What's in the Adaptive Toolbox and How Do People Choose From It? Rational Models of Strategy Selection in Risky Choice

Florian Mohnert
Max Planck Institute for Intelligent Systems, Tübingen, Germany
University of Amsterdam, Netherlands

Thorsten Pachur
Max Planck Institute for Human Development, Berlin, Germany

Falk Lieder
Max Planck Institute for Intelligent Systems, Tübingen, Germany
Bernstein Center for Computational Neuroscience, Tübingen, Germany

Abstract

Although process data indicate that people often rely on simplifying processes when choosing between risky options, current models of heuristics cannot predict people's choices very accurately. To address this apparent paradox, it has been proposed that people might adaptively choose from a toolbox of simple strategies. But which strategies are contained in this toolbox? And how do people decide when to use which decision strategy? Here, we develop a model according to which the decision maker selects a decision strategy for a given choice problem rationally from a toolbox of strategies; the content of the toolbox is estimated for each individual decision maker. Using cross-validation on an empirical data set, we find that this model of strategy selection from a personal adaptive toolbox predicts people's choices better than any single strategy (even when it is allowed to vary across participants) and better than previously proposed toolbox models. Our model comparisons show that both inferring the content of the toolbox and rational strategy selection are critical for accurately predicting people's risky choices. Furthermore, our analysis reveals considerable individual differences in the set of strategies people are equipped with and how they choose among them; these individual differences could partly explain why some people make better choices than others. These findings represent an important step towards a complete formalization of the notion that people select their cognitive strategies from a personal adaptive toolbox.

Keywords: decision making; bounded rationality; strategy selection; heuristics; computational modeling

Introduction

How do people make decisions under risk? This question is commonly studied by asking people to choose between gambles as in "Would you prefer a 20% chance of winning \$1000 (Gamble A) or a 95% chance of of winning \$200 (Gamble B)?" According to expected utility (EU) theory (von Neumann & Morgenstern, 1944), people should evaluate all possible outcomes that each available action might have and weight them by their respective probabilities. Empirical research, however, has demonstrated that human decision making systematically deviates from EU theory (e.g., Kahneman & Tversky, 1979). These deviations are commonly interpreted as an indication of human irrationality. Recent work, however, suggests that they could also reflect people's rational use of limited cognitive resources (Lieder & Griffiths, 2019; Griffiths, Lieder, & Goodman, 2015).

To date, the most prominent descriptive theory of risky choice is cumulative prospect theory (CPT; Tversky & Kahneman, 1992). CPT accounts for many violations of EU theory by postulating that people's decision mechanisms systematically distort the stated probabilities (i.e., overweighting rare and underweighting common events) and payoffs (diminishing sensitivity to additional increases in the outcome as the outcome gets larger, and an amplification of losses relative to gains). Interpreted as a cognitive process model, CPT predicts that information is processed exclusively within each option and that the information processing is identical across all problems. Process-tracing studies, however, show that people often compare options along individual attributes and that the processing varies across problems. These process data are instead consistent with processing policies of simple heuristics (Payne & Braunstein, 1978; Pachur, Hertwig, Gigerenzer, & Brandstätter, 2013) such as the lexicographic heuristic, that usually only looks at each gamble's most probable outcome while ignoring all other possible outcomes. Yet, model comparisons have found that assuming that people use a single heuristic across all problems, no single heuristic predicts risky choices nearly as well as CPT (Glöckner & Pachur, 2012).

One way to resolve this apparent paradox is to postulate that people are equipped with a toolbox of several, often heuristic, strategies and that they use different strategies on different trials. This raises the question of which strategies their toolbox is equipped with and how people select between them. Previous work on strategy selection has found that people adapt their strategy use to the structure of individual choice problem and the situation's requirements for speed versus accuracy (Payne, Bettman, & Johnson, 1988). The rational strategy-selection model by Lieder and Griffiths (2017) captures this adaptive flexibility as well as the variability and the learning-induced changes in people's strategy selection. It does not, however, specify the set of strategies from which people select among. To address this question, Scheibehenne, Rieskamp, and Wagenmakers (2013) developed a hierarchical Bayesian measurement model for inferring the contents of the cognitive toolbox. This model, however, assumes that people's tendency to select a given strategy is not systematically related to the choice problem at hand and the requirements of the current situation. Strategy selection, however, has been shown to be sensitive to problem-specific features (Payne et al., 1988).

