

Taking the easy way out? Increasing implementation effort reduces probability maximizing under cognitive load

Christin Schulze^{1,2} · Ben R. Newell¹

Published online: 16 February 2016
© Psychonomic Society, Inc. 2016

Abstract Cognitive load has previously been found to have a positive effect on strategy selection in repeated risky choice. Specifically, whereas inferior probability *matching* often prevails under single-task conditions, optimal probability *maximizing* sometimes dominates when a concurrent task competes for cognitive resources. We examined the extent to which this seemingly beneficial effect of increased task demands hinges on the effort required to implement each of the choice strategies. Probability maximizing typically involves a simple repeated response to a single option, whereas probability matching requires choice proportions to be tracked carefully throughout a sequential choice task. Here, we flipped this pattern by introducing a manipulation that made the implementation of maximizing more taxing and, at the same time, allowed decision makers to probability match via a simple repeated response to a single option. The results from two experiments showed that increasing the implementation effort of probability maximizing resulted in decreased adoption rates of this strategy. This was the case both when decision makers simultaneously learned about the outcome probabilities and responded to a dual task (Exp. 1) and when these two aspects were procedurally separated in two distinct stages (Exp. 2). We conclude that the effort involved in implementing a choice strategy is a key factor in shaping repeated choice under uncertainty. Moreover, highlighting the importance of

implementation effort casts new light on the sometimes surprising and inconsistent effects of cognitive load that have previously been reported in the literature.

Keywords Decision making · Probability matching · Working memory load · Rational choice theory

Solving complex problems requires effort, time, and resources. Conventional wisdom suggests that the more difficult the problem at hand, the more effort must be invested to achieve success: Proverbially, Rome was not built in a day, and there is no gain without pain. Paradoxically, however, people's performance on simple choice tasks—which seem to be surprisingly challenging to begin with—can apparently be boosted by *increasing* the task demands.

Consider the simple task of choosing between two options that offer the same payoff with unequal odds (e.g., with $p = .70$ and $1 - p = .30$). For example, imagine you are to choose between two casino slot machines: Slot machine A will pay out \$10 with a probability of .70, and slot machine B will also pay out \$10, but with a probability of .30. Which slot machine would you prefer to play? It is easy to see that a rational decision maker should select the alternative with the higher payoff probability (slot machine A) to maximize her chances of success. The same holds when she faces this choice repeatedly, provided that the outcome probabilities remain stationary and are serially independent. Thus, in repeated risky choice, a simple reward-maximizing strategy is to always select the more probable option—that is, to *probability maximize*. Yet many people faced with this task apply a more complicated strategy that involves aligning their choice frequencies to the relative probabilities of the outcomes. In the slot machine example, this means switching repeatedly between the options and playing machine A 70 % of the time

✉ Christin Schulze
cschulze@mpib-berlin.mpg.de

¹ School of Psychology, University of New South Wales, Sydney, Australia

² Max Planck Institute for Human Development, Berlin, Germany

and machine B 30 % of the time. The expected payoffs of this strategy—called *probability matching*—are much lower (for a review, see Vulkan, 2000).

Surprisingly, under some conditions, people maximize *more* (and match less) when the overall task difficulty is increased, for example, by the introduction of a concurrent verbal memory task that competes for cognitive resources (Wolford, Newman, Miller, & Wig, 2004). Yet this finding is surprising only if one assumes that people really believe the structure of the sequential choice task to be simple and the outcome sequence to be random. If they do not believe that the outcomes are statistically independent—which seems a reasonable assumption, in light of everyday experiences of repeated events (see, e.g., Ayton & Fischer, 2004)—they might attempt to outperform the static maximizing strategy by finding a predictable pattern in the outcome sequence (Gaissmaier & Schooler, 2008; Peterson & Ulehla, 1965). Because any predictable pattern must match the outcome frequencies, probability matching would occur as a by-product of such an elaborate search, rather than as a strategy per se. By this account, probability matching represents an ecologically rational response associated with the search for patterns. Wolford et al. (2004) argued that the assumption that people search for patterns in outcome sequences is in line with the reduced probability matching rates observed when cognitive resources are taxed. That is, occupying the cognitive resources needed for vigilant pattern search would undercut such normally occurring search behavior, and thus reduce probability matching.

Other findings also support the pattern search account of probability matching. Gaissmaier and Schooler (2008) demonstrated that participants who probability matched in the absence of a fixed pattern in the outcome sequence were more likely to detect patterns introduced at a later stage. Unturbe and Corominas (2007) found that the complexity of the rules that participants reported to have followed during a sequential choice task was inversely related to probability maximizing behavior (see also McMahon & Scheel, 2010). Moreover, allowing participants to conclusively infer that the outcome-generating process is random—and thereby encouraging them to accept the true absence of patterns in the outcome sequence—reduces probability matching (Morse & Runquist, 1960; Peterson & Ulehla, 1965).

However, not all researchers agree with this pattern search interpretation of probability matching. In fact, probability matching has classically been viewed as a simple mistake that violates the assumptions of traditional (as opposed to ecological) rational choice theory (Vulkan, 2000). Recent research adopting this view has argued that people make this mistake because cognitive constraints motivate them to fall back on cognitively simpler heuristic choice strategies (e.g., Koehler & James, 2009, 2014; Kogler & Kühberger, 2007; West & Stanovich, 2003) or because limitations associated with the

task environment—for instance, lack of financial incentives or insufficient outcome feedback—prevent them from learning how best to respond (Newell & Rakow, 2007; Shanks, Tunney, & McCarthy, 2002). According to this account, probability matching is simply a cognitive error that arises from cognitive constraints of the decision maker and/or from structural inadequacies of the choice environment.

Findings from dual-task paradigms (e.g., Wolford et al., 2004) appear to be at odds with this “simple-mistake” interpretation of probability matching, however. If people match by mistake because their cognitive capacity is limited, why would they match *less* when cognitive resources are further taxed under dual-task conditions? One explanation for this apparent inconsistency relates to differences in the effort involved in implementing probability matching and maximizing (see Koehler & James, 2014). To implement a probability matching strategy, decision makers need to track their choice proportions throughout sequential choice tasks—an effort that can be assumed to involve cognitive mechanisms similar to those that would be required by a concurrent verbal memory task. Probability maximizing, on the other hand, involves only a simple repeated response to the same option and requires less effort to implement. Thus, a concurrent task might reduce the adoption of probability matching simply because it impedes people’s ability or willingness to track their own choices (rather than patterns in the outcome sequence), and consequently drives them to default to choosing the same option repeatedly (i.e., to probability maximize).

Thus, the observation of reduced probability matching under cognitive load does not necessarily discriminate between the “sophisticated pattern search” and “simple-mistake” accounts of probability matching. Moreover, the potentially pivotal role of strategy implementation effort has remained largely unexplored, which may explain the inconsistent findings on the effects of cognitive load in these settings. Specifically, some studies have demonstrated that probability matching decreases under cognitive load (Wolford et al., 2004), whereas others have failed to find differences in probability matching rates under dual- versus single-task conditions (Otto, Taylor, & Markman, 2011).

