–

Your goal in the game show is to earn as much money as possible. This money will be paid as a bonus, added to your show-up fee, at the end of the study.

–

All gambles in this experiment rely on an experimental currency, the E$. The bonus will be computed using an exchange rate (10E$ = 0.1EUR).

–

You will be presented with one gamble at a time. Each gamble consists of
- a reward (expressed in E$)
- a probability (expressed in %)

–

[gamble shown] In this example, the gamble offers a chance to win 100E$ with a probability of 50%. With a probability of 50%, you get 0 E$.

–

[gamble shown] Every gamble you see belongs to you.

–

[gamble shown] Sometimes the payoff is lower (here only 3E$). Sometimes the probability is lower (here 44%). Thus, there are more or less attractive gambles. Now imagine, that we would like to buy the gamble from you.

–

[gamble shown] We ask you to set a price for every gamble using your mouse and the scale below the gamble. You are not very keen on selling? Set a high price. You would like to sell the gamble by all means? Set a very low price (e.g. 1E$). You can move along the rating scale as often as you like.

–

[gamble shown] What is the best price you can attain for each gamble? The highest attainable price will always be between 0 and the maximum value in each gamble (here 100E$), because the reward is probabilistic. For the gamble on the left, well pay you a random price between 0 and 100E$.

–

[two gambles shown: one attractive, one unattractive] Thus, the better a gamble appears, the higher the price that you should indicate. The worse a gamble appears, the lower the price you should indicate. You can adjust the price you would like to set with the mouse. Confirm your final price by clicking on the value. Your selection will be saved.

–

[fixation cross shown] The task consists of multiple rounds. Between rounds, a fixation cross will appear on the screen. Please look at the fixation cross before moving on to the next round.

–

[many gambles shown] In each round you will be presented with a different gamble to set a price for. At the end, 10 gambles will be selected at random, for example: [selection shown]

–

[gamble shown] We will make you an offer for each of these gambles (random value between 0 and the reward in the gamble). If our offer is higher or as high as your price, you receive the price that you set. When the offer is lower, the gamble will not be sold. The gamble will be played out (e.g. 50% chance of winning 100E$).

–

[gamble shown] Your best strategy is to set the lowest price that you would like to obtain. If the price is too low, you may sell an attractive gamble underprice. If the price is too high, you may be stuck with an unattractive gamble.

–

You will learn about your bonus at the end of the experiment. The bonus from all 10 rounds will be added; and paid out to you together with your show-up fee.

–

The rounds that will be played out are random. Thus, each and every single round can affect your bonus.

–

After this part of the experiment, we will ask you some additional questions about the gambles.

–

We would like to thank you in advance for taking part in our study! Data quality is important to us, so please try to complete the tasks as faithfully as possible. If you cannot or would not like to take part anymore, you can leave the study at this point. This does not have any negative consequences for you and you will receive the full payment apart from the bonus. You may also leave the study at a later timepoint.

–

You have as much time as you want for each decision. After 20 rounds, you can take a break. You can resume the task at your own pace. The study will take 1-1.5 hours in total.

–

Good luck and enjoy the experiment “Keep or Sell!”. Please fetch the experimenter to start.

2 Choice phase

2.1 Gamble pair characteristics

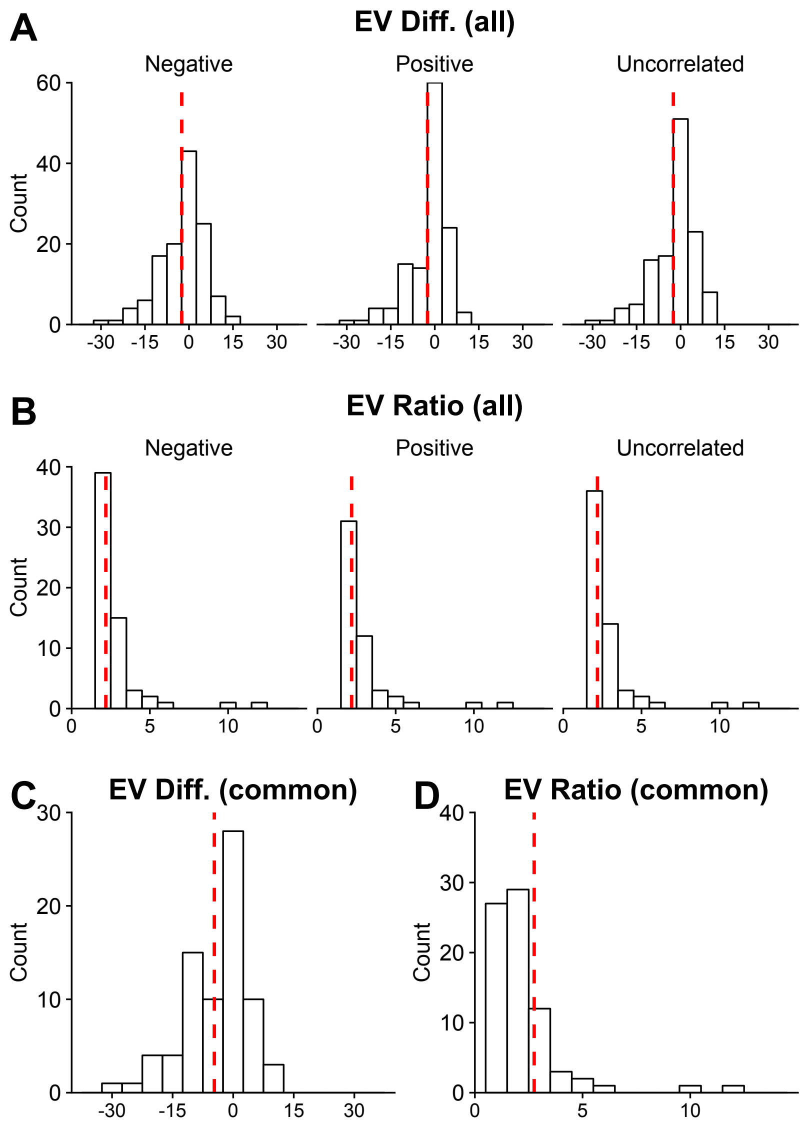

Figure S1: Characteristics of gamble pairs in the choice phase. (A, B) Some differences emerge due to the proportion of gambles that were condition–dependent—such as a larger proportion of gambles with very small EV differences in the positive condition. (C, D) Analyses were based on common gamble pairs.

