Supplementary Material for "The ecology of competition: A theory of risk–reward environments in adaptive decision making"

Timothy J. Pleskac[1,2], Larissa Conradt[2], Christina Leuker[2] and Ralph Hertwig[2]

[1]University of Kansas
[2]Max Planck Institute for Human Development, Berlin, Germany

July 23, 2020

Contents

1 Risk–Reward Estimation Study

Our goal was to replicate the risk–reward estimation studies in Pleskac & Hertwig (2014) with a German sample. We used the scenario (in German) described in Study 2. Moreover, we manipulated the order in which participants were asked to estimate how many balls they thought were in the basket and if they wanted to pay to play the gamble. Half of the participants completed the estimation question before the choice question, while the other half completed it in the other order. We also extended the range of payoffs (payoff levels, between participants: €2.5, €4, €10, €50, €100, €200, €500, €1,000, and €10,000).

1.1 Methods

1.1.1 Participants

A total of 455 participants were recruited using a web panel hosted at the Max Planck Institute for Human Development in Berlin, resulting in approximately 50 participants per payoff × order condition (between-participants). Each participant earned €2.0 for participation. The study was approved by the IRB at the Max Planck Institute for Human Development in Berlin.

1.2 Results

1.2.1 Estimation

	Coefficient	**95 % HDI**
Intercept (Mean)	-1.45	[-1.60, -1.29]
Reward (Mean)	-1.48	[-2.12, -0.79]
Order:Choice first (Mean)	-0.19	[-0.41, 0.03]
Order:Choice first × Reward (Mean)	1.43	[0.56, 2.28]
Intercept (Dispersion)	0.34	[0.14, 0.53]
Reward (Dispersion)	1.35	[0.50, 2.10]
Order:Choice first (Dispersion)	0.48	[0.20, 0.76]
Order:Choice first × Reward (Dispersion)	-2.10	[-3.11, -1.05]

Table S1: Bayesian beta regression table using normalized rewards (0–1), estimates (0–1), and question order (baseline: estimates, then choice) as predictors for both the mean and the dispersion of the estimates.

We modeled the effect of payoff on the mean and dispersion (variance) in the estimates using a beta regression with Bayesian estimation techniques (Kruschke, 2014), using Stan in R for regression analyses with the rstanarm package (*RStanArm Version 2.9.0-4*, 2016). In the regression, we used the reward value (normalized to be between 0 and 1), question order (estimate first vs. choice first), and the interaction between reward value and order, as predictors via a logit link of the estimates (normalized to be between 0 and 1). In addition, we used the same variables to simultaneously predict the dispersion (a proxy for the variance) of the estimates (Smithson & Verkuilen, 2006) via a log link.

Results of the beta regression are summarized in Table S1. Replicating past results, and consistent with participants using the risk–reward heuristic, as rewards increased the probability estimates decreased. There was also a credible order effect on the estimates as well as an interaction between order and reward level: Participants who chose whether or not to play first (and then estimated the chances of winning) generally showed a diminished effect of reward on their estimates. This also meant that participants who chose first appeared to estimate their chances of winning average, higher.

In terms of the dispersion, there was a credible decrease in the dispersion of the estimates as payoff increased (due to a negative log-link function). As with the mean estimates, there was a credible order effect on the dispersion of the estimates as well as an interaction between order and reward level. For participants who chose first (and then estimated the chances of winning), there was more dispersion in the estimates. Moreover, the effect of reward value changed from showing a credible decrease when estimating the probability first ($b = 1.38$, CI = $[0.52, 2.13]$) to a credible increase when making a choice first ($b = -0.76$, CI = $[-1.41, -0, 14]$; dispersion coefficient from a beta regression for estimate first condition).

It is difficult to draw strong conclusions about the order effect as the study was not designed to examine this effect. Nevertheless, the risk–reward estimation model (described in detail in section 2) may help understand these results. In this particular study, the risk–reward estimation model with an ignorance prior component showed a small increase the rate of participants using the principle of indifference (i.e. estimating a probability of .5) when choosing whether to play first compared to when estimating before choosing: Analyzing the data from Study 3 alone showed that the rate of the principle of indifference was $M = .28$ (HDI = $[.12, .41]$) for participants who chose before estimating, and $M = .24$ (HDI = $[.07, .37]$) for participants who estimated before choosing (difference: $M_{choose>estimate} = .04$, HDI = $[-.03, .12]$). An increased rate of using the ignorance prior for participants who chose first is consistent with a diminished effect of the reward on the mean probability estimate and can also lead to an apparent increase in dispersion at the higher levels of reward.

