

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Prospect Theory and Optimal Risky Choices with Goals

Permalink

<https://escholarship.org/uc/item/6845g1g5>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 42(0)

Authors

Jarecki, Jana B.

Rieskamp, Jörg

Publication Date

2020

Peer reviewed

Prospect Theory and Optimal Risky Choices with Goals

Jana B. Jarecki (jana.jarecki@unibas.ch, jj@janajarecki.com)
Center for Economic Psychology, University of Basel, Missionsstrasse 62a,
4055 Basel, Switzerland

Jörg Rieskamp (joerg.rieskamp@unibas.ch)
Center for Economic Psychology, University of Basel, Missionsstrasse 62a,
4055 Basel, Switzerland

Abstract

Decision making under risk is often studied as a preferential choice governed by stable individual personality characteristics, but risky choice can also be viewed as a dynamic problem of resource accumulation to survive. When decision makers aim to reach a particular goal in limited time, such as “earn at least \$100 in five choices,” risky choice becomes a non-trivial planning problem. This problem has an optimal solution that can differ from immediate expected-value maximization. We studied the optimality of risky choices under such minimum goal requirements experimentally and find that the observed choices under goals approximate the optimal solution. However, because the optimal model is very complex, we examine if simpler models can predict people’s choices better. We test an extended version of prospect theory, assuming a dynamic reference point that depends on the distance to the goal. This “dynamic prospect theory” was better than the alternative model in describing people’s decisions (i.e., for 63% of the participants, it was the best model). Our findings show that humans can excel in a highly complex, dynamic, risky choice problem and that a dynamic version of prospect theory provides one possible explanation for how people decide under risk when long-term goals matter.

Keywords: risky choice; energy budget rule; risk sensitivity; goals; choice modeling

Introduction

There is a long tradition in testing different theories to explain decision making under risk. These theories examine whether concepts besides people’s risk preferences need to be considered to provide an accurate description of the cognitive processes underlying risk-taking. However, this work has mainly ignored the situational aspects, such as long-term goals or requirements of the decision maker. Recent work has shown that situational goals, such as “you need \$100 in five risky choices”, significantly affect risky choices. High goals cause behavior change from risk-aversion (variance avoidance) to risk-seeking (Fujimoto & Takahashi, 2016; Korn & Bach, 2018; Mishra, 2014). We investigate whether people adaptively take goals into account in their risk-taking behavior and how optimal their behavior is. Furthermore, we explore how an extended, dynamic version

of cumulative process theory (Tversky & Kahneman, 1992) can describe people’s risky choices in comparison to alternative models.

In the context of evolutionary survival strategies, behavioral biologists have long argued that changes in risk taking with higher goals are fitness-maximizing behavior (Houston & McNamara, 1988). If energy needs are high, a safe option may yield too little energy to meet a minimum energy requirement. Decision makers should become risk-seeking and thereby bet on the high outcome of a riskier energy resource, given equal expected values of the options. This became known as risk-sensitive foraging (Caraco, 1981; Real & Caraco, 1986). Goal-dependent changes in risk taking may thus be an optimal response to meet resource goals, in line with a biological fitness-maximizing model for risky choice under requirements (Houston & McNamara, 1988).

Interestingly, people tend to show the optimal behavior. In a recent study on risk-taking with requirements, people indeed behaved in a near-optimal fashion (Korn & Bach, 2018). By contrast, in other higher-order cognitive tasks, optimal decision making is a rare result (e.g., Hecce Castañón et al., 2019). And animals, for instance, when engaging in risky choices with various sizes of goals, are not consistently optimal (Caraco, 1981; Kacelnik & El Mouden, 2013). Importantly, the optimal solution to risk-taking with requirement problems is computationally complex (see below), and people do not adhere to the predictions by the optimal model deterministically. Some authors have recently hypothesized that an extension of Kahneman and Tversky’s prospect theory (1979; 1992) can explain the observed behavior in risky choice tasks with goals (McDermott et al., 2008), which is a yet untested claim. Furthermore, it has been questioned on theoretical grounds to what extent cumulative prospect theory can represent decision making according to risk sensitive foraging theory (Houston et al., 2014). Therefore, our work examines a dynamic extension of prospect theory as a cognitive model to describe human decision making under risk with goals and a finite time horizon.

