Comparing Theories that Posit a Role for Task Features in Strategy Selection

Xinyu Xie (xx86@msstate.edu)
Jarrod Moss (jarrod.moss@msstate.edu)
Department of Psychology, Mississippi State University
Mississippi State, MS 39762 USA

Abstract

Salient features of a task play an important role in how people create task representations which then influence strategy selection for accomplishing the task. We examined two theories, Represent-Construct-Choose-Learn (RCCL) and Rational Metareasoning (RM), both of which incorporate task features into their models of strategy selection. RCCL theory posits that when a strategy's success rate is low, it indicates that the task representation is not useful and those represented features are irrelevant in this case so people tend to drop these features from the task representation. Conversely, RM theory posits that strategy selection is based on consideration of all available features, with no discrete changes in the features incorporated into the task representation. A study was conducted to examine how participants changed their strategy choices based on the success rate of using a specific task feature. The results showed that neither theory aligned closely with empirical data.

Keywords: feature; task representation; strategy selection

A strategy is a set of steps to solve a problem. Making choices about which strategy will succeed is not simple. Take driving as an example, there can be several routes from the starting point to the destination. How many traffic lights and stop signs are there on each route? What is the distance of each route? Will some routes include highways? All these factors can affect the choice of a certain route. How do we determine the optimal route if the goal is to reach the destination as quickly as possible? In other words, how do we choose the best strategy based on the information we have?

Some strategy selection theories focus on analysis of the costs and benefits of a strategy (Beach & Mitchell, 1978; Christensen-Szalanski, 1978; Payne et al., 1988, 1993). Some propose learning-based accounts of strategy selection (Erev & Barron, 2005; Rieskamp & Otto, 2006; Shrager & Siegler, 1998). However, these theories emphasize learning from feedback or association and focus on which strategy performs better on average overall. They neglect the influence of problem features used in the task representation. These problem features may serve as a basis for selecting strategies. The features used to represent a problem can influence strategy selection according to two existing theories: the Represent-Construct-Choose-Learn (RCCL; Lovett & Schunn, 1999) and Rational Metareasoning (RM; Lieder & Griffiths, 2017). Critically, the predictions of the two theories are different, which were tested in this study.

Represent-Construct-Choose-Learn Theory

The RCCL theory specifies how task representations can influence strategy choices and has four main stages: (1) Represent the task, (2) Construct a set of strategies based on features in the task representation, (3) Choose from among those strategies based on rates of success, and (4) Learn or update success rates based on performance (Lovett & Schunn, 1999).

The role of features in shaping task representations and their subsequent influence on strategy choices is central to the RCCL theory. A feature refers to a distinct characteristic or attribute of a task that can be used to represent and differentiate it from other tasks. Features can encompass various aspects, such as visual properties, object properties, spatial arrangements, or relational information. For example, when dealing with a puzzle-solving task, the shape, color, and position of puzzle pieces can all be considered features. A task representation refers to using a set of features to encode the task environment. Importantly, the salience of different features in a task impacts the initial task representation, by determining what features might be initially selected for the task representation.

In the next stage of the RCCL theory, the selected features of the task representation will be used to generate different strategies for use. Strategies are not merely a result of feature selection but are constructed utilizing the interplay of identified features, thereby tailoring problem-solving approaches to the task's unique attributes. For example, if someone selects the salient feature of position of puzzle pieces as input for initial task representation, someone may generate strategies such as finding all corner and edge pieces.

When a set of strategies are successfully formulated, individuals are tasked with choosing an optimal approach among them, guided by anticipated success rates of each strategy. In the example of puzzle-solving, the initially optimal strategy may be finding all pieces around corners under the task representation of position because the estimated success rate of this strategy is high. The reason is that it is likely to be easier to find all pieces around corners since those pieces have two straight edges that can be used to orient the pieces. So far, the success rate of find-corner-piece is relatively high, but the number of those pieces is limited. After utilizing this strategy to find all corner pieces, the success rate becomes low (near zero). So, the next step is to find a strategy that can continue solving the task. Under the given task representation, the alternative can be finding all pieces around the edges. However, due to the limited number of pieces, the success rate would eventually also become low, which means that individuals would seek for other strategies. Each strategy's success rate is learned with experience in the task. This learning mechanism leads to gradual changes of estimated success, and these changes in turn lead to strategy selection changes.