We aimed to close this gap in the literature by evaluating the role that strategy implementation effort plays in moderating the effects of cognitive load on sequential choice. To this end, we introduced a manipulation that reversed the typical experimental situation, by making the implementation of probability maximizing more taxing and, at the same time, allowing decision makers to probability match via a simple repeated response to a single option. Specifically, we manipulated the allocation of choice options to physical response options, so that repeatedly choosing the same physical option resulted in *either* probability matching *or* probability maximizing (see the Method section of Exp. 1 for details). By applying this manipulation in dual-task paradigms, we tested

whether cognitive load causes people to repeatedly choose the same physical option rather than the maximizing option. Our two experiments explored the significance of strategy implementation effort when decision makers learn about outcome probabilities and respond to a dual task simultaneously (Exp. 1) or in two procedurally separated stages (Exp. 2).

Experiment 1

Method

Participants One hundred (51 female, 49 male) undergraduate students from the University of New South Wales with a mean age of 19.63 years ($SD = 2.56$ years) participated in this experiment in exchange for course credit. These figures exclude one additional participant who strongly favored the infrequent event throughout (indicating misinterpretation of the task) and whose data were therefore disregarded. In addition to course credit, participants could earn a performance-based monetary payoff. Earnings ranged from AU\$4.55 to AU\$10.95 (AU\$1 \approx US\$0.96 at the time of the experiment).

Design and procedure All participants completed a computer-based repeated binary choice task over 500 trials. We factorially crossed two between-subjects factors: the effort involved in implementing a probability maximizing strategy during the choice task (high vs. low) and the presence of an interleaved concurrent working memory load task (present vs. absent). The concurrent task was a 3-back memory task that asked participants to memorize the three numbers last seen on the screen. We randomly assigned 25 participants to each of the four resulting experimental conditions. Participants could earn performance-based payoffs in both the choice and the memory task (if present), which were paid in cash at the end of the experiment. Participants were instructed to earn as much money as possible and to treat both tasks as equally important. After every block of 100 choice trials, participants received feedback on their accuracy and earnings in both tasks, and the main instructions were reiterated for the subsequent block. There was no practice period before the main task, and participants were not informed in advance about the total number of trials.

Following the choice (and concurrent memory) task(s), all participants were asked to complete a short questionnaire assessing their understanding of the underlying probability structure and their strategy use during the choice task (see the [Appendix](#) for details). Specifically, they were asked to estimate the outcome probabilities for each choice alternative and to consider two prediction strategies for ten hypothetical choice trials: (a) choosing the dominant color for all ten trials (i.e., probability maximizing) and (b) choosing the dominant color for seven out of ten trials (i.e., probability matching).

Note that the labels “probability maximizing” and “probability matching” were not used in the questionnaire. Participants were then instructed to indicate which strategy, (a) or (b), their choices most closely resembled (1) early and (2) late in the experiment, and which strategy, (a) or (b), they (3) expected to earn them more money and (4) would use if they were to play the game again (see Koehler & James, 2010, for similar post-task strategy evaluation questions).¹

Choice task The choice task was adapted from Wolford et al. (2004) and involved repeated binary decisions over 500 choice trials. Each trial started with the presentation of either a fixation cross (working memory load absent) or a digit between 0 and 9 (working memory load present) in the center of the computer screen. Figure 1A illustrates the task screen shown to participants in the working memory load conditions. The presentation of the fixation cross/digit served as a cue for participants to predict which of two colored squares—either a red or a green square—would appear next. The green square appeared in 70 % of trials, and the red square in the remaining 30 % (randomized across participants for red and green majority outcomes). Participants were informed that the sequence of red and green squares was random. Each of the colored squares was mapped to a different location on the screen—either above or below the fixation cross/digit (again, randomized across participants)—and participants made their predictions by pressing the up or down arrow key on the computer keyboard. This color–key mapping was shown on the screen throughout the task, as is illustrated in Fig. 1A. Participants earned two cents for each correct color prediction, and they were encouraged to attempt to earn as much money as possible. Following each choice, either a red or a green square appeared on the screen (in the location indicated by the color–key mapping), participants received verbal feedback about the accuracy of their prediction, and earnings were updated on the screen (see Fig. 1A). The next trial then started with either the fixation cross or a new digit. The primary dependent measure was participants’ proportions of choices of the more probable—that is, dominant—color outcome.

The implementation effort of probability maximizing was manipulated by modifying the allocation of response keys to predicting a red or a green square, so that repeatedly choosing the same physical keyboard key resulted in either probability matching or maximizing. In the low maximizing effort conditions, the color–key allocation remained the same throughout the task, and solely pressing the key corresponding to the majority outcome resulted in probability maximizing, as is

¹ The questionnaire also included the Cognitive Reflection Test (Frederick, 2005), the Berlin Numeracy Test (Cokely, Galesic, Schulz, Ghazal, & Garcia-Retamero, 2012), and measures of self-reported mathematical ability and education. Since none of these measures were notably related to our dependent variables, they are not considered further.

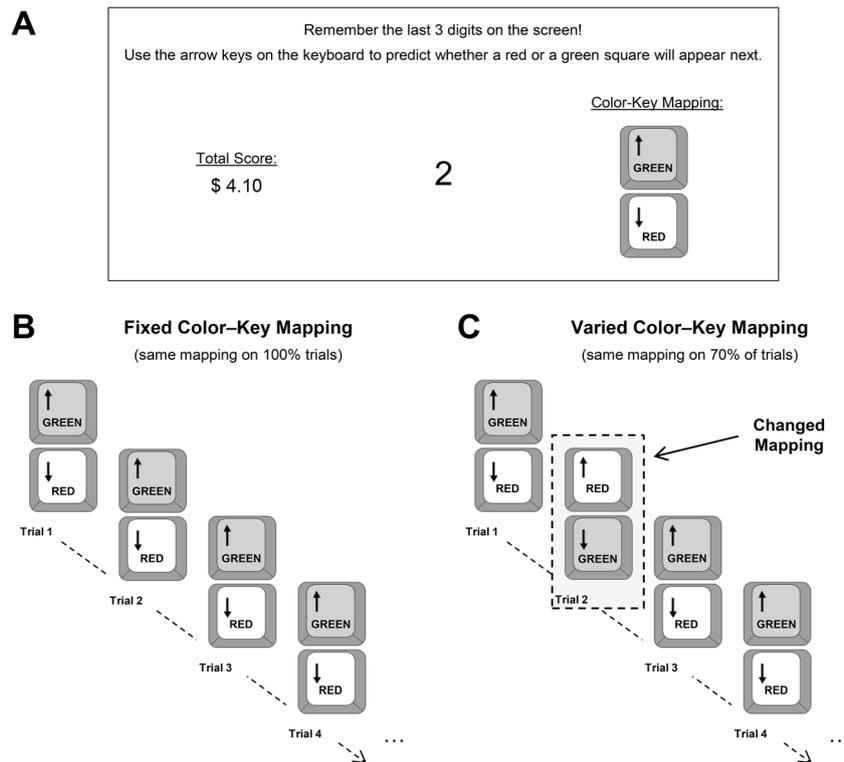


Fig. 1 Illustration of the experimental task screen (panel A) and the strategy implementation effort manipulation (panels B–C). Participants in the working memory load conditions saw a digit in the center of the screen, the current payoff on the left-hand side (updated after each choice), and the current color–key mapping on the right-hand side. In the actual experiment, this mapping illustration was colored (one key was red, the other green), and there was no verbal label. For participants in the no load conditions, the task screen looked identical, except that the digit was replaced by a fixation cross and the instructions to remember the last three digits were omitted. The mapping of response keys to colors was manipulated so that repeatedly selecting the same physical keyboard key resulted in either probability maximizing (panel