Here we develop an integrative model of risky choice with a *personal adaptive toolbox*. Our approach combines inferring the content of a person's cognitive toolbox with a rational model of strategy selection (Lieder & Griffiths, 2017). We validate this approach using a large empirical data set of risky choice data collected by Glöckner and Pachur (2012), testing it against single strategies, non-adaptive toolbox models, and CPT. Our model constitutes the first complete formalization of the notion that strategies are selected from a personal adaptive toolbox. It thereby enables more accurate inferences on people's cognitive toolbox than was previously possible, and we find that it predicts people's choices better than single strategies as well as other existing toolbox models.

The outline of this paper is as follows: We start by describing 11 extant (heuristic) strategies for risky choice, which might be contained in people's toolbox of decision strategies. We then introduce our computational model of the adaptive toolbox theory as well as several competitors. Next, we present a cross-validation method for inferring the set of strategies considered by an individual decision maker. We then evaluate our adaptive toolbox model against single strategies, non-adaptive toolbox models, and CPT. Finally, we apply our model to estimate the content of people's toolboxes—thereby elucidating why some people make better decisions than others. In closing, we discus the implications of our findings for the debate on human rationality as well as directions for future work.

Heuristics as Models of Risky Choice

A number of different strategies have been proposed as models of how people make decisions under risk. Following Glöckner and Pachur (2012), we consider the following ten heuristic strategies: the priority heuristic (PH), better-than-average (BTA), tallying (TALLY), probable (PROB), minimax (MINI), maximax (MAXI), lexicographic (LEX), equal-weight (EQW), least-likely (LL), and most-likely (ML). These heuristics cover a wide range of processing assumptions that differ in important aspects, such as whether they focus exclusively on the payoffs (BTA, TALLY, EQW, MINI, MAXI) or process both outcomes and probabilities (PH, PROB, LEX, LL, ML). For example, the minimax heuristic chooses the gamble with the highest minimum outcome and the least-likely heuristic identifies each gamble's worst outcome and then chooses the gamble with the lowest probability of the worst outcome.[1] Additionally, we include the weighted-additive strategy (WADD), which chooses the gamble with the highest expected payoff. Each of these strategies breaks ties between gambles by choosing randomly. We consider eleven simple models of risky choice according to which all people use one single strategy (either PH, BTA, TALLY, PROB, MINI, MAXI, LEX, EQW, LL, ML, or WADD) to make all their risky choices. Relaxing the assumption that all decision makers use the same strategy, we also tested a more flexible model (BEST), according to which each person might use a different strategy. That is, the BEST model has one parameter per person that encodes their strategy and has to be fitted to their choices.

Toolbox Models of Decision Making

According to the notion of an adaptive toolbox, each person is equipped with multiple strategies and employs them adaptively. In this section, we present three types of toolbox models that differ in whether the contents of the toolbox are inferred or assumed to be known and in their assumptions about how strategies are selected.

Strategy Selection Based on a Rational Cost-Benefit Analysis (RCBA)

Simulation studies by Payne et al. (1988) have shown that adaptively choosing between simple strategies can allow people to make many good decisions even when only little time is available. Assuming that decision makers are aware of the relevant properties of the choice problem (e.g., the magnitude of the possible outcomes), contextual factors (e.g., time pressure), and the speed and accuracy characteristics of the strategies in their toolbox, the adaptive decision maker (Payne et al., 1988) should choose strategies according to a rational cost-benefit analysis.