In L. K. Samuelson, S. L. Frank, M. Toneva, A. Mackey, & E. Hazeltine (Eds.), *Proceedings of the 46th Annual Conference of the Cognitive Science Society.*

Here, it is pivotal to note that success rates are intricately linked to the features embedded in the task representation. This connection between features and strategies' success rates underscores the theory's core principle. Central to the RCCL theory is the notion that success rates are only learned for strategies which are in turn based on the features incorporated into the task representation at the time the strategies are created. But what if the success rates of all strategies under the current task representation are low?

In instances where strategies yield suboptimal outcomes, the theory predicts a representation change. Low success rates with all available strategies prompt individuals to explore alternative features for their task representation. This adaptive process involves either adding or removing features from the task representation. For example, after finding all puzzle pieces around corner and edges, other strategies using position are relatively less successful, which drives individuals to seek other features to re-present the task. It can be the color of pieces in this example, where people can group pieces by different color. New strategies would be generated from this revised representation and their success rates learned by how well they perform. The key idea is that the representation of a task is not fixed but dynamic, with individuals actively modifying their feature selection based on their experiences and task performance. Critically, once a feature is dropped, it will not be used in the future task representation and strategy selection if successful strategies can be discovered that do not use that feature.

Rational Metareasoning Theory

A second theory that also uses features of the problem to guide strategy selection is the RM theory (Lieder & Griffiths, 2017). In the RM theory, there are a set of predefined strategies for a type of problem that the system must learn to select from based on the features of each problem. The process of strategy selection entails the estimation of each strategy's expected rewards (i.e., probability of successful problem resolution) and expected costs (i.e., execution time). The expected rewards and costs combine into an estimated value of computation (VOC) for each strategy with a person then opting for the strategy with the highest estimated value for a given problem (Lieder & Griffiths, 2017). The reward and cost of a particular strategy are estimated based on the entire set of problem features, and then these reward and cost estimates are combined into the VOC.

Sticking with the puzzle example, almost every puzzle has the same feature set including shape, color, and images. There could be a set of strategies based on the feature set. But the VOC of each strategy varies among the puzzle problems. If one puzzle has several distinctive colors in different locations, then sorting the pieces by color is likely to be most helpful (has the highest estimated VOC). But if in another puzzle which is about a picture of a forest, there is not a good mapping between color and location such that sorting the pieces by images or shape can be helpful instead of by color. In this manner, each problem of a specific type has a set of values for each feature that affects the estimated reward and cost of each strategy for that problem type. Strategy selection under RM is sensitive to how these rewards and costs change with different feature values.

In line with reinforcement learning principles (Sutton & Barto, 1998), the mapping from features to strategies is learned from experience via a set of weights between the features and the expected cost and reward. This learning mechanism can transfer the learned weights from prior problem solving to pick effective strategies for novel problems that have similar feature values. Thus, the RM theory posits that people select strategy based on the features of the problems. This model illustrates how people are capable of learning to predict each strategy's expected value with reward and cost from features of individual problems. In essence, the feature set of a problem is not changing, and only the weights between features and values are updated.

Comparison of RM and RCCL

The RCCL theory and the RM theory both emphasize the importance of task features for strategy selection. But they hold different views of how to use features to select strategies. In the RCCL theory, if the success rate of all strategies for a given task representation is too low, it indicates that the task representation is not useful, and the represented features are modified by adding or dropping features. So, people will tend to use other features in the task representation over time with some irrelevant features being dropped from the problem representation. However, the RM theory posits that strategy selection is based on all available features and those features will not change throughout the problem-solving process. In other words, the set of available features will change under the RCCL theory in some conditions possibly leading to a change in the set of strategies that can be considered, but under the RM theory the available features never change. In the case of RM, the weights between features could change, but the set of features are not changed in a discrete manner as in RCCL. Our experiment aims to examine whether the features of a problem impact strategy selection in a manner more consistent with the RM theory or the RCCL theory.

Current Study

The Building Sticks Task (BST) has been used in prior strategy selection work to examine support for some of the predictions of the RCCL theory (Lovett & Schunn, 1999). As shown in Figure 1, there are three building sticks (black) and a target stick (green) for one problem. The goal of the task is to add and subtract the lengths of the three sticks to match the target length. A participant's strategy can be categorized as either undershoot or overshoot based on the first click. The undershoot strategy starts with the longest stick that is shorter than the target (i.e., stick B) and then adds additional sticks to reach the target length. In contrast, the overshoot strategy starts with the longest stick (i.e., stick C) and then subtracts other sticks to reach the target length.