2.2 Instructions

Second part of the experiment

–

In the following, we are asking you to indicate which of two gambles you prefer.

–

The gambles are equivalent to the gambles you saw in the first part of the Experiment (reward 0-100E$, probability 0-100%).

–
Again you can earn a bonus that depends on your choices. We will now describe the task in detail.
–
You will now be presented with two gambles. Each gamble consists of
- a reward (expressed in E$)
- a probability (expressed in %)
–
[two gambles shown, nondominated] Please indicate which of the two gambles you would rather play. Example:
–
Please use the ‘P and ‘Q buttons to indicate your preferences for the left and right gamble, respectively.
–
Please put your left index finger on the ‘Q button and your right index finger on the ‘P button.
–
[two gambles shown, one gamble selected] Your selection (here the right gamble, in red) will then be saved.
–
Please carefully consider which gamble you would like to select. You cannot undo this choice once you press either P or Q.
–
However, there will be two types of rounds (again, there will be multiple rounds).
- rounds, in which you need to decide as quickly as possible (timelimit 1.5s)
- rounds, in which you are asked to make the best possible choice (no timelimit)
–
These different rounds appear in blocks (block with timelimit – without timelimit – with timelimit, etc.). Before you start, you will always be informed whether the upcoming round includes a timelimit or not.

- TIMELIMIT - Make a fast decision
- NO TIMELIMIT - Make the best possible choice

–
[two gambles shown] When there is a timelimit, you need to make a good, but fast, choice within 1.5s.
–
When the timelimit is surpassed, you will receive the following message: ‘Time expired! No payoff from this round. Please decide faster.
–
You can familiarize yourself with this timelimit in the upcoming practice rounds (no bonus during practice rounds yet).
–
When there is no timelimit, you have as much time as you like to make a decision. Please make the best possible choice.
–
These types of rounds will also be shown to you during the practice round.
–
[fixation cross shown] The task consists of multiple rounds. Between rounds, a fixation cross will appear on the screen. Please look at the fixation cross before moving on to the next round.
–
[many gamble pairs shown] In each round you will be presented with a different pair of gambles to choose between. At the end, 10 choices will be selected at random, for example: [selection shown]
–
[two gambles shown, one selection] Your choices from these rounds will be played out. In this ex-

ample, a 20% chance of winning 50E$ will be played out.
–
If a gamble pair from a timepressure trial is selected and you surpassed the timelimit in this round, you will receive 0E$.
–
The rounds that will be played out are random. Thus, each and every single round can affect your bonus.
–
The earnings from each round will be summed up and presented to you at the end of the experiment.
–
In this part of the experiment, your eye movements will be recorded. –
After this part of the experiment, we will ask you some additional questions about the gambles.
–
Good luck and enjoy! Please fetch the experimenter to continue.

3 Posttests

3.1 Setup

In the posttest phase, as an independent test of how well (individual) participants had picked up and maintained the different risk-reward relationships, participants completed three tasks. First, we asked participants to choose between an uncertain option and a sure thing (sure payoff = .3 × uncertain payoff). Second, we asked participants to estimate probabilities from payoffs. Specifically, we asked them to estimate the likelihood they would win 20 randomly drawn payoffs that fell within the range of possible payoffs within the experiment. Third, we reversed this task and asked them to estimate the payoffs that had been associated with 20 different probabilities (as in Leuker et al., 2018). We use the probability estimation task to estimate how well people learned (and maintained) the condition–dependent risk–reward structures over the experiment (see statistical analyses below).

3.2 Instructions

3.2.1 Choice Task (A)

In the following, we are asking you one more time to indicate which of two gambles you prefer.
–
The gambles are structured just like the gambles in the FIRST part of the experiment.
–
That is, possible rewards can again range from 0-100E$.
–
However, one gamble will not have probability information. The probability information is covered with a "?".
–
Consider how likely you would be to win this gamble, given the first part of the experiment.
–
The other gamble is a sure outcome (100%).
–
Select the gamble that you would rather play (with "P" or "Q").
–
You will make 40 choices in total.
–
You have as much time as you would like for this task.
–
5 of your choices from these rounds will be played out at the end of the experiment.
–

The rounds that will be played out are random. Thus, each and every single round can affect your bonus.
–
The earnings from each round will be summed up and presented to you at the end of the experiment.
–
Press the space bar to begin.

3.2.2 Payoff-probability estimation task (B)

In the following, you will only see one gamble at a time.
–
The gamble will not have probability information. The probability information is covered with a "?".
–
Your task is to estimate your chances of winning the reward in the gamble.
–
The gambles are similar to the gambles in the first part of the experiment.
–
Think back of these gambles and give your best estimate.
–
Click on the scale (0-100%).
–
Indicate a 0% probability if you think that the reward is extremely unlikely.
–
Indicate a 100% probability if you think that the reward is extremely likely.
–
As soon as you are happy with your estimate, log your choice by clicking on the box with the value.
–
Press the space bar to begin.