Models of finite-horizon risky choice with goals

Our optimal benchmark model for risky choice with goals and limited time formulates the choice as a discrete-time dynamic programming problem and solves it by backward induction (Houston & McNamara, 1988). The optimal choice in trial t depends on the number of remaining trials, the options, the goal g , and on the currently accumulated point state s . Figure 1 illustrates the model's solution strategy: The model spans the tree of all possible future states¹ and probabilities until after the very last trial. Given the goal g , the model computes a final reward based on the obtainable states after the very last trial in T ; usually, the reward function is binary (other forms are possible):

$$R(s, g) = \begin{cases} 0, & s < g \\ 1, & s \geq g \end{cases} \quad (1)$$

The model then determines the expected reward for all states s in the decision trial preceding the very last trial, $T-1$, for both options. The expected reward (ER) of the risky option is:

$$ER_r(o_R, s, T-1) = \sum_i R(x_i + s, g) \times p_i, \quad (2)$$

where x_i and p_i denote the risky gamble's i th outcomes and probabilities, g is the goal. The expected reward of the safe option, ER_s , in $T-1$, is computed analogously.

The model maximizes the expected reward for each state s in trial $T-1$, which gives the optimal action a^* :

$$a^*(s, T-1) = \text{Max}\{ER_r, ER_s\} \quad (3)$$

The model assumes a deterministic choice of a^* in the last trial $T-1$ for all states s . Given a^* in $T-1$, it determines the expected reward in the preceding trial $T-2$ for all states s for both options and defines the action a^* that maximizes this expected reward for each state in $T-2$. Given the resulting

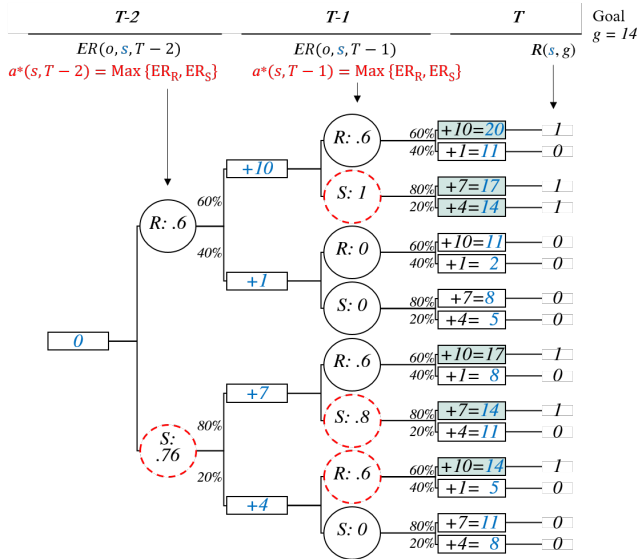


Figure 1. Illustration of the optimal solution and backward induction for a 2-trial task with a goal of 14 points and two options (R = 10 with 60% or 1; and S = 7 with 80% or 4).

¹Mathematically it suffices to span the tree using the unique states

deterministic optimal choices in $T-2$, the model determines the expected reward and optimal choice for all states in $T-3$. This process of backward optimal action selection given optimal future choices repeats until the model arrives at the expected rewards in trial 1, which it maximizes. The optimal solution is risk-neutral and finds the maximum expected reward.

Figure 1 exemplifies this optimal model. Assume a risky option that yields +10 with 60% or +1 and a safe option that yields +7 with 80% or +4, and a time horizon of 2, and a current state s of 0. In the final trial ($T-1$), four states are possible (1, 4, 7, 10), and for each of those, the model derives expected reward for both options (circles), which it maximizes (dashed red circles). For instance, in state +10 in $T-1$, the safe option S has an expected reward of 1, which exceeds the risky option's reward of .6. In the trial before, which is the first trial ($T-2$), the expected reward of the safe option, conditional on selecting the optimal choice in $T-1$, equals .76, which exceeds the expected reward of .60 of the risky option. The optimal response for trial 1 is "safe." The algorithm can compute the optimal choice (best action policy) for each possible time horizon, starting state, given gambles.

Computational Complexity. The optimal model requires the full future state-and-outcome tree (Figure 1). The number of final nodes in this tree grows exponentially with the number of trials, options, and outcomes (the curse of dimensionality, Bellmann (1961)). For a choice between two two-outcome gambles with 2 trials, the tree has 16 final nodes; with 3 trials, the number of final nodes equals 253; and with 5 trials, it equals 4,294,967,296. Beyond the optimal model's poor scalability, after each change in the trial count and each change in the state, the optimal model either requires the re-computation of the tree or the full tree must be held in memory as a lookup table. Therefore, the optimal model is cognitively complex from a computational resource standpoint, rendering it less plausible that people go through the necessary computations to solve the task.