Figure 1: BST examples. Participants click the button next to the stick to add or subtract the stick. The two BST problems shown give an example of building a stick starting with B or C respectively. If the stick being built is shorter than the target stick, its color will be red (left). If the stick being built is longer than the target stick, its color will be blue (right). When the stick participants are building equals the target stick length, participants will be told that the problem is done, and the next problem is presented. While solving a problem, if participants want to restart the problem from the beginning, they can click the reset button.

In BST problems, there is usually one stick closest to the target length, which is a salient feature of the task for participants and they will select this stick in accordance with a hill-climbing heuristic (Lovett & Anderson, 1996). However, this relative length cue may or may not be the correct strategy for a given BST problem. Using the relative length cue, we designed problems with a strong relative length cue or neutral problems where there is no relative length cue. Figure 2 shows three types of problems: (a) a problem with a strong relative length cue to choose undershoot; (b) a neutral problem; (c) a problem with strong relative length cue to choose overshoot. By manipulating the rates of success of the overshoot and undershoot strategies in different types of problems over time, our study attempts to determine if unhelpful features are dropped from the problem representation as explained by RCCL or will continue to be used to select strategies as explained by RM.

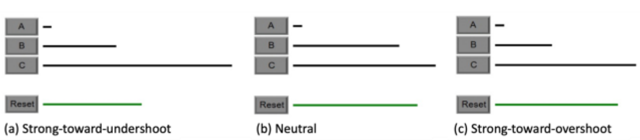

Figure 2: Examples of three types of problems. In the strong-toward-undershoot problem, stick B is much closer to the target stick than stick C. In the neutral problem, stick B and stick C are equally distant to the target stick. In the strong-toward-overshoot problem, stick C is much closer to the target stick than stick B.

Method

Design and Procedure

This study used a 2 (correct strategy during the experimental phase: undershoot vs. overshoot) * 4 (problem-group: cue-inconsistent, no-cue, cue-consistent, mixed) between-participant design. Participants were randomly assigned to each condition. The problem-group factor differed in the set of problems in the experimental phase (see a detailed description later).The correct strategy factor was primarily intended as a counterbalancing factor, and collapsing across the levels of this factor yielded the four primary conditions.

There were four phases presented in this order: relative-length-cue training, pretest, experimental, and posttest phases. All phases were identical for all participants except the experimental phase. In the relative-length-cue training phase, each participant was given 36 BST problems including 18 strong-toward-overshoot problems and 18 strong-toward-undershoot problems, where the relative length cue was 100% predictive of the correct strategy.

In the pre/posttest, six sets of three types of problems were used, 6 strong-toward-overshoot problems, 6 strong-toward-undershoot, and 6 neutral problems. All the problems in the pre- and posttest could be solved by either strategy, ensuring that there was no reason why one strategy would be selected over the other. Participants only selected the first stick in the pre- and posttest. The purpose of the pretest was to test if participants learned the relative length cue, while the posttest was compared to pretest to examine if the manipulations in the experimental phase affected strategy selection.

In the experimental phase, there were another 36 problems of the problem type for that problem-group condition. The number of BST problems for the relative-length-cue training phase and the experimental phase were determined in a pilot experiment. In these two phases, each BST problem could be solved only by one of the two strategies, and participants had to solve every problem to move on.

After consenting, participants began with instructions and practice trials to help participants understand how to use the undershoot strategy and the overshoot strategy to solve BST problems in the interface (refer to the interface description in Figure 1). Participants were guided through solving one BST problem for each strategy before having to practice solving one BST problem on their own before proceeding to the relative-length-cue training phase.

Participants

A sample of 259 participants were recruited from Prolific for compensation in exchange for participation. Since we had different conditions in this study, resulting in different hypothesized effect sizes, we estimated the sample size based on the smallest effect size that was anticipated to occur from pretest to posttest. Based on a linear mixed effects model examining the pretest to posttest change in proportion of one strategy being selected in a prior study using the BST, we used Monte Carlo simulation to conduct a power analysis showing that a sample size of 65 per problem-group condition would result in a power of .9.

Experimental Phase Conditions and Predictions

The manipulation in the experimental phase was to manipulate the type of problems participants experience after they had learned that the relative length cue was predictive of the correct strategy. Before the experimental phase and after the relative-length-cue training phase, participants' strategy choices on the pretest were expected to look like those in Figure 3 where they are relying on the relative length cue.