3.2.3 Probability-payoff estimation task (C)

In the following, you will again only see one gamble at a time.
–
The gamble will not have PAYOFF information. The payoff information is covered with a "?".
–
Your task is to estimate reward magnitudes.
–
The gambles are similar to the gambles in the first part of the experiment.
–
Think back of these gambles and give your best estimate.
–
Click on the scale (0-100E$).
–
Indicate a reward of 0E$ if you think that the probability is associated with a low reward.
–
Indicate a reward of 100E$ if you think that the probability is associated with a high reward.
–
As soon as you are happy with your estimate, log your choice by clicking on the box with the value.
–
Press the space bar to begin.

3.3 Results

Figure S2 plots how often people chose the uncertain option, participants' payoff–dependent probability estimates and their probability–dependent payoff estimates. Choices were impacted by the previously experienced risk–reward structure. In comparison to our earlier study (Leuker et al., 2018), the uncertain option was chosen more often overall since in this experiment, the sure thing was made less attractive, by offering only 1/3 of the payoff that the uncertain option offered; the ratio was 1/2 in our earlier studies.

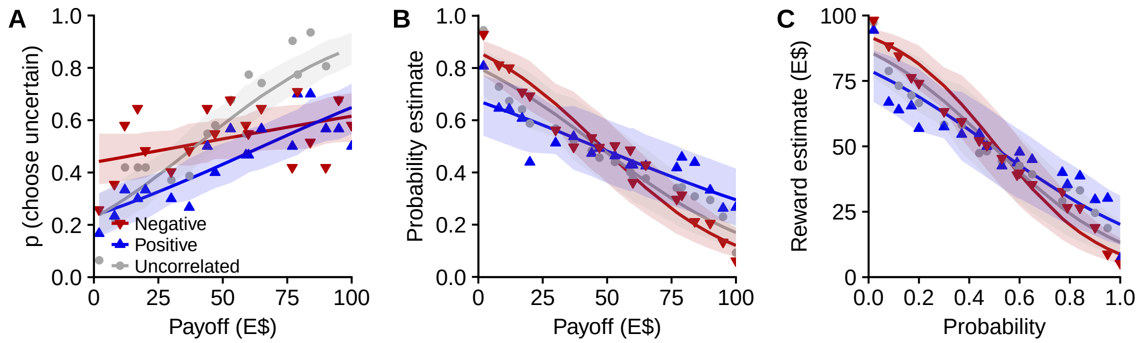

Figure S2: Posttests. (A) Participants decisions under uncertainty were impacted by the risk–reward structures they had been exposed to previously. (B, C) Payoff and probability estimates were influenced by the risk–reward structure from the incidental learning phase, but strongly biased towards an inverse relationship between risks and rewards.

The table shows the coefficients of the probability estimates in panel (C).

Condition	Slope (β)	Highest Density Interval (β)
Negative, Incidental	-0.65	$(-0.69; -0.61)$
Positive, Incidental	-0.28	$(-0.34; -0.22)$
Uncorrelated, Incidental	-0.58	$(-0.62; -0.53)$

Table S1: Risk–reward conditions reliably impacted participants estimates ($b_{uncorrelated} = .16$, CI $= [.09, .23]$, $b_{positive} = .37$, CI $= [.30, .44]$, model that compares estimates to those in the negative condition), but all estimates were biased to a negative risk–reward structure.

4 Drift diffusion modeling

4.1 Parameters

Parameter	Description
Nondecision time (τ)	τ models the part of the response time that is unrelated to the processing of the option itself (e.g., encoding, motor response).
Threshold separation (α)	Response caution/threshold. Here, the top boundary is set to the higher EV option. In the best condition, the two thresholds are set far apart, producing slower decisions but a higher rate of EV-maximizing choices. In the fast condition, boundaries are expected to be closer to the starting point, producing faster, more errant decisions that deviate from EV-maximizing choices.
Drift rate (δ)	The average change in preference at each unit of time, with $\infty < \delta_y < \infty$. The sign of the drift rate indicates the average direction of the incoming incremental change in preference, with positive values indicating preference in favor of the higher EV option. The magnitude of the drift rate characterizes the strength of the change in preference for that sample. Here, the drift rate is not a free parameter, but a linear function of the difference in EVs and the difference in the amount of time each gamble was fixated on (see below for coefficients), using a regression approach.
EV coefficient (β_{EV})	Effect of EV differences between the gambles on the drift rate. Positive values indicate that EV differences are a determinant of the drift rate. Higher values indicate stronger impact.
Gaze coefficient (β_{gaze})	Effect of gaze differences between higher and lower EV gambles on the drift rate. Positive values indicate that gaze differences are a determinant of the drift rate (i.e., a gaze "bias" towards the higher EV option). Higher values indicate stronger impact.
Interaction ($\beta_{gaze \times EV}$)	Interaction effect between gaze and EV differences as part of the drift rate. Positive values indicate a reliable interaction between gaze and value. Higher values indicate stronger interaction effect.
Relative start point (z)	Start point. If responses are unbiased, $z = .5$. Here, the higher EV option is not detectable from the beginning and the location of options has been counterbalanced. We therefore fix $z = .5$ instead of estimating it as a free parameter.

Table S2: Parameters of the extended DDM. Only τ, α, β_{gaze}, and β_{EV} are estimated. δ is a function of β_{gaze} and β_{EV} and (in model 4) of $\beta_{gaze \times EV}$; z is fixed.