B) or probability matching (panel C). In the fixed color–key mapping conditions, the color–key allocation remained the same throughout the task, and solely pressing the up arrow key to predict that a green square would appear above the fixation cross implemented probability maximizing for green majority outcomes. In the varied color–key mapping conditions, the color–key allocation remained the same on 70 % of trials, but flipped on the other 30 % (matching the outcome frequencies), so that solely pressing the up arrow key implemented probability matching for green majority outcomes. For half of the participants, the majority outcome was a green square; for the other half, it was a red square

shown in Fig. 1B. For example, solely pressing the up arrow key to predict that a green square would appear implemented a probability maximizing strategy for green majority outcomes. In the high maximizing effort conditions, the color–key allocation remained the same on 70 % of trials, but flipped on the other 30 %, which matches the frequencies with which the two colors appeared. The switch was shown in the mapping illustration displayed on the screen. Now, solely pressing the key mostly corresponding to the majority outcome resulted in probability matching, as is shown in Fig. 1C. For example, a participant solely pressing the up arrow key throughout the task would predict a green square to appear on 70 % of trials and a red square to appear on 30 % of trials, and would thus implement a matching strategy for green majority outcomes. Consequently, implementation of probability maximizing was made more difficult. Throughout the article, we refer to these two conditions as *fixed color–key mapping* (in which probability maximizing was easy to

implement) and *varied color–key mapping* (in which probability maximizing was difficult to implement).

Working memory task The memory task, which was interwoven with the choice task, required participants in the working memory load conditions to remember the last three numbers shown on the screen. Numbers between 0 and 9 were randomly selected and displayed in the center of the screen at the start of each choice task trial, as is illustrated in Fig. 1A. Once participants had made a choice (there was no time limit), the number was replaced by a fixation cross, followed by feedback on the outcome and earnings. At the start of the next trial, a new digit appeared. Participants were asked to maintain the last three numbers in memory, updating the set of numbers remembered with the appearance of each new digit. At four times at random intervals during each block of 100 choice trials, participants were tested and asked to recall the last three numbers they had seen as accurately as possible. Each correctly recalled digit raised participants’ earnings from the choice task by 5 %.

Results

For all parametric inferential statistics, we conducted Bayesian analyses in addition to using conventional methods of hypothesis testing. On the basis of the default Bayesian analyses of variance (ANOVAs) suggested by Rouder, Morey, Speckman, and Province (2012) and the default Bayesian t tests suggested by Rouder, Speckman, Sun, Morey, and Iverson (2009), we report Bayes factors (BF) that quantify the strength of evidence in favor of the presence of an effect.²

Choice task performance The proportions of participants' dominant color choices for each block of 100 trials, shown in Fig. 2, were subjected to a 2 (working memory load) \times 2 (mapping) \times 5 (block) mixed model ANOVA. The main effect of learning across trial blocks was significant, $F(2.65, 254.17) = 58.77, p < .001, \eta_p^2 = .380, BF = 5.76 \times 10^{34}$, and is illustrated by the upward trajectory of all group lines in Fig. 2.³ We found a main effect of mapping, $F(1, 96) = 10.29, p = .002, \eta_p^2 = .097, BF = 17.84$; participants who experienced fixed color–key mappings chose the dominant color more frequently ($M = .78$ across blocks of 100 trials) than did those who experienced varied color–key mappings ($M = .70$ across blocks of 100 trials). Additionally, the mapping effect significantly interacted with the within-subjects Block factor, $F(2.65, 254.17) = 3.43, p = .022, \eta_p^2 = .034, BF = 2.49$; the learning slopes of participants in the varied color–key mapping conditions remained flatter across blocks than did those of participants in the fixed color–key mapping conditions (see Fig. 2). Participants under working memory load had a slight tendency to select the dominant color *less* often ($M = .72$ across blocks of 100 trials) than did nonloaded participants ($M = .76$ across blocks of 100 trials); however, neither the main effect of working memory load, $F(1, 96) = 2.40, p = .125, \eta_p^2 = .024, BF = 0.73$, nor any of the

² All Bayesian analyses were carried out in R. For the Bayesian ANOVAs and t tests, we used the `anovaBF` and `ttestBF` functions included in the `BayesFactor` package (version 0.9.9; Morey & Rouder, 2014) with their respective default settings, with one exception: For the `anovaBF` function, the number of Monte Carlo samples used to estimate BF s was increased to 50,000. For the Bayesian ANOVAs, we computed BF s by comparing the performance of a model including the relevant effect with that of one omitting (only) that effect. All models were constructed hierarchically, such that the presence of an interaction term always involved the presence of all lower-order interactions and main effects involving the components of that higher-order interaction. Follow-up simple effects were analyzed by carrying out Bayesian t tests. For the Bayesian t tests, BF s in favor of the alternative hypothesis are reported. A BF of 10, for instance, suggests that the data are ten times more likely to have occurred under the model assuming the relevant effect than under a model omitting this effect, whereas a BF of 0.10 indicates that the data are ten times more likely to have occurred under the model omitting the relevant effect than under a model including it.

³ For all conventional ANOVAs, the degrees of freedom were corrected using the Greenhouse and Geisser (1959) coefficient when the sphericity assumption had been violated.

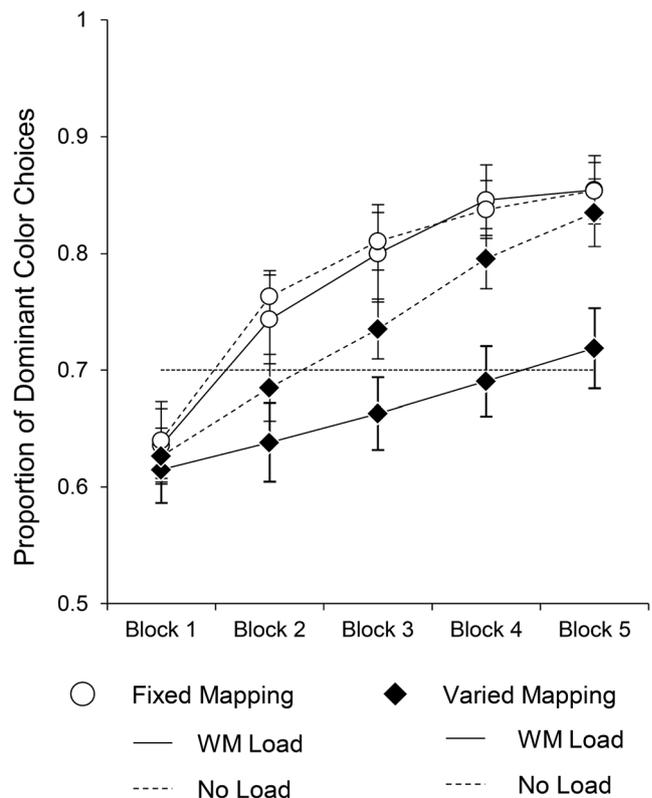


Fig. 2 Mean (\pm standard errors) proportions of participants' dominant color choices in each block of 100 choice trials in Experiment 1 by working memory (WM) load and mapping condition. The dashed line at .70 indicates probability matching

interactions with this factor reached statistical significance (all p s $\geq .184$, all BF s ≤ 0.65).