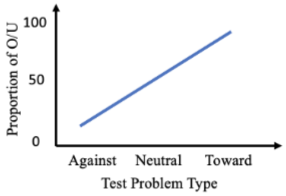

Figure 3: Expected strategy selection for pretest. For these figures, the dependent measure is the proportion of time participants selected the strategy (overshoot/undershoot) that was the strategy that solved all problems in the experimental phase. The assignment of overshoot/undershoot was counterbalanced across participants. There were three types of problems on the pretest: strong-toward-undershoot, neutral, and strong-toward-overshoot. These problem types were recoded based on the learned strategy as designated by the counterbalancing condition the participant was assigned to. For example, when assigned to overshoot learning in the experimental phase, strong-toward-undershoot problems were classified as against problems because the assigned successful strategy for that condition was overshoot but the relative length cue suggested the undershoot strategy would be successful. Thus, recoded problem types were against, neutral, and toward problems.

In the cue-inconsistent condition, participants experienced problems solved by one specific strategy (undershoot or overshoot counterbalanced across participants) but in which the correct strategy was the opposite of what the relative length cue indicated. For example, participants would only see problems for which the relative length cue strongly indicated overshoot, but the correct strategy was undershoot. In these problems, participants now had learned to use the relative length cue feature with a very high success rate and now this feature was no longer useful. In this condition, RCCL predicts that people would learn to ignore the relative length feature because the strategy that uses it is not successful. If that feature was dropped, then participants should not show sensitivity to the relative length cue during the posttest. RM predicts that people would use the strategy learned in recent trials only for similar problems because the relative length cue was not dropped from the representation. In this case, participants who saw strong-toward-overshoot problems but for which the correct strategy was undershoot would select undershoot on future problems that are strong-toward-overshoot. However, their preferences for problems with other features such as strong-toward-undershoot problems should not be altered by their experiences with strong-toward-overshoot problems in the experimental phase. This pattern of predictions can be seen in the first row of Figure 4.

In the no-cue condition, people only encountered neutral problems with no relative length cue during the experimental phase. RCCL's prediction was the same as before because the relative length cue strategy was not relevant. RM predicts that people would use the strategy learned in recent trials for neutral problems, and also use the original relative length cue feature for other problems not encountered in the experimental phase (second row of Figure 4).

In the mixed condition, participants experienced strong-toward-overshoot, neutral, and strong-toward-undershoot problems that are all solved by one strategy. For RCCL, the relative length cue again became unpredictive for most problems and was dropped from the representation. For RM theory, the prediction is that the slope is reduced as people experience problems with all values of the relative length cue being predictive of only one strategy being successful. But the weights between features and strategy success estimates should be a weighted combination of experiences such that the slope is not 0.

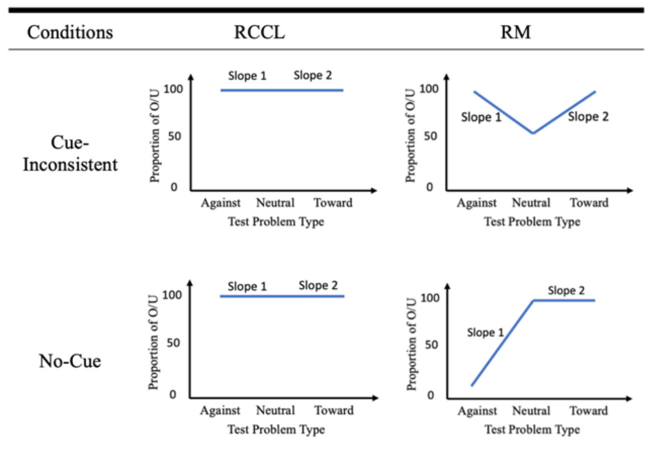

Figure 4: Illustration of predictions for posttest of the two conditions. The rows correspond to conditions. Then RCCL and RM's different predictions for how the pattern of strategy selection changes after the experimental phase are in the columns. Pretest data from all conditions are predicted to start out looking like the hypothetical pretest results shown in Figure 3. The figure also shows that the line can be decomposed into the slope from Against to Neutral (slope 1) and the slope from Neutral to Toward (slope 2). These slopes were compared in the analyses.