Dynamic Prospect Theory. Prospect theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992) assumes risk aversion for gains (outcomes above a reference point) and risk-seeking for losses (outcomes below a reference point). If people would take the goal as their reference point, prospect theory can describe shifts from risk-averse to risk-seeking behavior with higher goals (McDermott et al., 2008). This goal-based reference-point theory, however, is not dynamic and fails to track a continually changing resource state (Houston et al., 2014). To model multiple-trial choices with requirements, we will make the novel contribution of extending prospect theory to assume a dynamic reference point that depends on the remaining number of trials ($T-t$), and the distance of current state s to the distant goal g :

$$r_t = \frac{(g-s)}{T-t}.$$

in each trial t ; for illustrative purposes we used the full tree here.

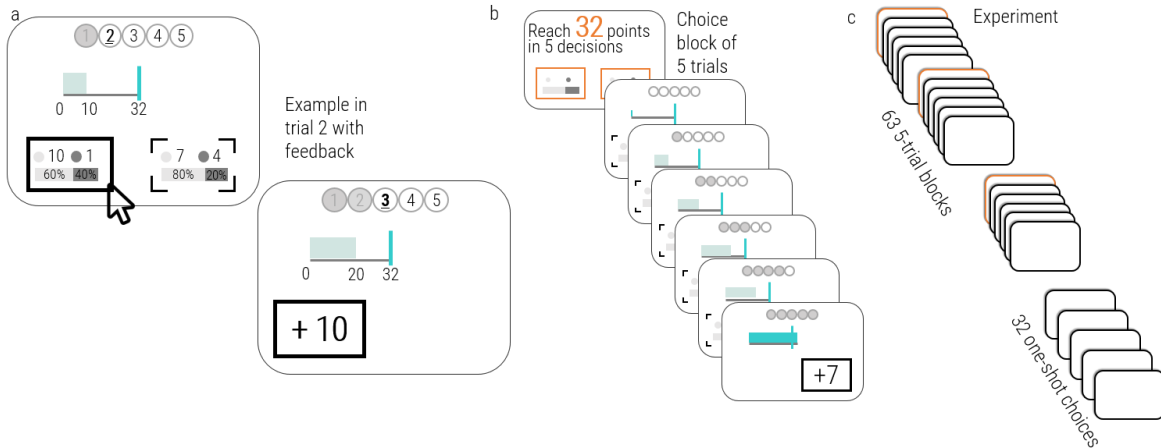


Figure 2. The 5-trial risk and requirements task. (a) Example choice in a block with a goal of 32. Shown is trial 2 of 5 in which the left gamble is chosen. After the choice, outcome feedback appears based on a draw from the chosen gamble’s outcome distribution; the outcome raises the point total, shown as bar, from 10 to 20, and the trial counter increases from 2 to 3. (b) A 5-trial block is characterized by a goal, and two gambles. (c) Experimental procedure, see text.

The possible outcomes x of a gamble are then transformed relative to the reference point r_t as $x - r_t$. We modeled the choices with the five-parameter version of prospect theory (power utility, specified as in Tversky and Kahneman, 1992), in which the utility of a gamble g is given by

$$u(g) = \sum_i \pi(p_i) v(x_i),$$

where x is the i^{th} outcome of a gamble minus the dynamic reference point, r_t . The value function is given by $v(x) = x^\alpha$ for positive or zero x and $-\lambda(-x)^\beta$ for negative x . The probability weighting function is $w(p_i) = p^\gamma / (p^\gamma + (1-p)^\gamma)^{1/\gamma}$ and we used two probability weighting parameter, γ^+ for positive and zero x and γ^- for negative x . Based on $w(p_i)$, the cumulative probability weight is then given by $\pi(p_i) = w(\sum_{x_i \geq x} p_i) - w(\sum_{x_i > x} p_i)$ for positive and zero x , and $w(\sum_{x_i \leq x} p_i) - w(\sum_{x_i < x} p_i)$ for negative x .

The idea behind the dynamic reference point is that, all else equal, high goals or a short time horizon set higher reference points. For high reference points, gain outcomes are more likely to lie below the reference point (a loss compared to the reference point), and prospect theory predicts risk-seeking behavior. In contrast, lower goals and long time horizons set lower reference points. For low reference points, gain outcomes are more likely to exceed the reference point, and prospect theory predicts risk-aversion.