Turning to the individual-level responses, Fig. 3 displays the full range of participants' proportions of dominant color choices in each trial block and for all conditions. To assess strategy selection in individual participants toward the end of learning, we classified participants' response proportions in the final trial block as either probability maximizing or probability matching. Participants who selected the dominant color on no less than 95 % of trials in the last block were defined as probability maximizers; participants who allocated their choices within 5 % of the average reward probability of the more probable option ($.70 \pm .05$) were defined as probability matchers (see, e.g., Schulze, van Ravenzwaaij, & Newell, 2015). We carried out three-way chi-square tests to evaluate whether the adoption of probability maximizing and probability matching in the final trial block was associated with the color–key mapping manipulation, contingent on working memory load condition. Under working memory load, participants who experienced varied color–key mappings were 7.67 times *less* likely to probability maximize in the final trial block than were participants who experienced fixed color–key mappings, $\chi^2(1) = 7.02, p = .008$. For nonloaded participants, we found no association between maximizing and color–key

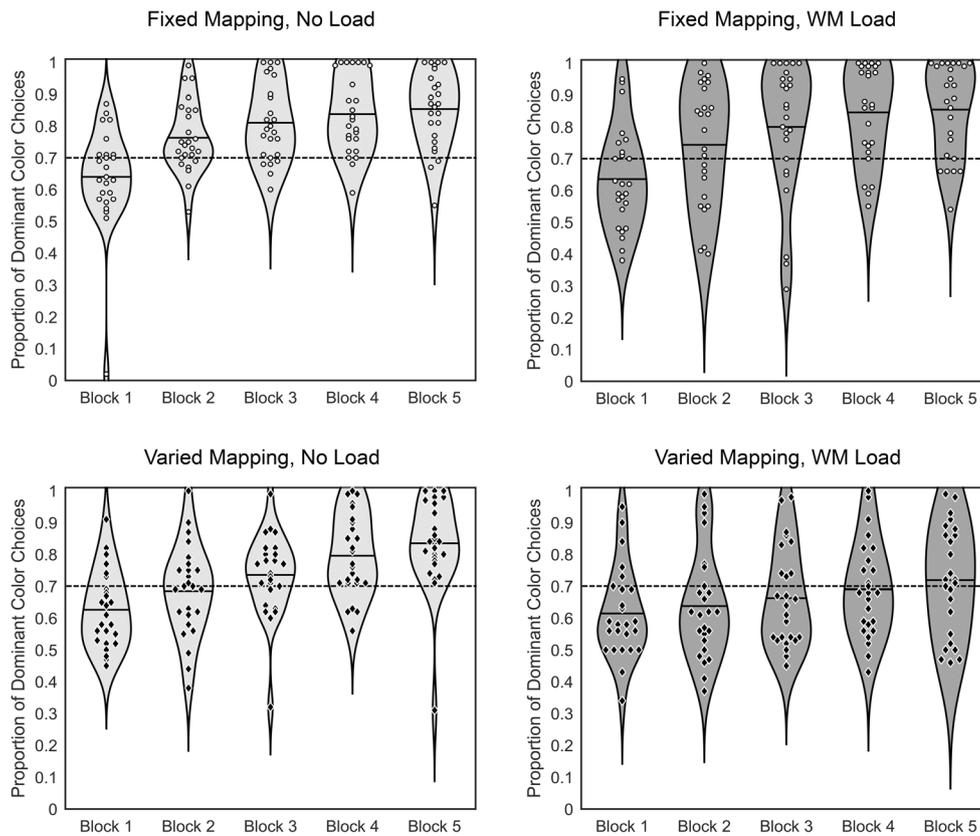


Fig. 3 Individual participants’ proportions of dominant color choices for each working memory (WM) load and mapping condition during each block of 100 choice trials in Experiment 1. Each white circle/black diamond represents the choice proportion of one participant in the fixed/varied mapping condition, respectively, in a specific trial block.

The violins around the circles/diamonds are density estimates for the distributions of participants’ responses in each condition and trial block; black lines represent the means across participants. The dashed lines at .70 indicate probability matching

mapping condition during the final trial block; in fact, the same number of participants (seven out of 25) were classified as probability maximizers in both mapping conditions.⁴ There was no relationship between the mapping manipulation and the use of probability matching in the final trial block, either for participants under working memory load, $\chi^2(1) = 0.10, p = .747$, or for nonloaded participants, $\chi^2(1) = 0.14, p = .713$.

proportion converged on probability matching by Block 5. The full range of individual choice proportions shown in Fig. 3 indicates that, in this condition, the proportion of probability matchers (28 %) equaled the proportion of participants selecting colors at random ($.50 \pm .05$ choices to either option) in the final trial block. Across all other conditions, only two other participants responded randomly in the last block.

Moreover, as Fig. 3 highlights, less than 30 % of the participants in each condition were characterized as probability matchers by the final trial block. Critically, this included participants who experienced varied color–key mappings under working memory load, and whose *average* response

Working memory task and questionnaire responses The proportion of correctly remembered numbers in the working memory load task for each block of the choice task (see Table 1) was subjected to a 2 (mapping) \times 5 (block) mixed model ANOVA. Memory task performance varied significantly across choice task blocks, with lower accuracy scores during both early and late blocks, although the Bayesian evidence was somewhat ambiguous, $F(2.69, 129.24) = 3.09, p = .035, \eta_p^2 = .060, BF = 1.53$. Participants who experienced varied color–key mappings were slightly less accurate on the memory task ($M = .88$ across blocks) than were participants who experienced fixed color–key mappings ($M = .92$ across blocks). However, neither the main effect of color–key mapping nor the mapping by block interaction was

⁴ When we analyzed the adoption rates of probability maximizing contingent on color–key mapping, we found no significant association between working memory load condition and maximizing during the final block for participants who experienced varied color–key mappings, $\chi^2(1) = 3.39, p = .066$, or for participants who experienced fixed color–key mappings, $\chi^2(1) = 0.80, p = .370$. That is, although we observed the same qualitative trend in strategy selection under fixed color–key mappings as was reported by Wolford et al. (2004)—ten maximizers under load versus seven maximizers without load—the effect was less pronounced and did not reach statistical significance. We also note that we had 2.5 times more participants per condition (25 vs. 10) than did Wolford et al.

Table 1 Performance on the memory task: Mean proportions (and standard deviations) of correctly remembered numbers in each block of 100 choice trials

	Color–Key Mapping	
	Varied	Fixed
Block 1	.85 (.16)	.87 (.16)
Block 2	.92 (.17)	.96 (.08)
Block 3	.89 (.21)	.92 (.11)
Block 4	.89 (.20)	.96 (.09)
Block 5	.88 (.22)	.91 (.15)

statistically significant. In fact, the Bayesian analysis provided evidence in favor of the absence of an interaction; $F(1, 48) = 1.41, p = .241, \eta_p^2 = .029, BF = 0.49$, for the main effect of mapping, and $F(2.69, 129.24) = 0.24, p = .847, \eta_p^2 = .005, BF = 0.04$, for the mapping by block interaction.

The data from the post-task questionnaire indicated that participants’ outcome probability estimates were least accurate when they experienced varied color–key mappings under working memory load (averaging at $M_s = .62$ and $.37$ for the two choice alternatives). By contrast, the mean probability estimates in all other conditions deviated no more than .03 points from the programmed outcome probabilities (.70 and .30). A 2 (working memory load) \times 2 (mapping) ANOVA on the absolute distance between the probability estimates for both choice alternatives revealed a significant main effect of mapping, $F(1, 96) = 8.38, p = .005, \eta_p^2 = .080, BF = 8.52$; the participants who experienced varied color–key mappings during the choice task discriminated less well between the two values ($M_{diff} = .32$) than did the participants who experienced fixed color–key mappings ($M_{diff} = .42$), judging them to be closer together than they actually were. This suggests that probability learning progressed less accurately when participants experienced varied color–key mappings. No other effects in this analysis were significant; in fact, the Bayesian analyses provided evidence in favor of the absence of these effects (all $p_s \geq .169$, all $BFs \leq 0.28$).