In the cue-consistent condition, participants experienced problems solved by the strategy predicted by the relative length cue (problems solved by overshoot when the relative length cue indicated overshoot). According to both theories, nothing had changed so pretest and posttest strategy selection would be the same. This condition served as a control that encountering a block of trials all solved by the same strategy did not result in participants only selecting the most recently successful strategy.

Analysis Approach

All analyses were conducted with generalized linear mixed effects (LME) models to analyze data. Random intercepts for both participants and problems were included. Random slopes for all within-participant and within-item

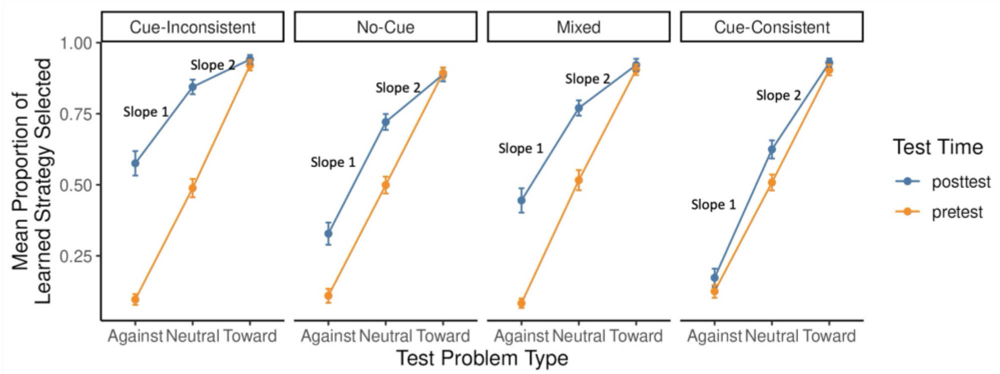

Figure 5: Proportion of choosing learned strategy in pretest and posttest under four conditions.

manipulations were also included. If the model did not converge or reported a singular fit, then the random effects structure was simplified by removing random slopes that accounted for little to no variance (Matuschek et al., 2017).

Results

First of all, we need to determine that if participants learned to use the relative length cue during the initial training phase. A logistic generalized linear mixed effects model was used to detect the effect of the relative length cue on the proportion of one certain strategy being selected. Since the strategy choice was categorized as either the overshoot strategy or the undershoot strategy based on the first click, the dependent measure chosen for this analysis was just the proportion of time that overshoot was selected. The relative length cue was coded from -1 = strong-toward-undershoot to 1 = strong-toward-overshoot. With the proportion of overshoot strategy as the dependent variable, the relative length cue was found to predict participants' strategy selection (b = 3.12, SE = .22, z = 14.25, p < .001). As can be seen by the pretest lines for each condition in Figure 5, the overshoot strategy proportion was nearly a linear function of the relative length cue, and the relative length cue had a positive effect on the proportion of one strategy being selected. The relative-length-cue training phase therefore served its purpose.

The main hypotheses concerned how the experimental phase altered strategy preference from pretest to posttest. We used a logistic linear mixed effects model to compare data, with the dependent measure as whether the the learned strategy was selected and the predictors of test problem type and test time. A separate model was used for each of the four conditions.

In order to better assess the extent to which our data aligned with the predictions of the RCCL and RM theories, we employed an approach focusing on the direction comparison of slopes (positive, zero, or negative). This approach enabled us to evaluate the change in strategy selection across different problem types. Specifically, we examined whether the changes in strategy selection from against problems to neutral problems (i.e., Slope 1 shown in Figures 4 and 5) and from neutral problems to toward problems (i.e., Slope 2) in the posttest align with the predictions. By focusing on the slope direction of these changes, we aimed to identify whether the data supported the theoretical predictions of both theories.

Figure 5 shows observed data for all conditions. It demonstrates that the posttest data did not exactly match either of the predictions displayed in Figure 4 for the cue-inconsistent condition. The results for the cue-inconsistent condition showed that there was a significant difference in strategy selection from pretest to posttest for against problems (b = -2.89, SE = .29, z = -9.9, p < .001) and neutral problems (b = -2.46, SE = .28, z = -8.68, p < .001), but no significant difference in strategy selection from pretest to posttest for toward problems (b = -1.06, SE = .38, z = -2.8, p = .06), and a significant difference in strategy selection for the posttest data between against problems and neutral problems (b = -1.91, SE = .21, z = -9.02, p < .001) and between against problems and toward problems (b = -3.21, SE = .28, z = -11.46, p < .001). By comparing the directions of slope 1 and slope 2 between Figure 5 and Figure 4, it shows that our data did not support any of the predictions in the cue-inconsistent condition though the direction of slope 2 was consistent with the RM prediction.