Model complexity. Although prospect theory is a more complex model in comparison to the optimal model in terms of its number of parameters—it has five parameters compared to zero parameters in the optimal model—in terms of the processing complexity, the optimal model scales less well than prospect theory. Prospect theory does not suffer from the curse of dimensionality, because it requires the computation of the reference point and a transformation of outcomes and probabilities, which scale linearly with growing numbers of outcomes and probabilities.

Experiment

To examine risky choices with requirements, we conducted an incentivized choice experiment.

Participants. Sixty participants recruited from Prolific Academic completed an online study, one participant was excluded (for self-reported low data quality), leaving a final sample of $N = 59$; 29 males, 28 females and 2 preferred not to state (49%, 47%, and 3%, respectively), the mean age was 29 years ($Med = 25$, $SD = 9$, range 19-58 years), data were collected from December 2018 to January 2019, the study was approved by the ethics committee of the Department of Psychology at the University of Basel. The mean study duration was 40 minutes ($Med = 36$, $SD = 26$). The study was incentivized (five randomly drawn blocks were paid, given that in a block, the goal was reached).

Procedure. Participants engaged in a risky choice under requirement task (Figure 2, following Pietras, Searcy, Huitema, & Brandt (2008)). The task consisted in collecting points to reach a point goal by making five choices among a pair of risky gambles. The gambles were two-outcome gambles in the gain domain; risky gambles had higher variances than safe gambles (the Appendix lists the stimuli). The goal, point status, the gambles, and a trial counter were displayed on the screen, as shown in Figure 2a. After participants made their choice, one outcome was drawn from the chosen gamble and displayed on the screen, and the point total increased by the value of the outcome. Each block consisted of five choices between the same gambles, after which the point total had to meet the goal (Figure 2b). After a 5-trial block, the point total was re-set to zero; no points were carried over between blocks. New blocks involved various gambles and goals; all blocks started with 0 points; all blocks consisted of 5 trials to reach the goal. We presented nine combinations of goals and gambles, repeated seven

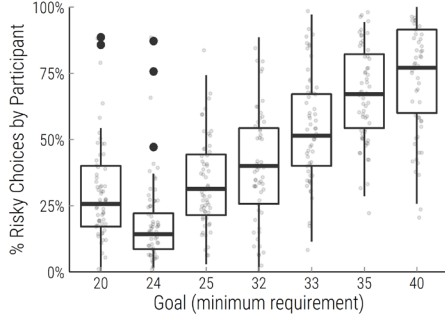


Figure 4. Increase in Risky Choice with the Goal

times, resulting in 63 blocks (Figure 2c). The block order was randomized.

Further, participants made 32 “one-shot” choices: The stimuli for these choices were a combination of a state, a goal, a trial number, and options (like Figure 2a). Participants made precisely one choice, asked to act as if they found themselves in this situation, after which no outcome feedback appeared, and they did not finish the block. The one-shot choices were not incentivized. The reason for including these choices was that during the 5-trial phase, the states that participants experienced were not comparable between participants, except for the initial state. We wanted comparable choices to test model predictions.

The goals in the experiment had various difficulties, measured as best chance to reach the given goal by the end of the block according to the optimal strategy (the expected reward in the first trial after backward induction, Eq. 3) Easy

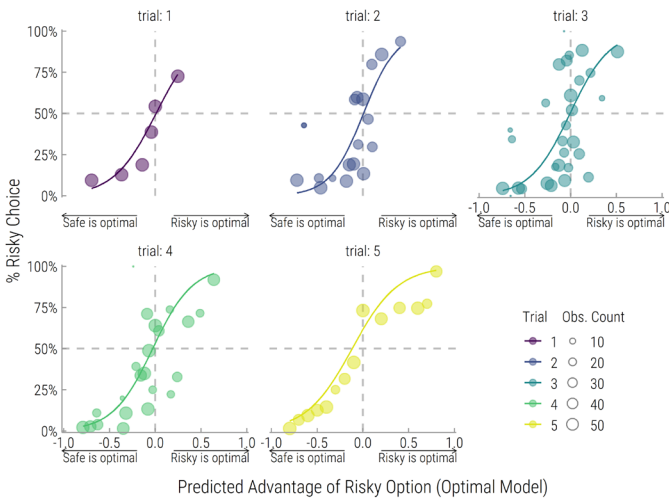


Figure 3. The observed risky choice share as a function of the advantage (see text) of the risky option as predicted by the optimal model. Lines are best-fitting logit regressions; colors separate trials in the 5-trial blocks.

Obs. Count = observation count at each advantage.