Building on the theory of rational metareasoning (Russell et al., 1991), the rational cost-benefit analysis (RCBA) model assumes that the expected payoff of making decision i using strategy h is integrated with the expected cost of the time $T(h,i)$ that it would take to do so. Together, they yield an estimate of the Value of Computation (VOC), defined as

$$\mathrm{VOC}(h,i) = \mathbb{E}\left[R(i,h(i))\right] - \delta \cdot T(h,i), \qquad (1)$$

where $R(i,h(i))$ is the payoff of decision $h(i)$ that strategy h would make in situation i, and $T(h,i)$ is the time it takes strategy h to make that decision. The balance between these two factors is determined by the relative opportunity cost δ. To model how long it takes to execute each strategy (i.e., the cost), we decompose the strategy into elementary information processes (EIPs) as introduced by Johnson and Payne (1985). Specifically, when a strategy is used to make a decision in a given choice problem, the number of EIPs required is recorded as $T(h,i)$. The RCBA model has two free parameters that can be estimated to accommodate individual differences: the set of available strategies H in the toolbox and the relative opportunity cost δ. For a given choice problem the strategy with the highest VOC in the toolbox is selected to make the choice.

[1]The equiprobable heuristic was not considered as it makes the same choice predictions as the equal-weight heuristic

Rational Strategy Selection Learning (RSSL)

The assumption of a full cost-benefit analysis for each strategy, as assumed by the RCBA, may be unfeasible for a boundedly rational mind. However, it might be possible to approximate the VOC. As one possible approach to do such an approximation, the rational strategy selection learning (RSSL) model assumes that the mind learns to predict each strategy's VOC based on the features of the choice problem at hand (Lieder & Griffiths, 2017). Specifically, the RSSL model assumes that people predict both the expected payoff and the expected time cost for each strategy (which are important for then determining the strategy's VOC) at a given problem based on a weighted sum of the features of the choice problem, such as the maximum probability or the range of outcomes; the weights for the estimation, in turn, are learned from the payoffs and decision times of past choices (with the latter is determined based on the number of EIPs the chosen strategy performed). The learning process is simulated using Bayesian linear regression and stochastic predictions are made by sampling from the posterior distribution.

The free parameters of the RSSL model are the number of samples drawn to predict the performance of each strategy, ζ, the set of strategies H, the opportunity cost δ and the amount of prior experience Λ (i.e., on how many choice problems the predictive models were trained on). For the latter parameter we assume that participants are equipped with some amount of prior experience in making choices using their strategies; hence we let the RSSL model learn from Λ randomly generated pairs of gambles prior to applying it to our participants' choices.

Toolbox Models Without Adaptive Strategy Selection

To assess how the assumption of rational strategy selection contributes to the predictive accuracy of the adaptive toolbox models introduced above, we evaluate them against simpler toolbox models that chooses strategies randomly for a given choice problem (rather than adaptively based on characteristics of the problems). In our first null model (NULL-TB1), every time a decision is made a strategy is selected from the set of 11 strategies introduced above. Our second null model (NULL-TB2) is like the first one except that the set of strategies it selects from is estimated on a participant-by-participant basis. Our third null model (NULL-TB3) extends the second one by allowing some strategies to be chosen more frequently than others. Specifically, following Scheibehenne et al. (2013), each strategy h is selected with probability θ_h, which is estimated from the participant's choices.

Cumulative Prospect Theory

According to CPT, the outcomes x_i of a gamble are transformed into subjective values according to the value function

$$v(x_i) = x_i^{\alpha} \text{ if } x_i \geq 0 \tag{2}$$

$$v(x_i) = -\lambda \cdot x_i^{\alpha} \text{ if } x_i < 0, \tag{3}$$

with an outcome sensitivity parameter $\alpha \in [0,2]$ that modulates the curvature of the value function and captures that people's sensitivity to changes in a payoff depend on its magnitude. Values of $\alpha < 1$ entails a concave value function with diminishing sensitivity to larger outcomes.