A similar pattern of results was observed for participants’ strategy endorsements. We carried out three-way chi-square tests to evaluate the association between endorsements of maximizing on three survey items (which strategy was used late in the experiment, which was expected to yield the highest payoff, and which would be used again) and color–key mapping, contingent on working memory load condition, as is summarized in Table 2. Under working memory load, relative to participants who had experienced fixed color–key mappings, participants who had experienced varied color–key mappings were 4.33 times less likely to identify probability maximizing as the strategy that would earn them more money, and 3.69 times less likely to say they would use maximizing in future games. No such relationship was found for self-

Table 2 Mean proportions of endorsements of probability maximizing as the strategy used toward the end of the experiment, the strategy with the highest expected payoff, and the strategy that would be used again; associations between maximizing endorsements and color–key mapping, contingent on working memory (WM) load condition

Survey Item	Varied Mapping <i>M</i>	Fixed Mapping <i>M</i>	Statistics		
			χ^2	<i>p</i>	Odds Ratio
Self-reported					
Load	.52	.72	2.12	.145	0.42
No load	.80	.68	0.94	.333	1.88
Highest pay					
Load	.48	.80	5.56	.018	0.23
No load	.76	.76	0.00	1.00	1.00
Play again					
Load	.52	.80	4.37	.037	0.27
No load	.76	.68	0.40	.529	1.49

For all chi-square tests, $df = 1$. Significant differences are shown in bold

reported strategy use toward the end of learning (“late in the experiment”) or under single-task conditions (see Table 2).

Discussion

Experiment 1 showed that increasing the effort involved in implementing probability maximizing led to *decreased* adoption rates of this strategy. This effect was most severe when participants’ cognitive resources were taxed by a concurrent working memory task. These findings suggest that the cognitive effort of attending to a concurrent task does not (paradoxically) cause people to choose the maximizing option more readily per se (as was suggested by Wolford et al., 2004). If that were the case, we would expect to have seen *increased* probability maximizing in the presence of a concurrent task, regardless of the mapping manipulation. The fact that we did not indicates that other factors, such as the effort involved in implementing a strategy, moderate people’s engagement in probability matching and maximizing under cognitive load. A potential caveat to this conclusion needs to be considered, however. The effect of implementation effort on strategy selection may have been confounded by impaired probability learning when color–key mappings varied under working memory load. Participants in this experimental condition estimated the outcome probabilities least accurately: Many responded at random during the choice task, and half failed to recognize probability maximizing as the superior strategy afterward. In Experiment 2 we addressed this issue and aimed to rule out impaired probability learning as a possible confound by separating probability learning and the imposition of cognitive load.

Increasing the implementation effort of probability maximizing by varying color–key mappings also simplified the implementation of probability matching. This is because the

color–key allocation changed with probabilities that matched the outcome frequencies. Therefore, exclusively selecting the key mostly corresponding to the majority outcome would have resulted in “easy” probability matching. Unlike probability maximizing, however, matching can be implemented in various ways over a long sequence of choices, and it is unlikely that the programmed color–key mapping changes strictly corresponded to participants’ concepts of a probability matched outcome sequence. Thus, it is not surprising that the majority of participants who probability matched when experiencing varied color–key mappings under working memory load did so by implementing their own representations of this strategy, whereas only one participant implemented probability matching by pressing a single key repeatedly throughout multiple blocks.

Experiment 2

Because the results of Experiment 1 suggested a possible confound in the mapping manipulation—namely, impaired probability learning when color–key mappings varied under working memory load—we designed a second experiment to address this issue. Experiment 2 replicated the basic task design of Experiment 1, but separated learning about outcome probabilities from responding under cognitive load versus no load by dividing the task into two distinct parts. Through this procedural change, we aimed to isolate the effects of implementation effort on strategy selection from those on probability learning.

Method

Participants One hundred (73 female, 27 male) undergraduate students from the University of New South Wales, with a mean age of 19.20 years ($SD = 1.73$ years), participated in this experiment in exchange for course credit. Additionally, participants could earn a performance-based monetary payoff. Earnings ranged from AU\$2.10 to AU\$5.60 (AU\$1 \approx US\$0.90 at the time of the experiment).

Design and procedure We again factorially crossed the effort involved in implementing probability maximizing during a sequential choice task with the presence/absence of a concurrent working memory load task, and randomly assigned 25 participants to each condition. The sequential choice task, the implementation effort manipulation (fixed vs. varied color–key mapping), and the 3-back working memory task were identical to those aspects of Experiment 1, with the following exceptions. The choice task was shortened to 200 trials, the concurrent memory task was introduced midway through the choice task, and the provision of outcome feedback was manipulated between two distinct parts of the task (see, e.g., Newell & Rakow, 2007). Part 1 was completed under

single-task conditions by all participants. During this part, outcome feedback was available following each choice (as in Exp. 1), and participants were able to learn the outcome probabilities via trial-by-trial feedback, while experiencing either fixed or varied color–key mappings. In Part 2, the mapping manipulation remained in place, and the concurrent 3-back working memory task was introduced for half of the participants (load vs. no load factor). Additionally, outcome feedback was removed for all participants, but instructions between the parts repeatedly emphasized that the outcome frequencies would remain exactly the same as in Part 1.

To separate these two parts more plausibly into distinct learning and decision-making stages, we did not remunerate choices in Part 1 (but provided outcome feedback). In contrast, accurate choices in Part 2 were rewarded with five cents (but no outcome feedback was given until all choices had been completed). Accuracy on the concurrent memory task again raised earnings from the choice task, as was described for Experiment 1. Note that, in order to keep the overall payoffs comparable between experiments, we increased the earnings per accurate choice from two to five cents (since the total number of potentially remunerated choice trials was reduced to 100).

Results

Choice task performance Figure 4 shows participants’ mean proportions of dominant color choices in each part of the choice task and for all experimental conditions. In Part 1, we found a significant effect of mapping on dominant color choices, $t(98) = -2.83$, $p = .006$, $BF = 6.81$; participants who learned under fixed color–key mappings chose the dominant color more frequently than did participants who experienced varied color–key mappings (Fig. 4A). This mapping effect persisted into Part 2, $F(1, 96) = 9.37$, $p = .003$, $\eta_p^2 = .089$, $BF = 12.12$; in the absence of outcome feedback, participants in the fixed mapping conditions again selected the dominant color more often ($M = .77$ across load conditions) than did those who experienced varied mappings ($M = .68$ across load conditions; see Fig. 4B). We found no significant effect of working memory load; in fact, the Bayesian analysis provided evidence in favor of the absence of an effect, $F(1, 96) = 0.94$, $p = .335$, $\eta_p^2 = .010$, $BF = 0.32$. There was also no significant mapping by load interaction, $F(1, 96) = 1.90$, $p = .172$, $\eta_p^2 = .019$, $BF = 0.62$. Figure 4B shows a slight trend toward higher maximizing rates in loaded relative to nonloaded participants who experienced fixed color–key mappings; however, this simple effect of working memory load at the fixed mapping factor level was not significant, $F(1, 96) = 2.75$, $p = .100$, $\eta_p^2 = .028$, $BF = 0.93$.