² Mixed-effects logit regression of the risky choices with the predictor goal and a by-subject random intercept and slope for goal.

³ Excluding trials in which the value of the risky equals the value of the safe option according to the optimal model, because in these trials both risky and safe choices are optimal. Including these trials, the observed choices in the 5-choice phase are 85% optimal. An

blocks in the 5-trial phase had 74%-92% chance to reach the goal, medium blocks had 60%-68%, and hard blocks had only 44%-48%. The chances in the one-shot choices ranged from 6% to 96%; for the one-shot choices, difficulties were given by the expected reward in the state and trial that were shown.

Results

The analyses were conducted in R (v3.6.3, R Core Team, 2019) using the lme4 package (v1.1, Bates et al., 2015); cognitive modeling used the cognitivemodels package (v0.0.7, Jarecki & Seitz, 2020).

Behavioral Results. Figure 3 shows that with higher goals, participants shifted from risk-averse to risk-seeking behavior. The proportion of risky choices in the blocks with goals as low as 20 ($M=.29, SD=.17$) or 24 ($M=.17, SD=.17$) were lower than the risky choice proportions in the blocks with a high goal of 40 ($M=.75, SD=.20$); with a significant effect² of goal on risky choice, $F(1, 59) = 254, p < .001$. This goal-dependent change in risk-seeking behavior is in line with previous results (Fujimoto & Takahashi, 2016; Korn & Bach, 2018; Pietras et al., 2008). The subsequent analysis will analyze the data by goal, trial, point total, against the optimal model as a gold standard.

Approximate Optimality of Human Risky Choices. Overall, 75% of the observed choices in the 5-trial blocks equaled the optimal choice—across all experienced point totals, goals, and trials.³ In the one-shot phase, 56% of the choices equaled the optimal model. For a detailed analysis of the choice optimality, we contrasted the observed risky choices with the predicted metric advantage of the risky option over the safe option from the optimal model. To this end, we used the data from the 5-trial blocks. The advantage of options depends on the expected reward (ER) of both options by state and trial (Eq. 2, the chance to meet or exceed the goal in the remaining trials if choosing an option now and following the optimal policy strictly after that). For all experienced states and trials, the expected rewards of a risky choice ER_R and a safe choice ER_S were compared, yielding the advantage of the risky choice, $a = ER_R - ER_S$. Higher values favor the risky option more strongly, values of 0 mean no advantage. Optimal agents would deterministically select “risky” for a > 0 , and select “safe” for a < 0 , else randomize.

Figure 4 shows that human behavior follows the computed advantage of the risky option not in a deterministic but in a soft-max fashion (S-shaped curve). For zero advantage, the choice proportions are close to what we expect from random choices, which is the optimal behavior. A regression analysis⁴ showed that a higher advantage significantly increased the

optimal value maximizing decision maker would achieve 70% accuracy.

⁴ Mixed-effect logit regression; predicting choices from the fixed-effect predictors advantage + difference in the expected values (risky-safe), and the random effect predictors by-participant random intercept and slope for advantage, using a logit link.

risky choice likelihood ($b = 4.52$, 95% CI [3.79, 5.29], $p < .001$); risky choices also depended on the expected value of the risky compared to the safe gamble ($b = 0.27$, 95% CI [0.25, 0.31], $p < .001$). A comparison of models showed that including the trial number as a predictor improved the model fit relative to the increase in model complexity ($AIC_{+trial} = 19,669 < AIC_{-trial} = 19,812$ and $BIC_{+trial} = 19,724 < BIC_{-trial} = 19,859$, $\chi^2 = 144$, $p < .001$). Removing the (fixed-effect) of the optimal advantage of the risky option as a predictor while keeping the expected value of the risky compared to the safe option worsened the model fit ($AIC_{+adv} = 19,669 < AIC_{-adv} = 19,740$ and $BIC_{+adv} = 19,724 < BIC_{-adv} = 19,787$, $\chi^2 = 72$, $p < .001$).

The result that risky choice tracks the optimal advantage of the risky option is particularly interesting from a cognitive processing perspective. Notably, at no point during the task were participants informed about the numerical advantage of the risky option, but they processed the trial count, the gambles, the state, and the goal. Three aspects are worth highlighting in this respect: Foremost, deriving the optimal solution to the risky choice under requirement problem is computationally complex, as outlined in the introduction. As a second point, it is also clear that participants' choices are imperfect compared to the optimal choices, approximating optimality in a soft-max fashion but not maximizing deterministically, in line with probabilistic choice models (Rieskamp, 2008). Thirdly, and potentially surprisingly, the observed risky choices in trial five do not seem to become more deterministic than in earlier trials, which is interesting because, in the last trial, the choice problem is easier than in the first trial.