The probabilities p of the cumulative probability distribution function are transformed according to the probability weighting function

$$w(p) = \frac{p^{\gamma}}{(p^{\gamma} + [1-p^{\gamma}])^{1/\gamma}}, \tag{4}$$

whose shape is determined by the parameter $\gamma \in [0,2]$, which is defined separately for gains and losses. The shape of the probability weighting function reflects the degree of nonlinear distortion when the probabilities are mapped onto decision weights. Values of $\gamma < 1$ entail an inverse S-shaped probability weighting function, indicating a reduced sensitivity to probabilities in the middle range and a relative amplification of the sensitivity to differences among extreme probabilities. The overall valuation of a gamble is determined by multiplying each of the subjective values of the gamble's outcomes x_i by a decision weight π_i that follows from the weighted cumulative probabilities of obtaining an outcome at least as good as x_i if the outcome is positive, and at least as bad as x_i if the outcome is negative (for details see Tversky & Kahneman, 1992), and then summing the products:

$$V = \sum_i \pi_i \cdot v(x_i). \tag{5}$$

To derive the probability that gamble A is chosen over gamble B we apply the softmax choice rule to the gambles' subjective values V; this choice rule which has a choice sensitivity parameter ϕ (for details see Glöckner & Pachur, 2012).

Next, we describe the data set and our approach to evaluate the models introduced in the previous sections.

Data

We evaluated our models using data collected by Glöckner and Pachur (2012), who presented 64 participants with a set of 276 two-outcome gamble problems. The payoffs of the gambles ranged from -1000 to 1200 and the set of gambles consisted of pure gain (all payoffs> 0), pure loss (all payoffs < 0), and mixed (both positive and negative payoffs) gambles. The presentation of the gamble problems was distributed over two sessions that were one week apart (i.e., there are 138 choices from each session). For more information, see Glöckner and Pachur (2012).

Model Evaluation

We evaluated the predictive accuracy of each of the models using a simplified cross-validation method (Friedman, Hastie, & Tibshirani, 2001). Specifically, for each model a score was calculated indicating how often it correctly predicted the participants' choices on a held-out test set, that was not used to fit the model's parameters. The predictive accuracy for a given participant was computed by averaging the model's performance in forward prediction (i.e., fitting the model on choice data from Session 1 (t_1) and testing it on data from Session 2 (t_2)) and backward prediction (i.e., fitting choices from t_2 and testing on data from t_1). To perform forward-prediction and backward-prediction, the data set was split into three subsets: a *training set*, a *validation set*, and a *test set*. The training set was used to fit the parameters (e.g., the subjective time cost δ) of a given sub-model (e.g., a strategy selection model with a particular set of strategies). The validation set was used to select among sub-models based on unbiased estimates of their generalization errors (e.g., to select the model's toolbox). The test set was used to obtain an unbiased estimate of the selected sub-model's generalization error that could be compared to the performance of the other models.

Model Fitting and Prediction

Given a set of choice problems and the corresponding choices made by an individual, we fitted each model's parameters by maximizing the proportion of gambles from the training set for which the model's predicted choice agreed with the participant's choice. The model parameters were estimated using participants' choices from t_1 and then used to predict choices from t_2—and vice versa. For forward-prediction, we used the 138 gamble problems and choices from t_1 (training set) and split the gamble problems and choices from t_2 into a validation set comprising 103 problems and a test set comprising 35 problems. Backward prediction was performed in the same way as forward prediction except with t_1 and t_2 reversed.

BEST model For the BEST model, according to which each participant uses a single strategy across all choice problems, we determined for each participant the strategy that achieved the highest accuracy (in terms of overlapping choices) on the training set choices and the validation set.

RCBA We estimated each participant's set of strategies H along with their subjective time cost δ using the following procedure: In the first step, H included only the strategy h_1 with the highest accuracy on the validation set. Next, we determined which strategy h_2, if added, would result in the set of two strategies with the highest predictive accuracy on the validation set. In doing so, we estimate δ by optimizing the accuracy of each candidate sub-model on the training set using Bayesian adaptive directed search (BADS) (Acerbi & Ma, 2017). We then proceeded to evaluate toolboxes that added a third strategy to the toolbox and re-estimated δ until toolboxes containing up to 11 strategies had been evaluated. That is, we estimated a set H_k of k strategies for each $1 \leq k \leq 11$ and estimated each participant's toolbox by the set $H_{k_{\max}}$ for which our model achieved the highest predictive accuracy on the validation set.