4.2 General approach

We used a Bayesian Hierarchical Modeling approach to estimate individual and group parameters simultaneously. We used JAGS in R to sample from the posterior distributions. In Bayesian parameter estimation, parameter estimates are represented as prior distributions and then updated into posterior distributions based on the observed data. The advantage of a hierarchical approach is that it can account for individual variation while simultaneously pooling individual estimates into group-level distributions. The joint posterior parameter distributions were estimated using Monte Carlo Markov Chain methods implemented in JAGS, called from R. We ran 25 chains, each with 4,000 recorded samples, which were drawn from the posterior distributions after a burn-in period of 500 samples. Model selection was based on deviance information criteria, with lower values indicating better fit; and on posterior predictive checks in which we assessed how well the model aligned to behavioral results (across learning conditions).

The experiment had a risk–reward (between-participants) × timepressure (within-participants) design that we fit simultaneously. We inspected the quality of the posterior distributions by visually inspecting the mixing of the chains and autocorrelation, and on the basis of the Gelman-Rubin statistic. We compared model fits using DIC values (lower values indicate better model fit) and creating posterior predictions for choices and response times using the mean parameter values obtained from the "winning" models, on both the condition and the participant–level.

We compared four models. The first only takes into account the value differences between the options, with larger value differences leading to higher drift rates. The second model only takes into account gaze differences, with more gaze towards the higher value option leading to higher drift rates toward that option. The third third model tests for main effects of gaze and value on the drift rate. The fourth model also allows for an interaction between gaze and value (similar to Cavanagh et al., 2014), including additive effects of value and gaze. We did not include an aDDM–like model in which the drift rate *solely* depends on the interaction between value and gaze, without additive effects—the reason for this is that we wanted to partial out and compare "value sensitivity"/"value usage" for the three risk–reward conditions, which is not possible in an interaction–only model. The best–fitting model was an extended DDM in which the drift rate depended on an additive effect of gaze differences and value differences (i.e. each of which was quantified by a free parameter, or regression coefficient).

We operationalized the value of an option in two different ways. First, we used used "Expected Value" as an approximation for the better option ($Expected\,Value_i = p_i \times x_i$). Second, we allowed risk preferences of our participants to vary, and used "Expected Utility" as an approximation for the better option ($Expected\,Utility_i = p_i \times x_i^{\alpha}$). We applied the same model specifications in both approaches (just the definition of "value" is different: it either refers to EV or EU). The DDM + EU model is reported in the main manuscript. Generally, and as is shown below, both approaches support the same conclusions.

4.3 Drift diffusion model specifications

Model 1: Value differences impact drift rate

We estimated α and τ, and fix the response bias at .5. In addition, we estimated the drift rate using a regression approach, such that the drift is described by an intercept $\beta 0$ and a regression coefficient, β_{EV}. The intercept is participant–dependent and therefore represents individual differences in the ability to detect the higher value option.

$$\delta = \beta_0 + \beta_{EV} \times (EV_H - EV_L) \tag{1}$$

Model 2: Gaze differences impact drift rate

As above, but in addition, we estimated the drift rate using a regression approach, such that the drift is described by an intercept δ_0 and a regression coefficient, δ_{gaze}. The gaze coefficient models to what extent choices depend on pure gaze towards a particular option. Gaze is entered as the proportion of total gaze to the higher vs. lower EV gamble.

$$\delta = \beta_0 + \beta_{gaze} \times (gaze_H - gaze_L) \tag{2}$$

Model 3: Value and gaze differences impact drift rate

As above, but such that the drift is described by an intercept δ_0 and two regression coefficients, δ_{EV} and δ_{gaze}. The value coefficient models peoples' sensitivity to differences in expected values. The gaze coefficient models to what extent choices depend on pure gaze towards a particular option. Gaze is entered as the proportion of total gaze to the higher vs. lower EV gamble.

$$\delta = \beta_0 + \beta_{EV} \times (EV_H - EV_L) + \beta_{gaze} \times (gaze_H - gaze_L) \tag{3}$$

Model 4: Value, gaze and their interaction impact drift rate (aDDM)

As model 3, but with an interaction between value and gaze ($\delta_{interaction}$). Cavanagh et al. (2014) modeled the interaction without separate additive effects, but here we aimed to compare "EV use" across risk–reward conditions.

$$\delta = \beta_0 + \beta_{EV} \times (EV_H - EV_L) + \beta_{gaze} \times (gaze_H - gaze_L) + \\ \beta_{interaction} \times (gaze_H \times EV_H - gaze_L \times EV_L)(4)$$

4.4 Results: Drift diffusion modeling with EV

4.4.1 Comparison of DICs

Model	DIC
DDM 1: Drift: EV Diff.	36406.23
DDM 2: Drift: Gazediff.	36632.80
DDM 3: Drift: EV Diff. + Gaze Diff.	36064.39
DDM 4: Drift: EV Diff. + Gaze Diff. + EV × Gaze	36179.30

Table S3: Deviance Information Criterion (DIC) for four different formalizations of the Drift Diffusion model using the gamble with the higher Expected Value at the threshold. Model 3 has the lowest DIC.