In the no-cue condition, Figure 5 shows that the posttest data did not follow either of the predictions as shown in Figure 4. The results showed that there was a significant difference in strategy selection from pretest to posttest for against problems (b = -1.35, SE = .22, z = -6.10, p < .001) and neutral problems (b = -1.11, SE = .18, z = -6.29, p < .001), no significant difference in strategy selection from pretest to posttest for toward problems (b = -0.14, SE = .26, z = -0.55, p = .99), and a significant difference in strategy selection for the posttest data between against problems and neutral problems (b = -1.96, SE = .18, z = -10.95, p < .001) and between neutral problems and toward problems (b = -1.25, SE = .21, z = -5.99, p < .001). By comparing the directions of slope 1 and slope 2 between Figure 5 and Figure 4, it shows that our data did not support any of the predictions in the no-cue condition though the direction of slope 1 was consistent with the RM prediction.

In the mixed condition, RCCL's predictions were still same as before. However, for RM's prediction, we expected to see a decrease of both slopes. If the slope changed significantly but not equaled to zero, it was consistent with RM theory.

As shown in Figure 5, there was a significant difference in strategy selection from pretest to posttest for against problems ($b = -2.35$, $SE = .22$, $z = -10.72$, $p < .001$) and neutral problems ($b = -1.28$, $SE = .17$, $z = -7.69$, $p < .001$), and no significant difference in strategy selection from pretest to posttest for toward problems ($b = -0.21$, $SE = .26$, $z = -0.81$, $p = .97$). Critically, there was a significant difference in strategy selection between against problems and neutral problems at posttest ($b = -1.60$, $SE = .17$, $z = -9.54$, $p < .001$) and between neutral problems and toward problems ($b = -1.37$, $SE = .23$, $z = -5.93$, $p < .001$). There was indeed a slope change from pretest to posttest, indicating that the data in the mixed condition supported RM theory.

In the cue-consistent condition, both theories predict there should be no significant differences in all problems from pretest to posttest. However, as shown in Figure 5, our results showed that there was a significant main effect of test time ($b = 0.43$, $SE = .13$, $z = 3.45$, $p < .001$). Pairwise comparisons revealed a significant difference in strategy selection from pretest to posttest for neutral problems, $b = -0.54$, $SE = .16$, $z = -3.46$, $p = .007$, contradicting the theoretical expectations though the directions of slope 1 and slope 2 were identical to both theories' predictions.

Discussion

To summarize, in the pretest we found that the relative length cue significantly predicted participants' strategy selection, confirming that participants learned to use this cue after training. However, when examining the change in strategy selection from pretest to posttest in different problem types under four conditions, our data did not entirely align with the predictions of either the RCCL or RM theories.

In the cue-consistent condition, the intention of this control condition was to investigate whether a block of trials all solved by the same strategy would lead participants to mechanically select that strategy. Our data indicated that there was not such a mental set (Ollinger et al., 2008). However, contrary to the predictions made by both theories that there are no differences from pretest to posttest, we observed a small but significant change in strategy selection from pretest to posttest. One possible reason is that the difference is caused by a change in base rates. The base rate of a specific strategy's success rate increases from the training phase (50%) to the experimental phase (100%). Consequently, the increase in the proportion of choosing the learned strategy from pretest to posttest could be driven by the base rate increase. This could be the reason why our data do not completely match the prediction though it accomplishes the goal of the control condition.

In the mixed condition, there was a significant slope change from pretest to posttest, providing support for the RM theory. As mentioned before, RM predicts that the slope is expected to be reduced as people experience all types of problems. The decreased slope in our data is the evidence that people consider both the relative length cue and the most successful strategy for recent trials.

In the no-cue condition, however, our data did not fully support any theory. The significant differences among three test problem types at posttest contradicted RCCL's prediction. One possibility is that the strategies utilizing the relative length cue may not be applicable in the experimental phase since there are all neutral problems with no relative length cue, so the estimated success rates of using relative length cue remain unchanged. In this case, the prediction of RCCL should be revised to be identical to the prediction of RM that only for neutral problems people would use the strategy learned in recent trials. However, the significant differences in strategy selection from pretest to posttest for the against problems contradicted this prediction again.