Figure 4 plots the distributions of individual participants’ choice proportions alongside the group averages, and we again categorized participants as probability matchers ($.70 \pm .05$ dominant color choices) or probability maximizers ($\geq .95$

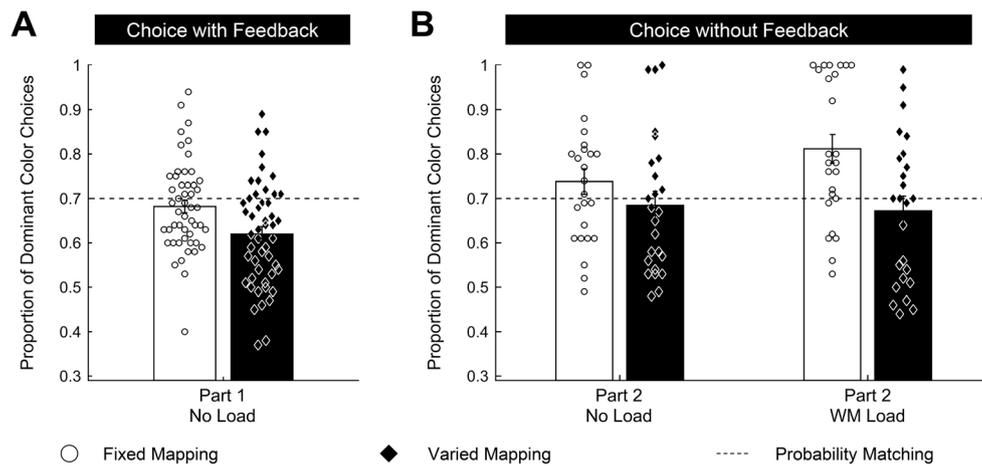


Fig. 4 Choice behavior when outcome feedback was present in Part 1 (A) and absent in Part 2 (B) of Experiment 2. Bar graphs plot the mean \pm standard error proportions of dominant color choices, averaged across participants in each experimental part and condition. The small

white circles/black diamonds plot the dominant color choice proportions of individual participants in the fixed/variied mapping conditions, respectively, of both experimental parts. The dashed lines at .70 indicate probability matching

dominant color choices) on the basis of these individual response proportions in Part 2 (i.e., during the final 100 trials of Exp. 2). We used chi-square analyses to assess the relationship between the mapping manipulation and the adoption of probability maximizing and matching in Part 2, contingent on working memory load. Under working memory load, participants who experienced varied color–key mappings were 6.47 times *less* likely to probability maximize in Part 2 than were participants who experienced fixed color–key mappings, $\chi^2(1) = 5.71, p = .017$. For nonloaded participants, we observed no association between maximizing and color–key mapping in Part 2: The same numbers of participants (three out of 25) probability maximized in both mapping conditions. We again found no relationship between the mapping manipulation and the adoption of probability matching in Part 2 for either loaded participants, $\chi^2(1) = 0.50, p = .480$, or nonloaded participants, $\chi^2(1) = 0.44, p = .508$.

Working memory task and questionnaire responses In Part 2, the proportions of correctly remembered digits in the working memory task again did not differ between color–key mapping conditions; in fact, the Bayesian analysis provided evidence in favor of the null hypothesis, $t(48) = -0.09, p = .930, BF = 0.28$. On average, participants who experienced varied color–key mappings correctly recalled 91 % of the probed digits; those who experienced fixed mappings remembered 92 %.

Although all participants learned the outcome probabilities in the absence of working memory load (Part 1 of the choice task), introducing load in Part 2 of the choice task may nevertheless have affected the subsequent probability estimates in the questionnaire. We therefore subjected the absolute distance between probability estimates for both choice alternatives to a 2 (mapping) \times 2 (load) ANOVA. We found a significant main effect of mapping, although the Bayesian evidence

was ambiguous, $F(1, 96) = 4.67, p = .033, \eta_p^2 = .046, BF = 1.68$; participants who experienced varied color–key mappings discriminated less well between the two values ($M_{\text{diff}} = .29$) than did participants who had experienced fixed color–key mappings ($M_{\text{diff}} = .38$), judging them to be closer together than they actually were. That is, outcome probability estimates were slightly less accurate in the varied color–key mapping conditions ($M_s = .65$ and $.35$) than in the fixed color–key mapping conditions ($M_s = .69$ and $.31$, for the respective choice alternatives across load conditions). We found no effect of working memory load; in fact, the Bayesian analysis provided evidence in favor of the absence of an effect, $F(1, 96) = 0.15, p = .695, \eta_p^2 = .002, BF = 0.23$. There was also no interaction between the load and mapping factors, $F(1, 96) = 1.30, p = .257, \eta_p^2 = .013, BF = 0.49$.

Table 3 summarizes the association between participants' endorsements of probability maximizing on each of three post-task questionnaire items and the color–key mappings they experienced in the choice task, contingent on working memory load condition. In contrast to Experiment 1, the fixed color–key mapping was not related to more frequent endorsements of probability maximizing as the strategy that would earn more money and should be used in future games, regardless of working memory load. Self-reported strategy use, on the other hand, accurately reflected the observed choice behavior; participants who had experienced varied color–key mappings were less likely to report that their choices in Part 2 resembled probability maximizing, irrespective of working memory load condition.

Discussion

Experiment 2 separated learning about the outcome probabilities from decision making under cognitive load, while again

Table 3 Mean proportions of endorsements of probability maximizing as the strategy used in Part 2, the strategy with the highest expected payoff, and the strategy that would be used again; associations between maximizing endorsements and color–key mapping, contingent on working memory (WM) load condition

Survey Item	Varied Mapping <i>M</i>	Fixed Mapping <i>M</i>	Statistics		
			χ^2	<i>p</i>	Odds Ratio
Self-reported					
Load	.28	.68	8.01	.005	0.18
No load	.32	.64	5.13	.024	0.26
Highest pay					
Load	.60	.72	0.80	.370	0.58
No load	.60	.80	2.38	.123	0.38
Play again					
Load	.56	.68	0.76	.382	0.60
No load	.72	.68	0.10	.758	1.21

For all chi-square tests $df = 1$. Significant differences are shown in bold

manipulating the effort required to implement probability maximizing. The general pattern of results replicated our findings from Experiment 1. Part 1 of the choice task showed that varying the mapping of physical response keys to choice options slowed probability learning, thus supporting the caveat that the mapping manipulation affected learning as well as the effort needed to implement probability maximizing. However, varied color–key mapping by no means abolished probability learning. Although slightly less accurate than under fixed color–key mappings, the probability estimates did not resemble random guesses, and probability maximizing was identified as the optimal response by similar proportions of participants in both mapping conditions after the choice task. More critically, Part 2 of the choice task demonstrated that the detrimental effect of varied color–key mappings on maximizing persisted in the absence of learning opportunities and was unaffected by the introduction of working memory load. Thus, we replicated the main results from Experiment 1, showing that increasing the effort involved in implementing probability maximizing leads to decreased adoption rates of this strategy, whereas the additional effort of attending to a concurrent memory task did not substantially affect responding, irrespective of strategy implementation effort.