RSSL As described above, to define a toolbox of strategies, the RSSL model estimates each strategy's VOC based on previous experience with gamble problems. To simulate this experience, we first randomly generated pairs of two-outcome gambles; their payoffs and probabilities were samples from the uniform distributions $\mathrm{Unif}([-1000, 1200])$ and $\mathrm{Unif}([0,1])$ respectively. The amount of prior experience (Λ) was set to 20000 gamble problems. Each choice problem was represented by a feature vector comprising the maximum probability of each gamble, the payoffs associated with the maximum probability (i.e., the most likely outcome), the ranges of payoffs within each gamble, and the range of payoffs across both gambles. These features were then used to predict the strategy's accuracy and effort for the problem at hand. The number of predictions ζ sampled from the posterior was set to 3. The parameters H and δ for the rational cost-benefit analyses model (which the RSSL shares with the RCBA model) were estimated following the same iterative procedure as described for the RCBA model.

Null models The Null-TB1 model has no free parameters. For the models NULL-TB2 and NULL-TB3 we estimated the set of strategies H using the same procedure as for the RCBA model. For NULL-TB3, we estimated the proportion parameters $\theta_1 ... \theta_{|H|}$ for a toolbox H by solving the constrained optimization problem to maximize the expected accuracy of participants' choices.

Cumulative prospect theory CPT's parameters were fitted to minimize G^2 based on the observed choices in the respective session. To reduce the risk of being stuck in local minima, we first conducted a grid search to identify the 20 best-fitting combinations of parameter values; these combinations were then used as starting values for subsequent optimization using the simplex method. For prediction, we derived deterministic choice predictions from CPT.

Results

Predictive Accuracy

Figure 1 shows how accurately single strategies, simple toolbox models, adaptive toolbox models, and CPT predicted the risky choices in the test set. Out of the eleven single strategies, WADD and minimax predicted people's choices best, with 65.8% and 61.3% accuracy, respectively. Relaxing the assumption that all participants use the same strategy and instead inferring a potentially different strategy for each participant (i.e., the BEST model) increased predictive accuracy to 69.4%, which is significantly higher than that of the best-performing single strategy WADD ($t(34) = 2.90$, $p = .004$). The simplest toolbox model Null-TB, which chooses randomly among all the strategies, was less predictive of people's choices than the BEST model (57.5% vs. 69.4%, $t(34) = -8.79$, $p < .001$). Its predictive accuracy increased, however, when we allowed the content of the toolbox to be estimated for each participant separately (65.7% vs. 57.5%, $t(34) = 6.66$, $p < .001$). Additionally estimating the relative frequency with which each strategy is selected

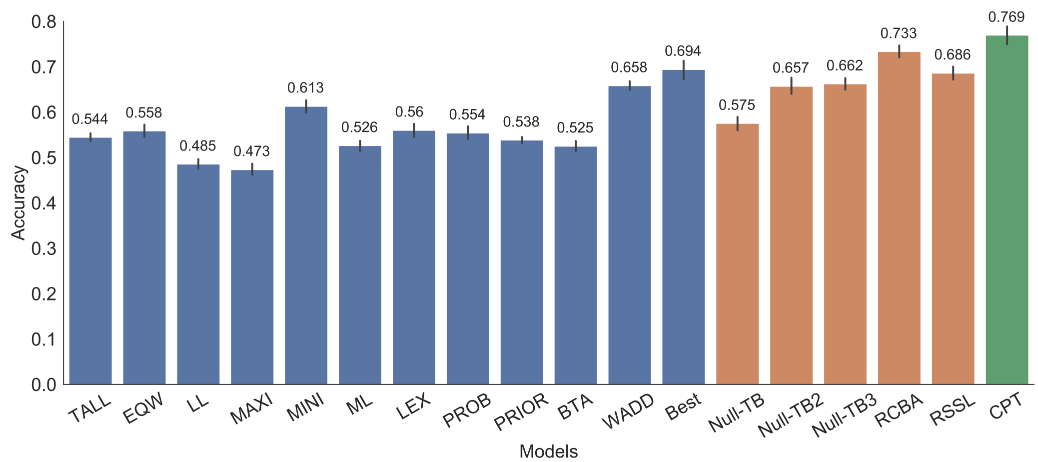

Figure 1: Comparison of how accurately each model predicted people's choices in the validation set. The single-strategy models are shown in blue, toolbox models in orange, and CPT in green.