Model Comparison. To test how well prospect theory describes the risky choices, we compared the extension of prospect theory to an extension of the optimal model. Because participants' choices smoothly followed the advantage of the risky option, we specified the optimal model with a soft-max choice rule (Sutton & Barto, 2018) to select among the expected rewards of the risky and safe option in Eq. 3 (Korn & Bach, 2018). The optimal model thus had one free choice rule parameter. Prospect theory also used the soft-

max choice rule and had a total of six free parameters. The model comparison included a random-choice model predicting $Pr(risky)=0.5$ as a baseline. All free model parameters were estimated on the data in the 5-trial phase by maximum likelihood at the individual level; and the one-shot choices were used for out-of-sample predictions. To this end we compared the models regarding their predictive performance in the one-trial choices. The predictions of the optimal model were derived, assuming the block would be played until the fifth trial. The predictive performance was assessed using evidence weights from the log likelihood of the models (Wagenmakers & Farrell, 2004) in an individual-level strategy classification in which an evidence strength > 0.5 defines the winning model. The model comparison results differ for fitting and out-of-sample prediction (Table 1). The results from the out-of-sample prediction of the one-shot choices show that $n=37$ participants were best predicted by the extended prospect theory, and only $n=11$ participants were best-described by the soft-max optimal model; the baseline model predicted 11 participants best. Table 1, which shows the mean model fits aggregated across participants, reveals that although the optimal model with a soft-max choice rule outperforms the dynamic prospect theory in fitting (due to the parameter-based complexity of CPT), prospect theory predicts substantially better.

Table 1. Performance of the Models

Model	N_{Par}	Fit		Prediction of	
		to 5-trial phase		one-shot phase	
		AIC	BIC	AICw	LLw
Dynamic CPT	6	358	911	0.41	0.57
OPT+soft-max	1	356	354	0.59	0.19
Baseline	0	437	437	0.01	0.24

Note: N_{Par} = number of free parameters, including the soft-max choice rule parameter. AIC/BIC = Akaike/Bayesian Information criterion, AICw = Akaike weight, values of 1 indicate strong relative evidence for a model, 0 indicates weak evidence (Wagenmakers & Farrell, 2004), LLw = like AICw based on log-likelihood for the out-of-sample prediction. CPT = prospect theory, OPT = optimal model.

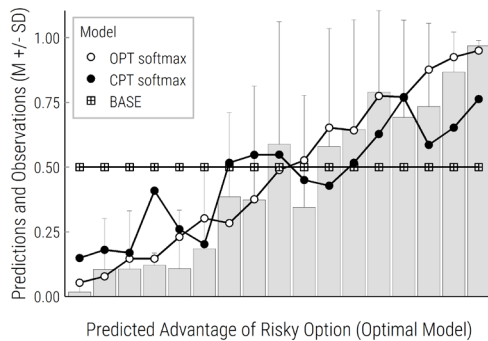


Figure 5. Model predictions against observed data (bars) error bars are SDs; CPT=prospect theory, OPT=Optimal model

Figure 5, which plots the predicted against the observed risky choices, shows that, at high advantage levels, the CPT model under-predicts the average riskiness slightly. The optimal model over-predicts only certain trials with intermediate advantage levels. The result that prospect theory outperforms the optimal model in out-of-sample prediction, despite the greater number of free parameters compared to the optimal model, might be interpreted as that people fail to calculate the complicated optimization on the fly when confronted with a single description-based scenario. By contrast, if participants had experience with the trial and their past choices, their performance became closer to the optimal

model.

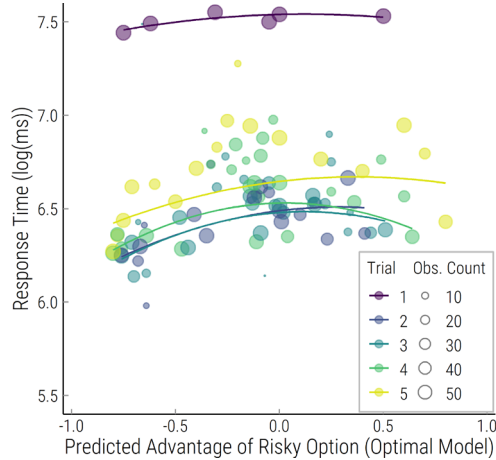


Figure 6. Response times (log(ms)) by advantage. Obs. Count = observation count at each advantage.