4.4.2 Parameter estimates and posterior predictions for the winning model

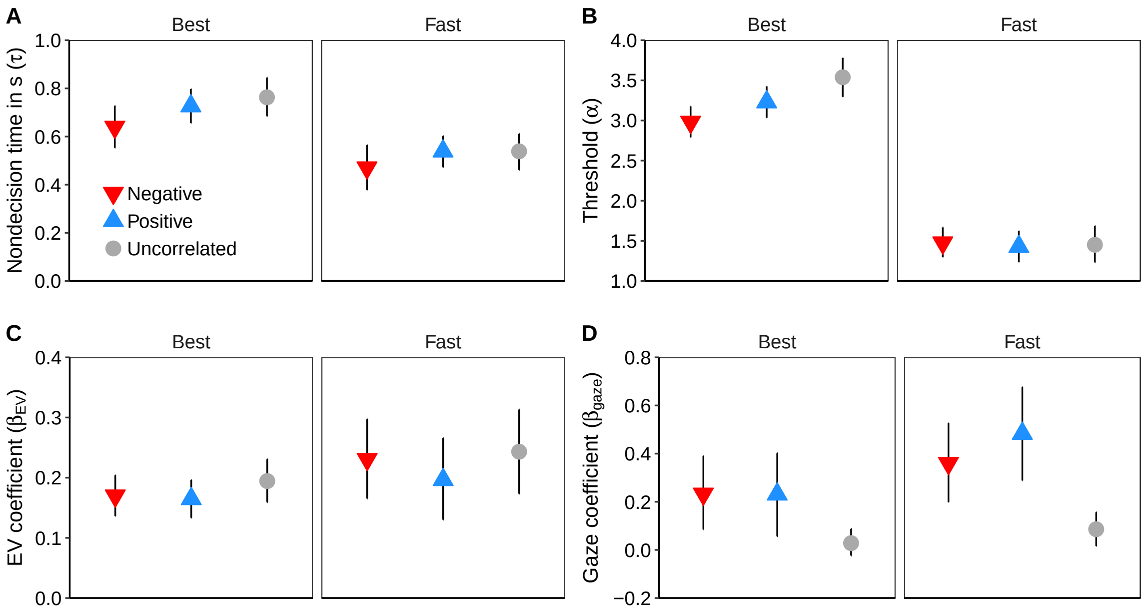

Figure S3: Parameter estimates after incidental learning for the winning model (Model 3), specified using higher and lower EV gambles. Group means and 95% highest density intervals of the posterior distributions. This plot is presented for the sake of completeness and also appears in the main manuscript.