independently of the problem (Null-TB3) further improved the toolbox model's predictive accuracy to 66.2% vs. 65.7% ($t(34) = 0.46$, $p = .64$). While the benefit of this choice problem-unspecific strategy-selection mechanism was rather small; adding an adaptive, problems-specific selection mechanism to the RCBA model drastically improved the accuracy of the toolbox approach to 73.3% vs. 65.7% ($t(34) = 6.50$, $p < .001$). The RSSL model, which approximates the rational cost-benefit analysis of strategy selection using the features of the choice problems as predictive cues, did not perform as well as the RCBA model (68.6%).

Critically, the RCBA model predicted people's choices better than the best-performing single strategy WADD ($t(34) = 8.80$, $p < .001$) and the BEST model ($t(69) = 3.02$, $p = .003$). This suggests that decision makers indeed adaptively choose from a personal toolbox of strategies when solving a sequence of different choice problems.

The RCBA model also achieved higher predictive accuracy than all null models, NULL-TB ($t(34) = 15.51$, $p < .001$), NULL-TB2 ($t(34) = 6.51$, $p < .001$) and NULL-TB3 ($t(34) = 7.56$, $p < .001$). These results corroborate the usefulness of combining inference about the content of the toolboxes with a model of how people's strategy choices are informed by the specific requirements of each individual decision. This finding strongly supports adaptive toolbox theories of human decision-making (Gigerenzer & Selten, 2002) in general and the idea of an adaptive *personal* toolbox in particular. Despite the substantial improvement in predictive accuracy we achieved by combining inference on the toolbox with adaptive strategy selection, the resulting RCBA model predicted people's choices not as well as CPT (73.3% vs. 76.9%, $t(34) = 3.03$, $p = .003$). While the RCBA model may thus not capture all aspects of how people make decisions, it being a process model still affords many practical advantages for understanding people's choices that cannot be obtained by modeling the choices with CPT (but see Pachur, Suter, & Hertwig, 2017). For example, the estimated contents of the toolbox and estimated parameters of the strategy selection mechanism provide a window onto the cognitive mechanisms underlying risky choice and how they vary across individuals.

Comparing Predicted and Actual Performance

Next, we compared the models and people in terms of their performance of their risky choices. Performance here is measured as the average expected value (EV) of the chosen gambles. WADD achieved an EV of 149.02, which therefore represents the upper bound on how well one could perform in this task. The RCBA model predicted a higher performance than what was actually observed for people's choices (143.41 vs. 130.83, $t(69) = 5.57$, $p < .001$). CPT, on the other hand, predicted a lower performance than people actually achieved (113.1 vs. 130.8, $t(69) = 5.68$, $p < .001$). The performance of the RSSL model and the toolbox model Null-TB3 fell in between, with 124.83 EV and 126.81 EV, respectively, and were closer to people's actual performance. These findings suggest that while people may not choose strategies optimally, they may still be substantially more resource-rational than CPT would make us believe.

Which Strategies Are In The Adaptive Toolbox?

Given our finding that the best-suited model to predict people's choices is the RCBA model, we next analyze its estimated parameters H and δ. Figure 3 shows how many strategies were in the estimated toolboxes of all participants. 28.91% of all toolboxes included 4 strategies, and 60.15% of all toolboxes included between 3 and 5 strategies. The average toolbox size was 4.3.

Next, we counted how often each of the eleven strategies was included in the estimated toolboxes (see Figure 2). Interestingly, with 79.68% and 71.09% WADD and minimax are the most frequently included strategies in the toolboxes. These two strategies also predicted people's choices most accurately (see Figure 1).

