Judgment and Decision Making (2023), **Vol. 18**:e15 1–21
doi:10.1017/jdm.2023.16

EMPIRICAL ARTICLE

Systematic metacognitive reflection helps people discover far-sighted decision strategies: A process-tracing experiment

Frederic Becker[1], Maria Wirzberger[2], Viktoria Pammer-Schindler[3], Srinidhi Srinivas[1], and Falk Lieder[1]

[1]Rationality Enhancement Group, Max Planck Institute for Intelligent Systems, Tübingen, Germany; [2]Department of Teaching and Learning with Intelligent Systems, University of Stuttgart, Stuttgart, Germany and [3]Institute of Interactive Systems and Data Science, Graz University of Technology, Graz, Austria

***Corresponding author.** Frederic Becker; E-mail: frederic.becker@tuebingen.mpg.de

Received: 16 March 2023; **Accepted:** 17 March 2023

Keywords: metacognitive learning; systematic reflection; boosting; planning; decision-making

Abstract
Short-sighted decisions can have devastating consequences, and teaching people to make their decisions in a more far-sighted way is challenging. Previous research found that reflecting on one's behavior can boost learning from success and failure. Here, we explore the potential benefits of guiding people to reflect on whether and how they thought about what to do (i.e., *systematic metacognitive reflection*). We devised a series of Socratic questions that prompt people to reflect on their decision-making and tested their effectiveness in a process-tracing experiment with a 5-step planning task ($N = 265$). Each participant went through several cycles of making a series of decisions and then either reflecting on how they made those decisions, answering unrelated questions, or moving on to the next decision right away. We found that systematic metacognitive reflection helps people discover adaptive, far-sighted decision strategies faster. Our results suggest that systematic metacognitive reflection is a promising approach to boosting people's decision-making competence.

1. Introduction

Consistent with the adage 'Failing to plan is planning to fail', research on judgment and decision-making has found that insufficient foresight is a major source of mistakes that people come to regret (Kinnier and Metha, 1989), such as neglecting their education and failing to save for retirement (Bruine de Bruin et al., 2007; Goda et al., 2019; Wang and Sloan, 2018). One way to address this problem is to increase people's decision-making competence (i.e., boosting; Hertwig and Grüne-Yanoff, 2017) by helping them learn to use more far-sighted decision strategies. Simply telling people about the importance of planning is not enough to achieve this (Larrick, 2004). However, when a person learns such a lesson from experience, it can fundamentally change how they make future decisions (e.g., Heller et al., 2016).

Learning from experience can produce valuable knowledge about the world that can inform future decisions (e.g., 'Debt compounds exponentially fast'.). Importantly, learning from experience can also improve the process of decision-making itself. For instance, a college freshman who failed to start studying for their first midterm exam on time and received a poor grade on it might subsequently learn to plan when, where, and how to study for the final exam. As a consequence, the way in which they decide how to spend their time can become more deliberate, more far-sighted, and more proactive. Learning how to decide can involve gaining explicit knowledge about the pros and cons of different ways of decision-making (e.g., 'I get better grades when I plan weekly study sessions'), the acquisition of implicit, procedural knowledge about how to make good decisions, or both. Both components of learning how to decide are instances of *metacognitive learning* (He et al., 2021; He and Lieder, 2022b; Jain et al., 2019; Krueger et al., 2017; Lieder and Griffiths, 2017) because they involve the acquisition of metacognitive knowledge and (meta)cognitive skills, respectively.

Here, we focus on 2 important types of metacognitive learning: discovering new decision strategies (He and Lieder, 2022b; Jain et al., 2022; R. Siegler and Jenkins, 2014) and learning when to select which decision strategy (Erev and Barron, 2005; Lieder and Griffiths, 2017; Rieskamp and Otto, 2006). While most previous research on metacognition investigated the metacognitive regulation of students' study behavior (Panadero, 2017; Veenman et al., 2006; Zimmerman, 1990) or reasoning and problem-solving (Ackerman and Thompson, 2017; Griffiths et al., 2019), metacognition also plays an important role in strategy discovery (Shrager and Siegler, 1998; R. S. Siegler, 1999) and strategy selection (Lieder and Griffiths, 2017).

The term *systematic reflection* (Ellis et al., 2014) refers to a structured procedure in which people (i) analyze what they have done, (ii) evaluate their performance, (iii) determine which behaviors had a positive impact on their performance outcomes and which behaviors had a negative impact, and (iv) plan how to improve. The structure of this procedure is often provided by asking a person or group to answer a series of questions. Those questions are generally designed to structure reflection on 'knowledge, values, behavior and practice' (Pammer-Schindler and Prilla, 2021). Previous research has found that asking people to systematically reflect on their behavior can boost their subsequent performance (Ellis et al., 2014). For instance, Anseel et al. (2009) found that asking people to reflect on what they did correctly and what they did wrong in a simulated management task significantly increased their subsequent performance on the second instance of the same management task. Moreover, it has been shown that systematic reflection can help people learn from both positive and negative outcomes (Ellis et al., 2006; Ellis and Davidi, 2005). Research in human–computer interaction has shown in many different ways that reflection can be supported by computing technology (Pammer-Schindler and Prilla, 2021), especially in workplace settings (Renner et al., 2020). Step-by-step guidance through the process of systematic reflection is specifically suitable to the implementation of reflection guidance via chatbots (e.g., Kocielnik et al., 2018; Wolfbauer et al., 2020).