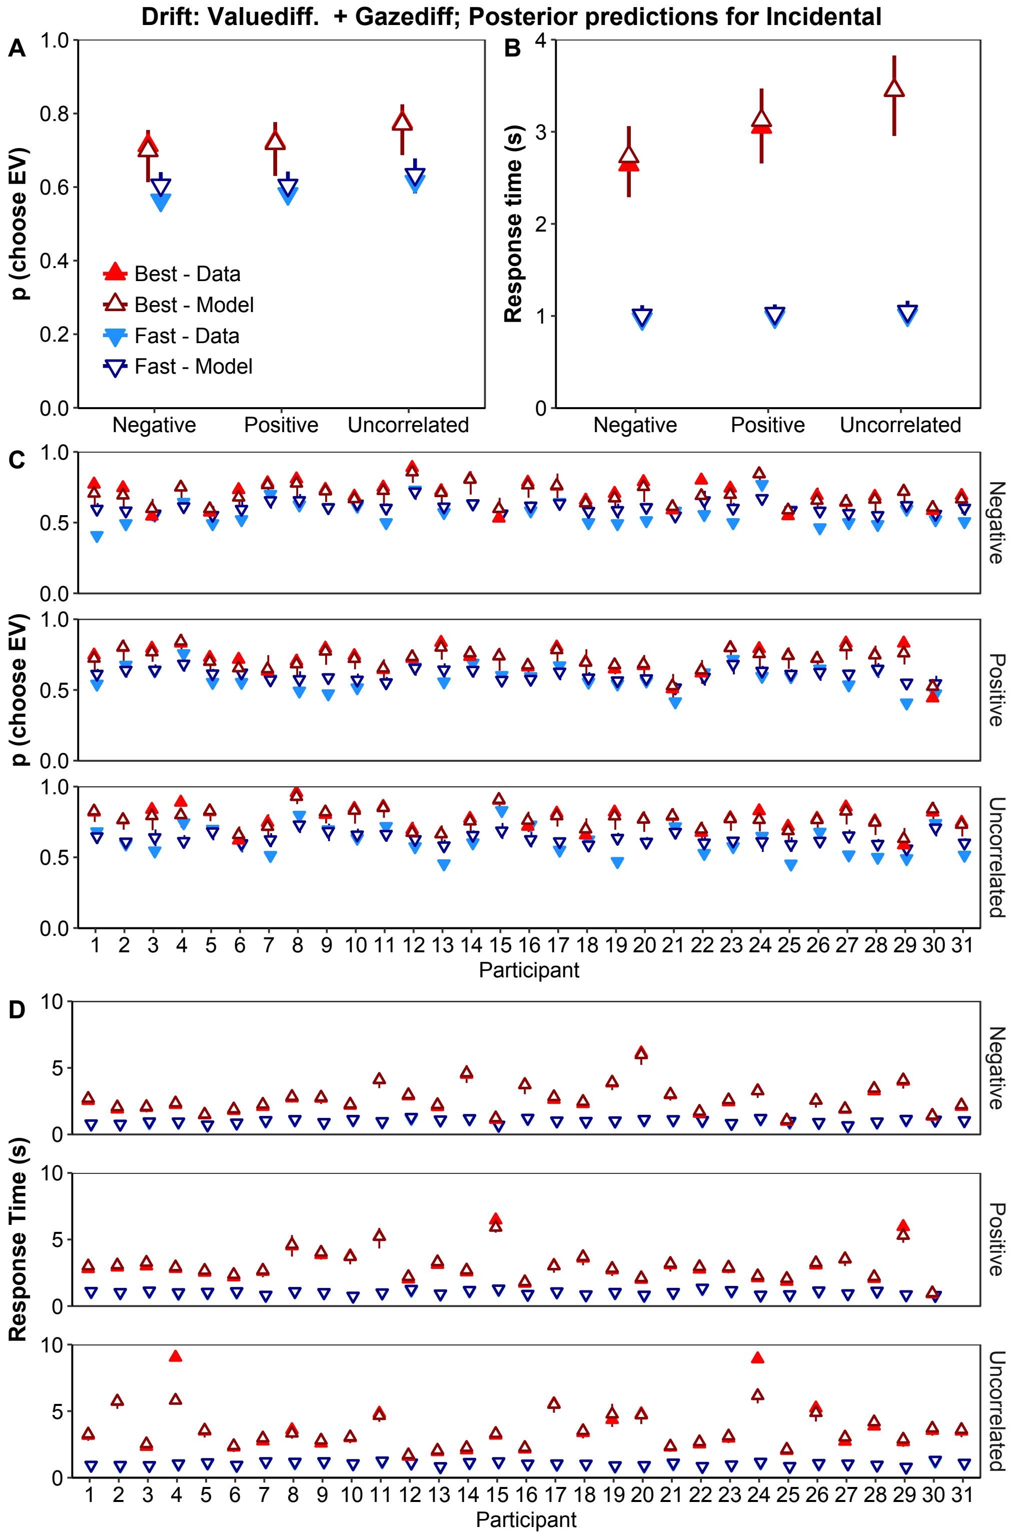

4.4.3 Model–choice data link

Beyond DICs and posterior predictive checks, one way to assess how well the model describes the data is to link individual differences in parameter estimates back to individual differences in the proportion of EV choices. Figure S7 shows that there were plausible links between individual parameter estimates and individual differences in the proportion of times the higher expected–value option was chosen.

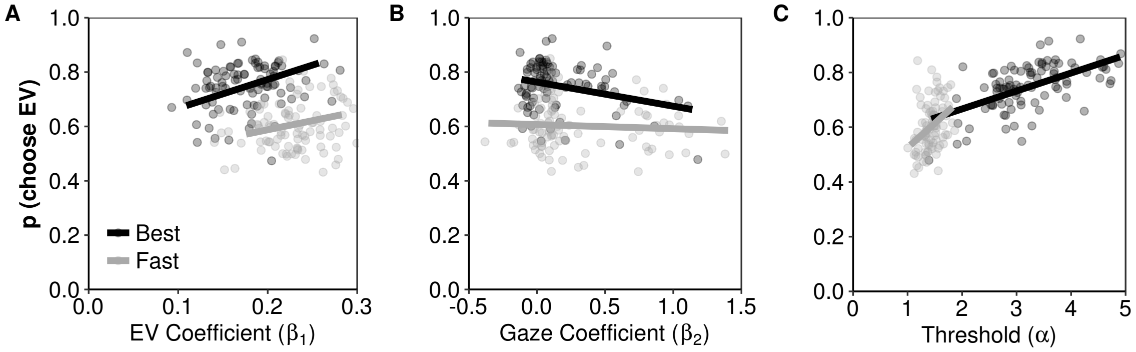

Figure S4: Relationship between DDM parameters and behavioral choice results. Each dot represents one participant. (A) Participants who were sensitive to EV differences chose the higher EV option more often. (B) Participants who distributed their attention more evenly (indicated by a gaze coefficient of 0) chose the higher EV option more often. (C) Participants who set higher thresholds chose the higher EV option more often.

4.5 Results: Drift diffusion modeling with EU

4.5.1 Risk preferences across risk–reward conditions

We used a Bayesian Hierarchical Modeling approach to estimate individual and group parameters for an Expected Utility model using a power value function on the payoff ($Expected\ Utility_i = p_i \times x_i^{\alpha}$), where $\alpha < 1$ signals risk averse; $\alpha = 1$ signals risk neutral and $\alpha > 1$ signals risk seeking behavior. We also implemented a Luce Choice rule, with lower θ values suggesting less deterministic choices.

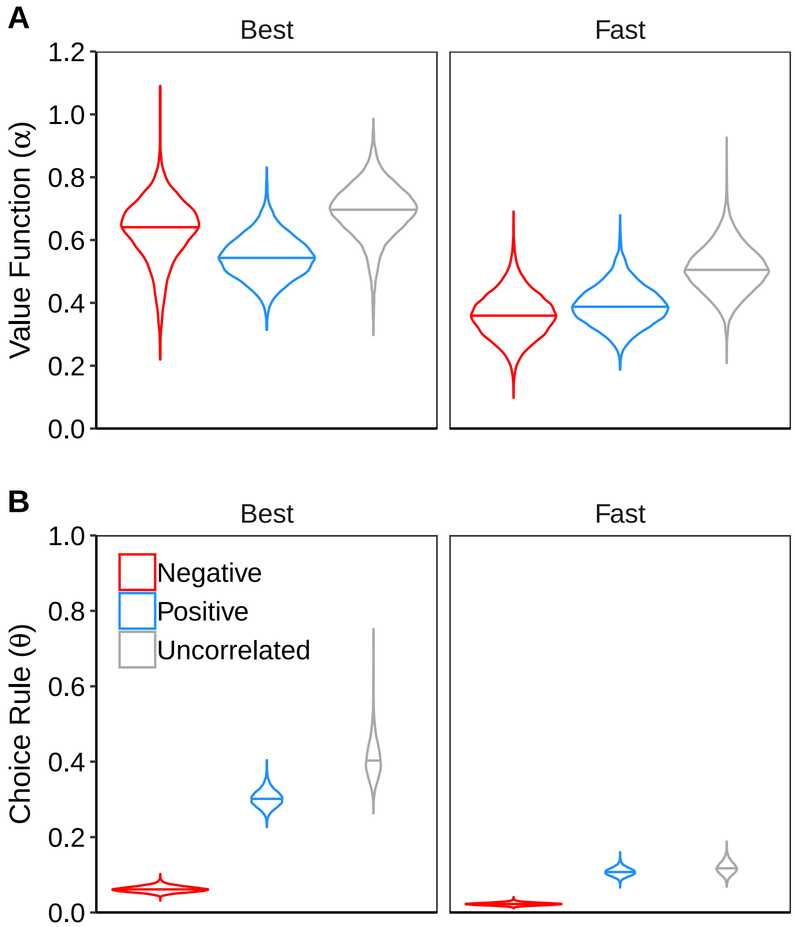

Figure S5: Parameter estimates suggest little differences in risk preference across the risk–reward conditions (A). Instead, choices appeared less deterministic in the negative risk–reward environment compared to the other two environments (B).