Response Times. Figure 6 shows the response times (log-transformed). The response times in trial 1 of the 5 available trials in a block ($M=1,845$ milliseconds, back-transformed) were slower than the decisions in the later trials 2, 3, 4, and 5 ($M_s = 628$ ms, 622 ms, 653 ms, 734 ms, respectively), with significant pairwise differences of the estimated regression⁵ coefficients $\Delta b_{1-2} = 1.077, p < .0001$; $\Delta b_{1-3} = 1.085, p < .0001$; $\Delta b_{1-4} = 1.032, p < .0001$; $\Delta b_{1-5} = 0.900, p < .0001$. Further, the response times were also faster the higher the advantage of either the risky or the safe option, and the advantage had a quadratic effect on response times ($b = -0.246, p < .001$), and this regression specification outperformed a regression with no quadratic term ($AIC_{+quadr} = 41,247 < AIC_{-quadr} = 41,280$ and $BIC_{+quadr} = 41,341 < BIC_{-quadr} = 41,366$; $\chi^2 = 34, p < .001$). In line with research on value-based decisions, we can interpret the inverted-u-shape of the response time curves as a proxy for choice difficulty: the trials with small or zero advantages were more complicated to process than trials where one of the options was the better choice.

Regarding concerns about the quality of response-time data from online studies, we designed the study following principles for high response time accuracy in web-based studies using CSS animations, which can cause a slight bias (30-100ms over-estimation) but low noise in the response times ($SD < 10$ ms) in modern browsers (Garaizar & Reips, 2018; Reimers & Stewart, 2015); and these measurement uncertainties are low compared to the 1,000 ms differences in response times between trial 1 and the remaining trials, that we found in our data.

Discussion

In this study, we have shown that goals have a significant

⁵ Mixed-effect logit regression; predicting $\log(\text{response times})$ from advantage + difference in expected values + trial + (trial)²; and

effect on risky choice. In the context of reaching a choice goal by making risky decisions, people can approximate a complex optimal planning solution for risky choice tasks that involve various goals and a 5-trial time horizon. Consistent with predictions by the optimal model for the task, people dynamically selected the risky option more often as the expected reward of the risky option increased relative to the expected reward of the safe option. A dynamic extension of prospect theory could describe slightly less than half of the participants better than the optimal choice model from behavioral biology, and prospect theory predicted the majority of the participants better than the optimal model. Theoretically, given the computational demands of the optimal model, prospect theory seems to provide a model with higher plausibility, at least in the early trials, where the optimal model is more complicated to compute than in later trials. However, qualitatively the dynamic prospect theory model could not consistently describe the aggregated choices (Figure 5), suggesting that other heuristics might describe the processes of goal-directed risky choice even better.

In general, goal-dependent changes in risk preferences are in line with a broader biobehavioral view on human risk-taking. Following this view, the cognitive processes underlying decision making under risk have evolved, at least partly, as a solution to resource and requirement problems. Following the biobehavioral view, an essential under-studied component of people's risk-taking is its function for the decision-maker. Risky decisions may be fruitfully viewed as a means to an end, for example, to reach financial, personal, or social requirements. We believe that the interplay of risk attitudes in interaction with situation-specific demands provides exciting new research avenues for risk researchers.

References

- Bates, D., Mächler, M., Bolker, B. M., Walker, S. C., Maechler, M., Bolker, B. M., & Walker, S. C. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. <https://doi.org/10.18637/jss.v067.i01>
- Bellmann, R. E. (1961). *Adaptive Control Processes: A Guided Tour*. Princeton University Press.
- Caraco, T. (1981). Energy budgets, risk and foraging preferences in dark-eyed juncos (*Junco hyemalis*). *Behavioral Ecology and Sociobiology*, 8(3), 213–217.
- Fujimoto, A., & Takahashi, H. (2016). Flexible modulation of risk attitude during decisionmaking under quota. *NeuroImage*, 139, 304–312. <https://doi.org/10.1016/j.neuroimage.2016.06.040>
- Garaizar, P., & Reips, U. (2018). Best practices: Two Web-browser-based methods for stimulus presentation in behavioral experiments with high-resolution timing requirements. *Behavior Research Methods*, 2. <https://doi.org/10.3758/s13428-018-1126-4>
- Herce Castañón, S., Moran, R., Ding, J., Egner, T., Bang,

with random effects by-participant intercept and slope for advantage; Tuckey p -value adjustment