1.2.2 Choice

	Coefficient	**95% HDI**
Intercept	-0.52	[-0.89, -0.16]
Reward	1.25	[0.35, 2.19]
Estimate	1.32	[0.15, 2.51]
Order:Choice first	0.69	[0.18, 1.20]
Order:Choice first × Reward	0.62	[-0.90, 2.29]
Order:Choice first × Estimate	-1.15	[-2.88, 0.58]

Table S2: Logistic regression table using normalized rewards (0–1), estimates (0–1), and question order (baseline: estimates, then choice) as predictors for both the mean and the variance estimates.

We also examined the choices participants made in terms of whether or not to pay to play. We did so using a logistic regression with Bayesian estimation techniques (Kruschke, 2014). The logistic regression was run using Stan in R for regression analyses with the rstanarm package (*RStanArm Version 2.9.0-4*, 2016). In the regression, we entered the normalized reward value, the normalized estimates, whether participants entered an estimate first (1) or choice first (0), and the interaction between order and the reward value and the interaction between the order and the estimate as predictors of the choice. The results are summarized in Table S2. They largely replicate the studies reported in Pleskac & Hertwig (2014). On average, participants were more likely to choose to pay

€2.0 to play the gamble as the reward value increased. Consistent with participants using the risk–reward heuristic to help decide whether to pay to play, their estimates were also credible predictors of their choice. Finally, note there was a credible order effect such that participants that chose first were more likely to pay to play.

2 Risk–Reward Estimation Model

2.1 Risk–reward model components

The risk–reward estimation model predicts the estimates participants provide in a pay-to-play gamble like the risk–reward estimation task. The model consists of three core assumptions: that people use the risk–reward heuristic, that ecological distortion of probability estimates occurs, and that estimates vary. These aspects are developed in the main article.

In addition, we had a hypothesis that, at least in the estimation task, some people might employ the principle of indifference. We came to this hypothesis when we noted that there was a small but noticeable cluster of estimates at .5 in the estimation task, which is consistent with people in the face of uncertainty treating it as equally likely to occur or not (Keynes, 1948; Laplace, 1814/1902). There is good evidence that people sometimes rely on this rule when judging the probability of an event (Fox & Clemen, 2005; Fox & Levav, 2004; Fox & Rottenstreich, 2003). With this in mind, we modified the risk–reward estimation model to capture the rate at which people use the principle of indifference rather than relying on the risk–reward heuristic. From an ecological perspective, we might expect an increased rate of the use of the principle of indifference when the environmental conditions are not expected to create a risk–reward structure (e.g., when the system is out of equilibrium or when the resource is unlimited). According to the model, with probability $0 \leq \lambda \leq 1$ people rely on the principle of indifference; otherwise they are assumed to use the risk–reward heuristic

$$p''(x) = \lambda \times .5 + (1 - \lambda) \times p'(x). \tag{1}$$

In the models that estimated the rate at which people used the risk–reward heuristic, we set the mean of the beta to $\mu = p''(x)$.

2.2 Model estimation

We estimated several different models in a model comparison. To do so, we estimated the posterior distributions over the parameters of the hierarchical models using Markov Chain Monte Carlo (MCMC) methods. Our model estimation procedure was implemented with JAGS 4.3 using Matlab via matjags to interface with JAGS. The general JAGS code used is in the OSF. The estimation of the model used three parallel chains. Each chain consisted of 1,000 burn-in steps (unrecorded samples to allow the chain to reach the reasonable parameter space) and 20,000 samples for a total of 60,000 samples.

The chains were evaluated for representativeness and accuracy using the procedures outlined by Kruschke (2014). Representativeness was evaluated using visual inspection of trace plots of the chains and density plots. All the chains at the group level were inspected visually with random samples of chains from the individual level. Representativeness was also evaluated numerically using the Gelman-Rubin statistic with the conventional heuristic that values of the Gelman-Rubin statistic above 1.1 were worrisome. All the chains met these standards suggesting representativeness of the posterior distributions. Accuracy was evaluated by examining the autocorrelation and the effective sample size. The effective sample size estimates the sample size of the chain after accounting for the autocorrelation present in the samples. As a rough standard we sought to have approximately an effective sample size of approximately 10,000. Our main focus in this paper was on comparing the group level mean estimates of the parameters.

2.3 Model comparisons

As a means for testing for the importance of different components of the RREM, we fit several variants. We compared the fits of the different models with the Deviance Information Criterion (DIC; Spiegelhalter et al., 2002) where smaller values are better. Table S3 lists the DICs for the different variants of the risk–reward estimation model where different models were included (1) or not (0). The comparisons show that the best model is the model with the distortion component and a change in dispersion, but does not include an ingorance prior component. We report this model in the paper.