We used individual parameter estimates from using the alpha parameter from the model to determine an individuals' subjectively "better" option. Next, we conducted the same analysis as before, but used "choosing the higher EU gamble" as the correct answer at the top boundary. Note that the two additional parameters from the EU model are not accounted for in the DICs below.

4.5.2 Comparison of DICs

Model	DIC
Model 1: Drift: Valuediff.	37492.37
Model 2: Drift: Gazediff.	34207.44
Model 3: Drift: Valuediff. + Gazediff.	33943.35
Model 4: Drift: Valuediff. + Gazediff. + Value × Gaze	33991.42

Table S4: Deviance Information Criterion (DIC) for four different formalizations of the Drift Diffusion model using the gamble with the higher Expected Utility at the threshold.

4.5.3 Parameter estimates and posterior predictions for the winning model

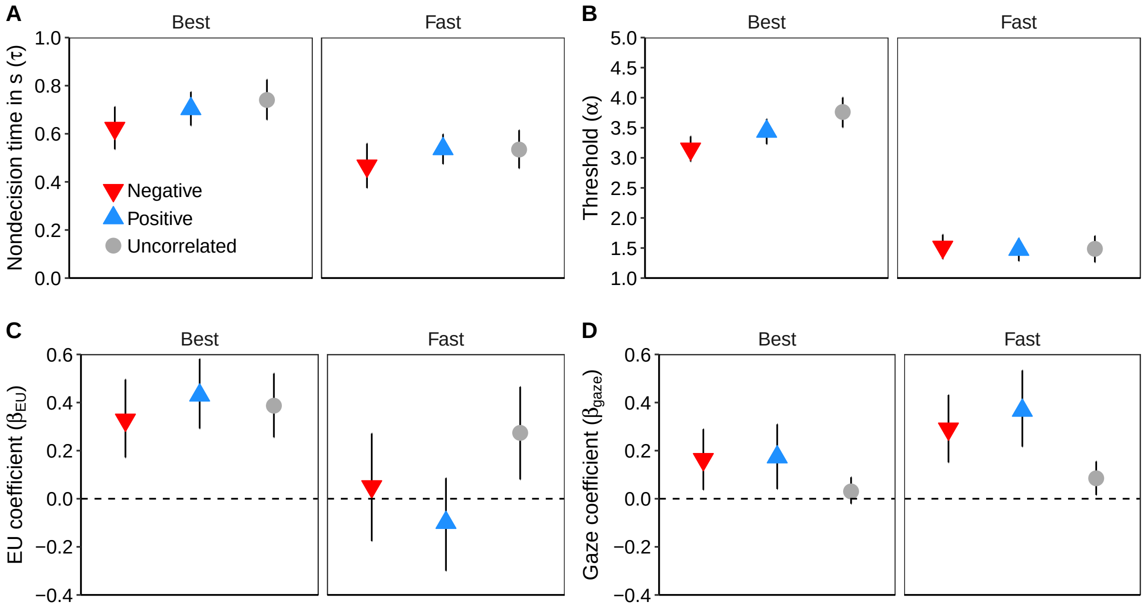

Figure S6: Parameter estimates after incidental learning for the winning model (Model 3), specified using higher and lower EU gambles. Group means and 95% highest density intervals of the posterior distributions.