Similarly, in the cue-inconsistent condition, significant differences among three test problem types at posttest contradicted RCCL's prediction, and the posttest curve showed participants' sensitivity to the relative length cue decreased. It suggests that participants do not drop the feature and still incorporate the relative length cue in the task representation. However, we also found significant differences in strategy selection from pretest to posttest for neutral problems, which did not align with RM's prediction that strategy selection only differed for against problems. Moreover, the slope direction of changes in strategy selection from against problems to neutral problems was positive, opposite to RM's prediction that the slope would be negative. It raises the question of why the slope direction of the data differs from that predicted by RM. To address this question, we revisited the predictions. Notably, when reasoning about the predictions for each experimental condition according to the RM theory, we regarded the relative length cue pointing to the overshoot strategy as a separate feature from the feature indicting the length cue pointing to the undershoot strategy, which implies that there are two weights for the expected value of these features mapped to this strategy. Each time the predictability of a feature is changed, only one of these weights should theoretically have been changed in this case. However, if the relative length feature is viewed as continuous, then there will only be one weight. And the prediction will be that each time after experiencing a problem with the success or failure of the relative length cue, the weight will be updated and the feature value tied with overshoot/undershoot together will be changed, which is different from the earlier prediction that the values tied with overshoot and undershoot vary separately. Under the newer prediction, the strategy selection will change in a manner more closely aligned with the data. But the problem lies in the fact that there is no specific reference to how a feature is defined in the RM theory. It will be dependent on how the researcher thinks the features are represented.

In conclusion, our data showed evidence against the RCCL theory, but it may support the RM theory when different ways to define a feature are considered. Given the existing questions and complexities within our findings, further research is needed to reveal the dynamics of strategy selection in the context of the BST and other tasks.

Acknowledgments

This work was supported by the Office of Naval Research grant number N000142112617.

References

Beach, L. R., & Mitchell, T. R. (1978). A Contingency Model for the Selection of Decision Strategies. *The Academy of Management Review*, *3*(3), 439–449. https://doi.org/10.2307/257535

Christensen-Szalanski, J. J. J. (1978). Problem solving strategies: A selection mechanism, some implications, and some data. *Organizational Behavior and Human Performance*, *22*(2), 307–323. https://doi.org/10.1016/0030-5073(78)90019-3

Erev, I., & Barron, G. (2005). On Adaptation, Maximization, and Reinforcement Learning Among Cognitive Strategies. *Psychological Review*, *112*, 912–931. https://doi.org/10.1037/0033-295X.112.4.912

Lieder, F., & Griffiths, T. L. (2017). Strategy selection as rational metareasoning. *Psychological Review*, *124*(6), 762–794. https://doi.org/10.1037/rev0000075

Lovett, M. C., & Anderson, J. R. (1996). History of Success and Current Context in Problem Solving: Combined Influences on Operator Selection. *Cognitive Psychology*, *31*(2), 168–217. https://doi.org/10.1006/cogp.1996.0016

Lovett, M. C., & Schunn, C. D. (1999). Task representations, strategy variability, and base-rate neglect. *Journal of Experimental Psychology: General*, *128*(2), 107–130.

Matuschek, H., Kliegl, R., Vasishth, S., Baayen, H., & Bates, D. (2017). Balancing Type I error and power in linear mixed models. *Journal of Memory and Language*, *94*, 305–315. https://doi.org/10.1016/j.jml.2017.01.001

Ollinger, M., Jones, G., & Knoblich, G. (2008). Investigating the effect of mental set on insight problem solving. Experimental Psychology, 55(4), 269–282. https://doi.org/10.1027/1618-3169.55.4.269

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*, 534–552. https://doi.org/10.1037/0278-7393.14.3.534

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker* (pp. xiii, 330). Cambridge University Press. https://doi.org/10.1017/CBO9781139173933

Rieskamp, J., & Otto, P. E. (2006). SSL: A Theory of How People Learn to Select Strategies. *Journal of Experimental Psychology: General*, *135*, 207–236. https://doi.org/10.1037/0096-3445.135.2.207

Shrager, J., & Siegler, R. S. (1998). SCADS: A Model of Children's Strategy Choices and Strategy Discoveries. *Psychological Science*, *9*(5), 405–410. https://doi.org/10.1111/1467-9280.00076