It is also informative that a beta regression on the estimates where the mean and variance are allowed to vary as a function of the payoff level has a DIC of -1,1139.2 That is, the Risk–Reward Estimation Model provides a better fit to the data than a statistical model, demonstrating the importance of accounting for the cognitive strategies used in making estimates.

Distortion	**Dispersion**	**Ignorance Prior**	**DIC**
1	1	1	-1,234.2
1	1	0	-1,237.3
1	0	1	-1,162.9
0	1	1	-1,069.3
1	0	0	-1,232.9
0	1	0	924.7
0	0	1	-1,060.0
0	0	0	942.9

Table S3: DICs for different variants of the Risk–Reward Estimation Model.

3 Apartment Study

3.1 Methods

We aimed for 160 participants for each of the 5 conditions who were paid 1,00GBP for their time. On average, participants took 8.5 minutes to complete the survey (IQR 6.4 – 11.5 minutes). Inclusion criteria were being a British citizen, fluent in English (self-rated) and an approval rate of 80% of studies previously completed on the platform (Prolific Academic). The final sample consisted of $N = 160$ participants who saw the new market condition first, $N = 156$ who saw the nonideal condition first, $N = 144$ who saw the unsaturated condition first, $N = 155$ who saw the standard condition first and $N = 159$ who saw the unequal competitors condition first (Ns are slightly uneven due to the randomization mechanism that did not account for dropouts). All participants saw the remaining four conditions in random order. The wording of the scenarios in these conditions is reported in the main manuscript. No participant was excluded from the analysis.

3.2 Results

	Coefficient	**95% HDI**
Intercept / Standard (Mean)	-0.24	[-0.26, -0.22]
Apartment Quality (Mean)	-0.63	[-0.67, -0.59]
New Market (Mean)	-0.09	[-0.12, -0.06]
Nonideal (Mean)	-0.29	[-0.32, -0.25]
Unsaturated (Mean)	0.01	[-0.01, 0.04]
Unequal Abilities (Mean)	0.03	[0.01, 0.06]
Apartment Quality x New Market (Mean)	0.14	[0.09, 0.20]
Apartment Quality x Nonideal (Mean)	-0.25	[-0.32, -0.18]
Apartment Quality x Unsaturated (Mean)	0.34	[0.30, 0.39]
Apartment Quality x Unequal Abilities (Mean)	0.36	[0.31, 0.40]
Intercept / Standard (Dispersion)	0.49	[0.43, 0.56]
Apartment Quality (Dispersion)	0.63	[0.53, 0.73]
New Market (Dispersion)	-0.37	[-0.46, -0.28]
Nonideal (Dispersion)	0.08	[-0.01, 0.17]
Unsaturated (Dispersion)	-0.29	[-0.39, -0.20]
Unequal Abilities (Dispersion)	-0.20	[-0.30, -0.10]
Apartment Quality x New Market (Dispersion)	0.11	[-0.03, 0.25]
Apartment Quality x Nonideal (Dispersion)	-0.02	[-0.16, 0.13]
Apartment Quality x Unsaturated (Dispersion)	-0.07	[-0.22, 0.08]
Apartment Quality x Unequal Abilities (Dispersion)	-0.10	[-0.25, 0.05]

Table S4: Within-participant results for the apartment study. Estimates from a Bayesian beta regression fit with *stan* (*RStanArm Version 2.9.0-4*, 2016). Consistent with the between-participant results, participants adjusted their mean estimates across different conditions; but there were no reliable effects on the dispersion (see interaction effects between condition and apartment quality).

References

Fox, C. R., & Clemen, R. T. (2005). Subjective probability assessment in decision analysis: Partition dependence and bias toward the ignorance prior. *Management Science*, *51*(9), 1417–1432. doi: 10.1287/mnsc.1050.0409

Fox, C. R., & Levav, J. (2004). Partition-edit-count: Naive extensional reasoning in judgment of conditional probability. *Journal of Experimental Psychololgy: General*, *133*(4), 626–42. doi: 10.1037/0096-3445.133.4.626

Fox, C. R., & Rottenstreich, Y. (2003). Partition priming in judgment under uncertainty. *Psychological Science*, *14*(3), 195–200. doi: 10.1111/1467-9280.02431

Keynes, J. M. (1921). *A treatise on probability.* London, United Kingdom: Macmillan and Co.

Kruschke, J. K. (2014). *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan* (2nd ed.). New York, NY: Academic Press/Elsevier.

Laplace, P. S. (1902/1814). *A philosophical essay on probabilities.* Wiley and Sons. (Original work published 1814)

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

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

Smithson, M., & Verkuilen, J. (2006). A better lemon squeezer? maximum-likelihood regression with beta-distributed dependent variables. *Psychological Methods*, *11*(1), 54–71. doi: 10.1037/1082-989X.11.1.54

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *64*(4), 583–639.