General discussion

Why does people's poor performance on simple choice problems in single-task settings sometimes improve when dual-task demands are added? We have shown that repeated binary choice under cognitive load is significantly influenced by the effort involved in implementing adequate versus suboptimal strategies. Whereas the optimal solution, probability

maximizing, is typically implemented via a simple repeated motor response, the more common but inadequate response, probability matching, requires more effort. In our study, increasing the effort required to implement probability maximizing in dual-task settings markedly reduced the adoption of this strategy. That is, when the mapping of choice options to physical response keys changed during the task, and simply repeating the same motor response no longer resulted in optimal maximizing, it was harder for participants to choose optimally. This was the case both when they simultaneously learned about the outcome probabilities and dealt with a dual task (Exp. 1) and when these two task demands were procedurally separated (Exp. 2). These results indicate that at least some people might maximize fortuitously rather than intentionally when a concurrent task is introduced, and cast new light on the paradoxical finding that increased task demands sometimes boost sequential choice performance (see Wolford et al., 2004). Moreover, by including strategy implementation effort as a main variable in our design, we found no general effects of cognitive load on choice behavior. That is, although we observed qualitative patterns in strategy selection similar to those reported by Wolford et al. (2004) in our control conditions—in which strategy implementation effort remained confounded with optimal responding (i.e., fixed color–key mapping)—these small load effects remained marginal at best.

We are not the first to fail to find effects of cognitive load on sequential binary choice (see also Otto et al., 2011). The inconsistency of findings may indicate that the link between cognitive load and probability maximizing is not particularly robust, but hinges on the specifics of the experimental setup, such as the difficulty of the concurrent task or the outcome probabilities in the choice task (Otto et al., 2011). In our experiments, for instance, we used less discriminable outcome probabilities than were used in previous studies that had reported load effects (70 %:30 % vs., e.g., 75 %:25 % in Wolford et al., 2004). Therefore, it is important to consider methodological alternatives that tax cognitive capacity by other means than imposing cognitive load. McMahon and Scheel (2010), for instance, depleted participants' blood glucose levels—a manipulation assumed to tax cognitive processing by limiting the metabolic supply available to the brain (e.g., Donohoe & Benton, 1999)—and found that depleted participants maximized more than control participants did. However, this effect may also be related to the differences in the effort required to implement the strategies of probability matching and probability maximizing, a possibility that has not yet been examined in conjunction with the manipulation of blood glucose levels.

Setting aside for a moment the robustness of the link between taxed cognitive resources and probability maximizing, our results call into question the interpretation of this effect as support for the pattern search hypothesis of probability matching (e.g., McMahon & Scheel, 2010; Wolford et al., 2004). Specifically, we have shown that people tend to

probability maximize less when the implementation of this strategy is made difficult by the manipulation of the color–key allocation. Likewise, it is possible that people might abandon intended probability matching when its implementation is made difficult by the introduction of cognitive load. In other words, taxed cognitive resources may limit people’s ability to track their own choices, rather than their ability to track patterns in the outcome sequence. Moreover, in replication of previous findings (Newell & Rakow, 2007), Experiment 2 demonstrated that even in the absence of any outcome feedback—and thus any opportunity to seek and test patterns in the outcome sequence—a substantial subset of participants continued to choose suboptimally, regardless of cognitive load. Because pattern search was not feasible under these conditions, the observed responses were likely generated by a different cognitive process.

It is important to stress, however, that the aim of our study was not to rule out that people sometimes search for patterns in outcome sequences or that this search behavior can contribute to probability matching. On the contrary, we subscribe to a multifaceted view of probability matching and believe that our findings contribute to a growing body of empirical research suggesting that a conglomerate of processes (rather than a single mechanism) underlie the matching phenomenon (for a review advocating this perspective, see Koehler & James, 2014). Gaissmaier and Schooler (2008), for instance, showed that some participants matched by following a simple “win–stay lose–shift” heuristic, whereas others seemed to match as a by-product of a more elaborate quest to identify outcome patterns. Similarly, Otto et al. (2011) found that probability matching can arise from a simple cognitive strategy that considers only the most recent outcome, or from integrating a longer window of past outcomes.

Manipulating strategy implementation effort affected not only choice, but also learning about task characteristics such as outcome probabilities. This result is not surprising, because the varied color–key mapping manipulation essentially burdened participants with an additional task feature to absorb during the experiment. This additional requirement affected learning most noticeably under cognitive load. In other words, the more people needed to learn, the slower their learning progressed and the less adequate their choices became. Thus, the overarching outcome of this work is an inverse relationship between *how much* needs to be managed during a task and *how well* people cope with these demands. This conclusion backs recent criticism of the notion that human cognition could benefit from explicit cognitive effort being withdrawn in favor of implicit processing (Newell, 2015). Specifically, some researchers have claimed that category learning and multi-attribute decisions may improve when deliberative thought is occupied or distracted, so that superior implicit processes can stir human cognition to better outcomes (Dijksterhuis, 2004; Filoteo, Lauritzen, & Maddox, 2010).

Reevaluating the recent research on this claim, Newell concluded that the empirical evidence for advantageous effects of unconscious or implicit thought on category learning and decision making has been overstated, and that engaging in effortful deliberation is often very useful in these situations.

Similarly, we have shown here that poor performance on a singular sequential choice task does not (sometimes) miraculously improve *because* a concurrent memory task is added. Rather, we have argued that a concurrent task increases overall difficulty and drives people to use simpler strategies. The optimal solution in typical sequential choice tasks is, incidentally, also the simplest to carry out. Varying the mapping of physical response keys to choice options removed the implementation advantage of maximizing and reduced people’s engagement in this strategy. Moreover, adding this attention-demanding task feature to an already challenging dual-task–probability-learning mix likely pushed some participants to give up on the task entirely and to respond randomly in Experiment 1 (see Fig. 3). Procedurally separating learning about the outcome probabilities from choice under cognitive load noticeably alleviated this “triple burden” in Experiment 2, but it is worthwhile to consider alternatives to probability learning paradigms when multiple concurrent demands are of interest.

Recent research has applied more descriptive versions of repeated choice problems that disclose the relevant outcome probabilities (and sometimes the outcome-generating process) to participants prior to the task—for example, by asking participants to predict the outcomes of repeated rolls of a fair die (Gal & Baron, 1996; James & Koehler, 2011; Newell & Rakow, 2007; Peterson & Ulehla, 1965). Probability matching is typically observed in both variants, but it tends to occur less when the probabilities are known from the start (Fantino & Esfandiari, 2002) and/or when the generating process can be identified as random (Morse & Runquist, 1960; Peterson & Ulehla, 1965). Instead of examining choice in a two-stage probability learning paradigm like the one used in Experiment 2, an alternative approach might be to apply a descriptive version of the problem to fully relieve participants from the onus of having to learn about outcome probabilities. Currently, we know of no research that has examined the effects of taxed cognitive resources on choice behavior within such descriptive paradigms.⁵

We conclude that the effort involved in implementing a choice strategy is a key factor in shaping sequential

⁵ In light of previous findings, however, we would expect that such paradigms run the risk of ceiling effects in probability maximizing, and thus diminish the potential for finding effects of cognitive load or strategy implementation effort on sequential choice. Preliminary experiments using a “dice paradigm” that were run in the same participant pool with similar sample sizes and identical independent variables corroborated this expectation. When the outcome probabilities were stated at the start, a pattern of choice trends emerged similar to those reported for Experiments 1 and 2; however, the effect of implementation effort diminished strongly (although, again, no effect of cognitive load was detected).

choice—one that deserves careful consideration when interpreting previously reported effects of cognitive load in this context (e.g., Otto et al., 2011; Wolford et al., 2004).

Author note C.S. is now at the Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany. We thank Derek Koehler, Greta James, and Tim Rakow for valuable discussions during the development of this research; Frank Ji for help with data collection; and Susannah Goss for editorial assistance. This research was supported by Australian Research Council grants to B.R.N. (DP110100797; FT110100151).