- D., & Summerfield, C. (2019). Human noise blindness drives suboptimal cognitive inference. *Nature Communications*, 10(1), 1–11. <https://doi.org/10.1038/s41467-019-09330-7>
- Houston, A. I., Fawcett, T. W., Mallpress, D. E. W., & McNamara, J. M. (2014). Clarifying the relationship between prospect theory and risk-sensitive foraging theory. *Evolution and Human Behavior*, 1–6. <https://doi.org/10.1016/j.evolhumbehav.2014.06.010>
- Houston, A. I., & McNamara, J. M. (1988). A framework for the functional analysis of behaviour. *Behavioural and Brain Science*, 11, 117–163. <https://doi.org/10.1017/S0140525X00053061>
- Jarecki, J. B., & Seitz, F. I. (2020). *Cognitivemodels – an Extendable R Software Package to Develop and Test Cognitive Models, Submitted to ICCM 2020*. <https://github.com/JanaJarecki/cognitivemodels>.
- Kacelnik, A., & El Mouden, C. (2013). Triumphs and trials of the risk paradigm. *Animal Behaviour*, 86(6), 1117–1129. <https://doi.org/Kacelnik & El Mouden, 2013>
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(3), 263–291. <https://doi.org/10.2307/1914185>
- Korn, C. W., & Bach, D. R. (2018). Heuristic and optimal policy computations in the human brain during sequential decision-making. *Nature Communications*, 9(2018), 1–41. <https://doi.org/10.1038/s41467-017-02750-3>
- McDermott, R., Fowler, J. H., & Smirnov, O. (2008). On the Evolutionary Origin of Prospect Theory Preferences. *The Journal of Politics*, 70(2), 335–350. <https://doi.org/10.1017/S0022381608080341>
- Mishra, S. (2014). Decision-making under risk: Integrating perspectives from biology, economics, and psychology. *Personality and Social Psychology Review*, 18(3), 280–307. <https://doi.org/10.1177/1088868314530517>
- Pietras, C. J., Searcy, G. D., Huitema, B. E., & Brandt, A. E. (2008). Effects of monetary reserves and rate of gain on human risky choice under budget constraints. 78, 358–373. <https://doi.org/10.1016/j.beproc.2008.01.016>
- R Core Team. (2019). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. <http://www.r-project.org/>
- Real, L., & Caraco, T. (1986). Risk and foraging in stochastic environments. *Annual Review of Ecology and Systematics*, 17(1), 371–390. <https://doi.org/10.1146/annurev.es.17.110186.002103>
- Reimers, S., & Stewart, N. (2015). Presentation and response timing accuracy in Adobe Flash and HTML5/JavaScript Web experiments. *Behavior Research Methods*, 47(2), 309–327. <https://doi.org/10.3758/s13428-014-0471-1>
- Rieskamp, J. (2008). The probabilistic nature of preferential choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(6), 1446–1465. <https://doi.org/10.1037/a0013646>
- Stott, H. P. (2006). Cumulative prospect theory’s functional menagerie. *Journal of Risk and Uncertainty*, 32(2), 101–130. <https://doi.org/10.1007/s11166-006-8289-6>
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: an introduction* (2nd ed.). MIT Press.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4), 297–323. <https://doi.org/10.1007/BF00122574>
- Wagenmakers, E., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, 11(1), 192–196. <https://doi.org/10.3758/BF03206482>

Appendix

Table A1. Stimuli used in the Experiment

Difficulty	Goal	Safe		Risky		EV _s	EV _R	Vars	Varr		
		x ₁	Pr(x ₁)	x ₂	x ₁					Pr(x ₁)	x ₂
Easy	24	5	0.9	4	7	0.2	1	4.9	2.2	0.09	5.76
Easy	32	7	0.8	4	10	0.6	1	6.4	6.4	1.44	19.44
Easy	40	8	0.8	2	10	0.8	0	6.8	8.0	5.76	16.00
Medium	24	5	0.7	4	7	0.2	1	4.7	2.2	0.21	5.76
Medium	25	5	0.8	4	9	0.4	2	4.8	4.8	0.16	11.76
Medium	35	7	0.7	5	9	0.7	2	6.4	6.9	0.84	10.29
Hard	20	4	0.8	3	8	0.2	0	3.8	1.6	0.16	10.24
Hard	33	7	0.7	3	9	0.6	1	5.8	5.8	3.36	15.36
Hard	35	7	0.7	3	10	0.6	1	5.8	6.4	3.36	19.44