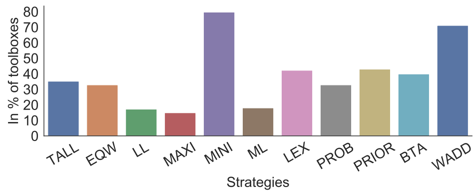

Figure 2: Percentage of cases each strategy was included in the toolboxes estimated by the RCBA model.

Minimax is especially useful as a risk-minimizing strategy when the probabilities of the possible outcomes are similar. The high inclusion rate of WADD suggests that at least when there are only two choices with only two possible outcomes, maximizing expected value is a viable and cognitively feasible strategy.

Furthermore, our results suggest that individual differences in decision quality might be due to the fact that different people are equipped with different toolboxes. For example, participants whose inferred toolbox included WADD performed better (144.59 EV) than participants whose inferred toolboxes did not include WADD (140.52 EV). Conversely, participants whose toolboxes were estimated to contain minimax achieved a lower performance than participants who did not use minimax (140.99 EV vs. 152.91 EV). These observations suggest that inferences obtained with the RCBA model can shed light on why and how people make the choices that they make. Additionally, our analysis identified another source of individual differences in decision performance: people's subjective cost of their time and effort. Specifically, our parameter estimates revealed a negative rank correlation between performance (in terms of EV) and the subjective opportunity cost δ (Spearman's $\rho(62) = -0.58$, $p < .001$), reflecting that higher opportunity costs favour less resource-intensive strategies even when they lead to less accurate decisions.

Finally, we found that the estimated size and content of the toolbox and the objective opportunity cost together explained 27% of the variance in individual differences in performance ($R^2 = 0.27$, $F(21, 106) = 2.25$, $p < .001$).

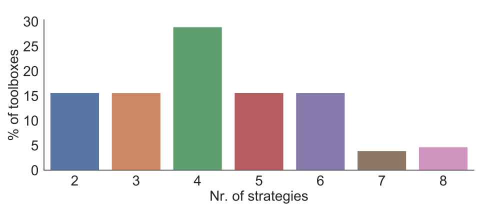

Figure 3: Toolbox sizes estimated by the RCBA model.

Discussion

We presented a model that represents the first complete formalization of the adaptive toolbox metaphor of human judgment and decision making (Gigerenzer & Selten, 2002). Our personal adaptive toolbox model predicted people's risky choices better than single strategies, non-adaptive toolbox models, or adaptive toolbox models that assume that all decision makers have the same strategies in their toolbox. Furthermore, the mechanistic nature of our model makes it possible to draw inferences about the cognitive architecture and processes underlying people's decisions. Furthermore, unlike CPT, our rational model of strategy selection can be applied to a wider range of domains, including inferential problems, such as those used by Gigerenzer and Goldstein (1996), by adapting the set of strategies (which can be deterministic or stochastic) and the reward function.

The success of the model that chooses strategies according to a rational cost-benefit analysis provides additional support for the view that people make rational use of their limited cognitive resources (Griffiths et al., 2015; Lieder & Griffiths, 2019). Our model is an important step towards reverse-engineering the mechanisms underlying the adaptive flexibility of human decision-making and individual differences in risky choice. But the mechanisms by which people efficiently approximate its rational cost-benefit analysis and the resulting suboptimalities need be investigated further before any definite conclusions can be drawn.

Future work will revisit the comparison with CPT using more complex decision problems, including problems with many alternatives and many possible payoffs (Payne et al., 1988), where people's selective processing of only a small subset of the available information might have a notable impact on their choices. We will also compare our models to other psychologically plausible models of risky choice including the utility-weighted sampling model (Lieder, Griffiths, & Hsu, 2018) and decision-field theory (Busemeyer & Townsend, 1993; Rieskamp, 2008; Bhatia, 2014) and apply likelihood-based model selection methods.