Appendix

The wording of the post-task questionnaire items was modified to account for the randomization of the dominant color outcome across participants. Here, we reproduce the questionnaire for participants for whom the red square represented the dominant color outcome. Also note that because Experiment 2 was split into two parts—and outcome feedback was provided only in Part 1—the wording of the questionnaire items was modified slightly, as indicated in square brackets.

Post-task questionnaire

1. Think back on your last 100 predictions. [Think back to the first part of the experiment.] What do you think were the chances of:
 - A. A red square appearing on the screen?
 - B. A green square appearing on the screen?
(Please estimate in percentages, e.g., if the respective colored square appeared on every trial, type 100 %.)
2. Think back on the number of times you predicted that a red square would appear on the screen during the experiment. Please estimate in percentages, e.g., if you predicted a red square would appear on every trial, type 100 %.
 - A. Early in [In the first part of] the experiment, over the first few trials?
 - B. At the end [In the second part] of the experiment, over the last few trials?
 - C. If you were to play this game again (on how many trials would you choose red)?
3. Imagine you could give the next person participating in the experiment advice on how to earn as much money as possible. What would you say to them? This is an open response, so feel free to type in whatever you feel is relevant. Use the keyboard to type in your response in the box below this question.
4. Consider the following two prediction strategies for 10 successive trials of the choice task you just completed:
 - 1) Predicting red for all 10 choice trials. 2) Predicting red for 7 choice trials and green for 3 choice trials.
 - A. Early in the experiment [Early in the first part of the experiment], over the first few trials, which strategy, 1 or 2, did your predictions most closely resemble?
 - B. At the end of the experiment [At the end of the second part of the experiment], over the last few trials, which strategy, 1 or 2, did your predictions most closely resemble?
 - C. Which strategy, 1 or 2, would you expect to earn you more money?
 - D. Which strategy, 1 or 2, would you use if you were to play this game again?

References

- Ayton, P., & Fischer, I. (2004). The hot hand fallacy and the gambler's fallacy: Two faces of subjective randomness? *Memory & Cognition*, 32, 1369–1378. doi:10.3758/BF03206327
- Cokely, E. T., Galesic, M., Schulz, E., Ghazal, S., & Garcia-Retamero, R. (2012). Measuring risk literacy: The Berlin Numeracy Test. *Judgment and Decision Making*, 7, 25–47.
- Dijksterhuis, A. (2004). Think different: The merits of unconscious thought in preference development and decision making. *Journal of Personality and Social Psychology*, 87, 586–598. doi:10.1037/0022-3514.87.5.586
- Donohoe, R. T., & Benton, D. (1999). Cognitive functioning is susceptible to the level of blood glucose. *Psychopharmacology*, 145, 378–385.
- Fantino, E., & Esfandiari, A. (2002). Probability matching: Encouraging optimal responding in humans. *Canadian Journal of Experimental Psychology*, 56, 58–63. doi:10.1037/h0087385
- Filoteo, J. V., Lauritzen, S., & Maddox, W. T. (2010). Removing the frontal lobes: The effects of engaging executive functions on perceptual category learning. *Psychological Science*, 21, 415–423. doi:10.1177/0956797610362646
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19, 25–42. doi:10.1257/089533005775196732
- Gaissmaier, W., & Schooler, L. J. (2008). The smart potential behind probability matching. *Cognition*, 109, 416–422. doi:10.1016/j.cognition.2008.09.007
- Gal, I., & Baron, J. (1996). Understanding repeated simple choices. *Thinking & Reasoning*, 2, 81–98.
- Greenhouse, S. W., & Geisser, S. (1959). On methods in the analysis of profile data. *Psychometrika*, 24, 95–112. doi:10.1007/BF02289823
- James, G., & Koehler, D. J. (2011). Banking on a bad bet: Probability matching in risky choice is linked to expectation generation. *Psychological Science*, 22, 707–711. doi:10.1177/0956797611407933
- Koehler, D. J., & James, G. (2009). Probability matching in choice under uncertainty: Intuition versus deliberation. *Cognition*, 113, 123–127. doi:10.1016/j.cognition.2009.07.003
- Koehler, D. J., & James, G. (2010). Probability matching and strategy availability. *Memory & Cognition*, 38, 667–676. doi:10.3758/MC.38.6.667

- Koehler, D. J., & James, G. (2014). Probability matching, fast and slow. In B. H. Ross (Ed.), *The psychology of learning and motivation* (Vol. 61, pp. 103–131). San Diego, CA: Elsevier.
- Kogler, C., & Kühberger, A. (2007). Dual process theories: A key for understanding the diversification bias? *Journal of Risk and Uncertainty*, *34*, 145–154. doi:10.1007/s11166-007-9008-7
- McMahon, A. J., & Scheel, M. H. (2010). Glucose promotes controlled processing: Matching, maximizing, and root beer. *Judgment and Decision Making*, *5*, 450–457.
- Morey, R. D., & Rouder, J. N. (2014). BayesFactor (Version 0.9.9). Retrieved from <http://bayesfactorpcl.r-forge.r-project.org>
- Morse, E. B., & Runquist, W. N. (1960). Probability-matching with an unscheduled random sequence. *American Journal of Psychology*, *73*, 603–607.
- Newell, B. R. (2015). “Wait! Just let me not think about that for a minute”: What role do implicit processes play in higher-level cognition? *Current Directions in Psychological Science*, *24*, 65–70. doi:10.1177/0963721414551958
- Newell, B. R., & Rakow, T. (2007). The role of experience in decisions from description. *Psychonomic Bulletin & Review*, *14*, 1133–1139. doi:10.3758/BF03193102
- Otto, A. R., Taylor, E. G., & Markman, A. B. (2011). There are at least two kinds of probability matching: Evidence from a secondary task. *Cognition*, *118*, 274–279. doi:10.1016/j.cognition.2010.11.009
- Peterson, C. R., & Ulehla, Z. J. (1965). Sequential patterns and maximizing. *Journal of Experimental Psychology*, *69*, 1–4. doi:10.1037/h0021597
- Rouder, J. N., Morey, R. D., Speckman, P. L., & Province, J. M. (2012). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, *56*, 356–374. doi:10.1016/j.jmp.2012.08.001
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian *t* tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, *16*, 225–237. doi:10.3758/PBR.16.2.225
- Schulze, C., van Ravenzwaaij, D., & Newell, B. R. (2015). Of matchers and maximizers: How competition shapes choice under risk and uncertainty. *Cognitive Psychology*, *78*, 78–98. doi:10.1016/j.cogpsych.2015.03.002
- Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, *15*, 233–250. doi:10.1002/bdm.413
- Unturbe, J., & Corominas, J. (2007). Probability matching involves rule-generating ability: A neuropsychological mechanism dealing with probabilities. *Neuropsychology*, *21*, 621–630. doi:10.1037/0894-4105.21.5.621
- Vulkan, N. (2000). An economist’s perspective on probability matching. *Journal of Economic Surveys*, *14*, 101–118. doi:10.1111/1467-6419.00106
- West, R. F., & Stanovich, K. E. (2003). Is probability matching smart? Associations between probabilistic choices and cognitive ability. *Memory & Cognition*, *31*, 243–251. doi:10.3758/BF03194383
- Wolford, G., Newman, S. E., Miller, M. B., & Wig, G. S. (2004). Searching for patterns in random sequences. *Canadian Journal of Experimental Psychology*, *58*, 221–228. doi:10.1037/h0087446