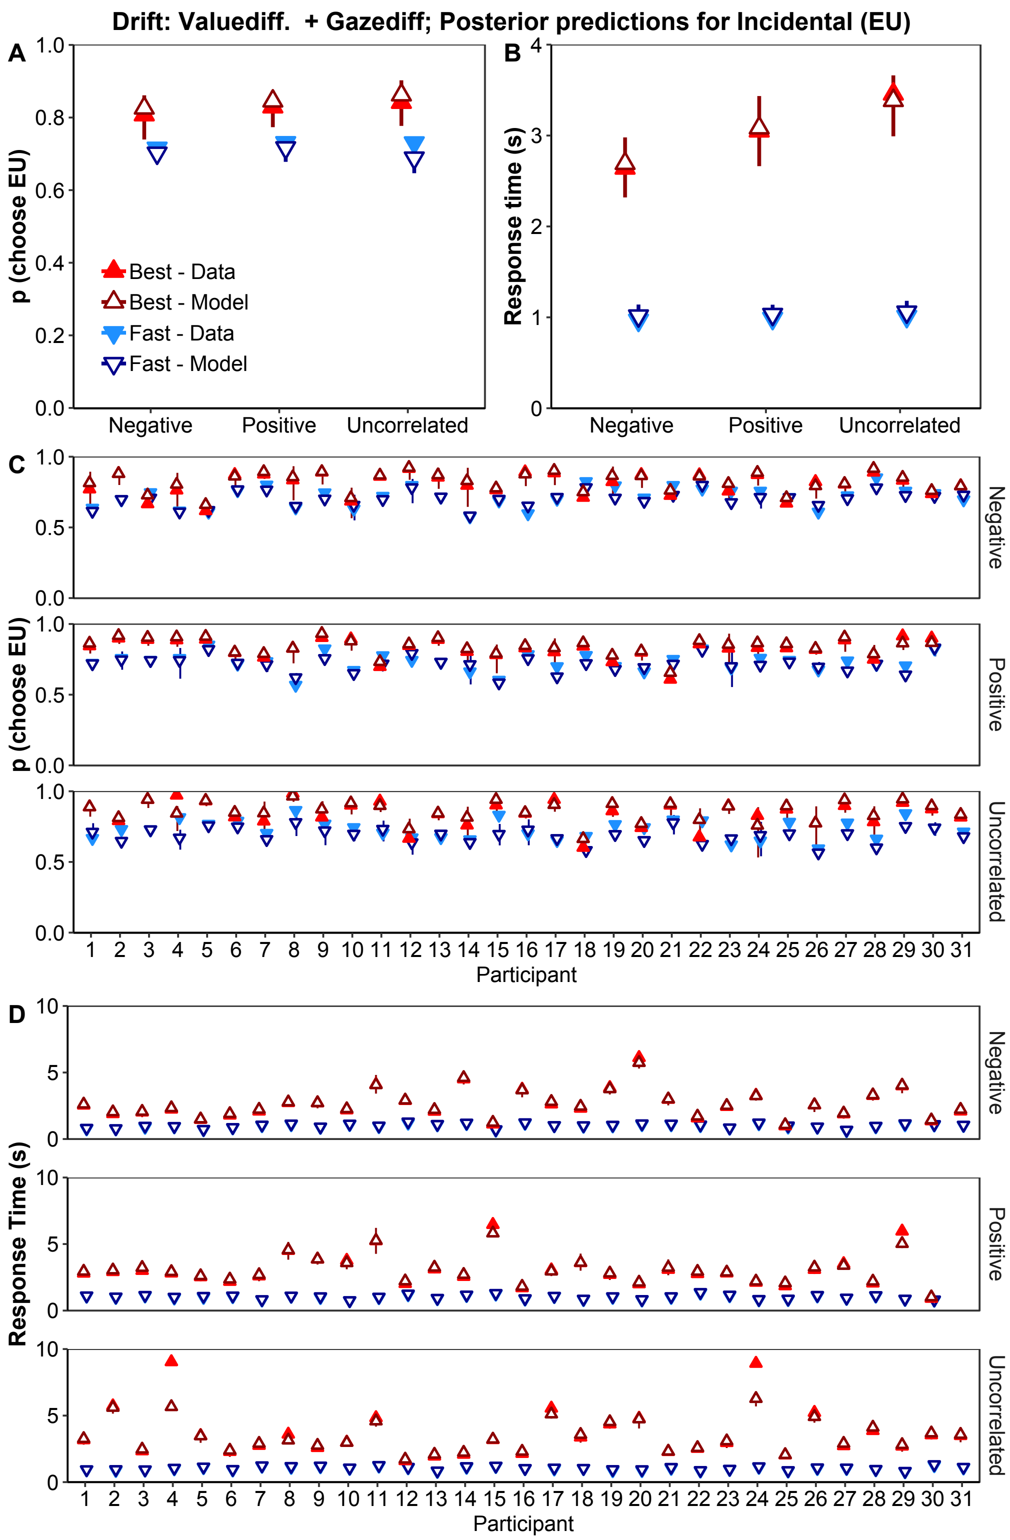
Drift: Valuediff. + Gazediff; Posterior predictions for Incidental (EU)
A
B
C
D
p (choose EU)
Response time (s)
Response Time (s)
Best - Data
Best - Model
Fast - Data
Fast - Model
Negative
Positive
Uncorrelated
Participant

4.5.4 Model–choice data link

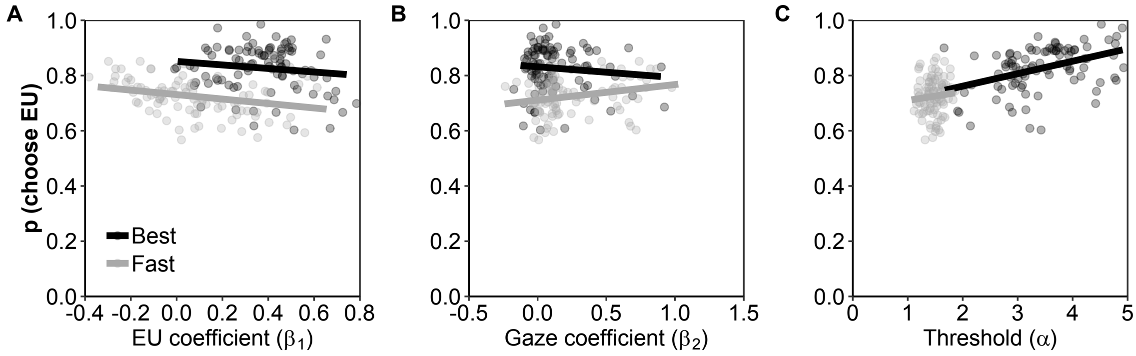

Figure S7: Relationship between DDM parameters and behavioral choice results. Each dot represents one participant. (A) Participants who were sensitive to EU differences did not necessarily choose the higher EU option more often. This result is inconsistent with the result from the EV model. (B) Participants who distributed their attention more evenly (indicated by a gaze coefficient of 0) chose the higher EU option more often, especially in the best condition. (C) Participants who set higher thresholds chose the higher EU option more often.

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

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

Leuker, C., Pachur, T., Hertwig, R., & Pleskac, T. J. (2018). Exploiting risk–reward structures in decision making under uncertainty. *Cognition*, *175*, 186–200. doi: 10.1016/j.cognition.2018.02.019