Future work will refine the strategy selection learning model with more realistic assumptions about decision makers' prior experience and the features they use to predict the performance of their strategies. In particular, future refinements of this model might take into account that people's strategy choices are informed by them learning from how well each strategy worked when they previously used it in the real world. This prior experience could be simulated by training the RSSL model on choice problems that are more like those that people encounter in everyday life (e.g., in having more possible outcomes and larger differences between the alternatives' expected values). The eleven strategies considered here are unlikely to cover all the decision mechanisms people use. Hence, we will consider additional strategies derived from resource-rational analysis (Lieder & Griffiths, 2019; Lieder, Krueger, & Griffiths, 2017; Gul, Krueger, Callaway, Griffiths, & Lieder, 2018) and process tracing.

References

Acerbi, L., & Ma, W. (2017). Practical Bayesian Optimization for model fitting with Bayesian adaptive direct search. In I. Guyon et al. (Eds.), *Advances in neural information processing systems 30* (pp. 1836–1846). Curran Associates, Inc.

Bhatia, S. (2014). Sequential sampling and paradoxes of risky choice. *Psychonomic Bulletin & Review*, *21*(5), 1095–1111.

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological review*, *100*(3), 432.

Friedman, J., Hastie, T., & Tibshirani, R. (2001). *The elements of statistical learning* (Vol. 1) (No. 10). Springer series in statistics New York, NY, USA:.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: models of bounded rationality. *Psychological review*, *103*(4), 650.

Gigerenzer, G., & Selten, R. (2002). *Bounded rationality: The adaptive toolbox*. Cambridge, MA: MIT Press.

Glöckner, A., & Pachur, T. (2012). Cognitive models of risky choice: Parameter stability and predictive accuracy of prospect theory. *Cognition*, *123*(1), 21–32.

Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, *7*(2), 217–229. doi: 10.1111/tops.12142

Gul, S., Krueger, P. M., Callaway, F., Griffiths, T. L., & Lieder, F. (2018, September). Discovering rational heuristics for risky choice. In *The 14th biannual conference of the German Society for Cognitive Science, GK.*

Johnson, E. J., & Payne, J. W. (1985). Effort and accuracy in choice. *Management Science*, *31*(4), 395–414.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291. doi: 10.2307/1914185

Lieder, F., & Griffiths, T. L. (2017). Strategy selection as rational metareasoning. *Psychological Review*, *124*(6), 762-794.

Lieder, F., & Griffiths, T. L. (2019). Resource-rational analysis: Understanding human cognition as the optimal use of limited ressources. *Behavioral and Brain Sciences*.

Lieder, F., Griffiths, T. L., & Hsu, M. (2018). Overrepresentation of extreme events in decision making reflects rational use of cognitive resources. *Psychological review*, *125*(1), 1.

Lieder, F., Krueger, P. M., & Griffiths, T. L. (2017). An automatic method for discovering rational heuristics for risky choice. In *Proceedings of the 39th annual meeting of the cognitive science society.* Austin, TX: Cognitive Science Society.

Pachur, T., Hertwig, R., Gigerenzer, G., & Brandstätter, E. (2013). Testing process predictions of models of risky choice: A quantitative model comparison approach. *Frontiers in Psychology*, *4*, 646.

Pachur, T., Suter, R. S., & Hertwig, R. (2017). How the twain can meet: Prospect theory and models of heuristics in risky choice. *Cognitive Psychology*, *93*, 44–73.

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

Payne, J. W., & Braunstein, M. L. (1978). Risky choice: An examination of information acquisition behavior. *Memory & Cognition*, *6*(5), 554–561.

Rieskamp, J. (2008). The probabilistic nature of preferential choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(6), 1446.

Russell, S., Wefald, E., Karnaugh, M., Karp, R., Mcallester, D., Subramanian, D., & Wellman, M. (1991). Principles of metareasoning. In *Artificial intelligence* (pp. 400–411). Morgan Kaufmann.

Scheibehenne, B., Rieskamp, J., & Wagenmakers, E.-J. (2013). Testing adaptive toolbox models: A Bayesian hierarchical approach. *Psychological Review*, *120*(1), 39-64.

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

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