Much is known about when and how much systematic reflection on behavior improves subsequent behavior (Ellis et al., 2014), and about helping students reflect on and regulate their study behavior through prompting and learning analytics (e.g., Azevedo, 2005; Azevedo et al., 2012; Bannert et al., 2009; Bannert and Reimann, 2012; Hilliger et al., 2020). However, there is virtually no research on the potential benefits of guiding people to systematically reflect on their decision-making strategies. We refer to the latter as *systematic metacognitive reflection*. Moreover, while we know that systematic reflection on behavior improves behavior, it remains unknown whether systematic metacognitive reflection on decision-making can foster the discovery of adaptive cognitive strategies. Finally, it remains unclear who benefits the most from systematically reflecting on how they reached their decisions, and under which conditions this is most beneficial.

We predicted that systematic metacognitive reflection improves how people make subsequent decisions by fostering metacognitive learning (H1). We hypothesized that those improvements would be driven by people learning to plan more and to use more far-sighted planning strategies (H2). Based on prior work (Sitkin, 1992), we additionally predicted that systematic metacognitive reflection would be especially beneficial after decisions that were made poorly (H3). Moreover, we predicted that reflection would be most beneficial after the first couple of decisions a person makes in a new domain (H4). Finally, because Anseel et al. (2009) found that how much people engage with reflection questions depends on their need for cognition (NFC; Cacioppo et al., 1984), we hypothesized that the effects of systematic metacognitive reflection would be moderated by the person's NFC (H5).

https://doi.org/10.1017/jdm.2023.16 Published online by Cambridge University Press

To test these hypotheses, we experimentally investigate the effects of systematic metacognitive reflection on people's planning strategies in a 5-step sequential decision-making task. We induce systematic metacognitive reflection by asking participants a series of Socratic questions that guide them to reflect on how they made their decisions.

We then measure the effect of metacognitive reflection on how people make subsequent decisions. Changes in people's decision strategies can be difficult to detect because any observed decision could have been generated by numerous different decision strategies. One of the earliest strategies that have been used to address this problem in decisions that require planning is asking participants to think aloud (Simon and Newell, 1971). Another method is to measure which pieces of information the decision-maker acquires during the decision-making process and in which order they acquire them (Callaway et al., 2017; Callaway et al., 2022; Ford et al., 1989; Payne et al., 1993; Willemsen and Johnson, 2011). Previous research has shown that such methods can discern between alternative decision strategies that would be indistinguishable based on participants' choices alone (Johnson et al., 2002; Johnson et al., 2008). Such methods can be used to quantify to which extent people adapt their strategies to the structure of the environment (e.g., Callaway et al., 2022) and to measure learning-induced changes in people's decision strategies (He and Lieder, 2022a; Jain et al., 2022; Payne et al., 1988). We, therefore, use a *process-tracing* method to measure whether and, if so, how reflection changes people's decision process. We then analyze the resulting data with a new computational method for inferring how the decision strategies of individual participants changed from each decision to the next (Jain et al., 2022).

Process-tracing allowed us to demonstrate that engaging people in systematic metacognitive reflection improves the quality of subsequent decisions by increasing how much they think about what to do and improving what they think about and how they think about it. Specifically, we found that asking participants to reflect on how they reached their decisions enabled them to discover adaptive far-sighted decision strategies. Moreover, we found that systematic metacognitive reflection was especially helpful for people who initially planned poorly. The Socratic questions we asked to prompt this reflection are only a small step away from a chatbot that fosters systematic metacognitive reflection about real-life decisions. Therefore, our findings could give rise to a cost-effective intervention for improving decision-making in the real world.

2. Method

We preregistered this experiment and our data analysis on osf.io. Our preregistration is available at https://doi.org/10.17605/OSF.IO/M6DFU. The analysis scripts, the data, and the code for the experiment are available at https://github.com/RationalityEnhancementGroup/ReflectiveLearning.

2.1. Participants

We recruited a total of 265 participants by advertising our study on the online participant recruitment platform Prolific. Participation was restricted to people who had not previously participated in other experiments we ran on Prolific using similar tasks. These sample sizes were chosen to achieve a statistical power of at least 80% for all of our hypotheses, assuming medium-sized effects and a significance level of .05. Participants received a guaranteed payment of £2 for about 20–25 minutes of work and could earn an additional performance-dependent bonus of around £1, which provided an incentive for good task performance. The average wage of the experiment was £7.29.

To be included in the analysis, participants had to meet the following preregistered inclusion criteria. They had to pass the task comprehension quiz in at most 3 attempts. In addition, we excluded participants who reported having previously participated in an experiment using a similar task. Applying these criteria led to the exclusion of 8 participants. The age of the participants

https://doi.org/10.1017/jdm.2023.16 Published online by Cambridge University Press

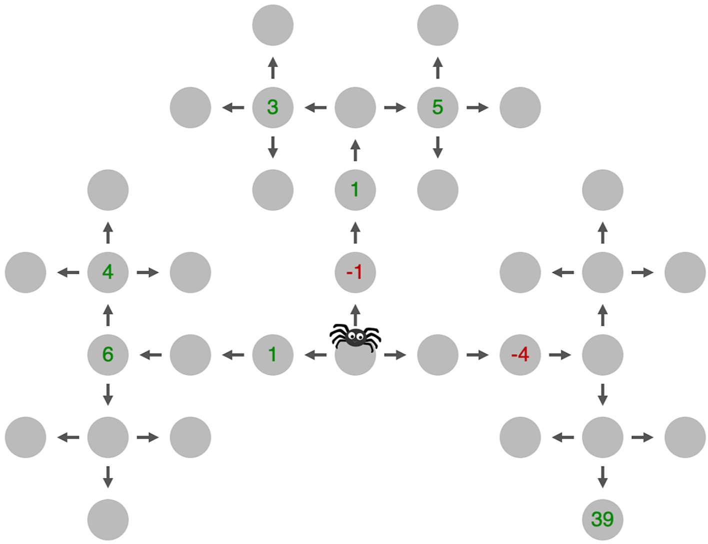

Figure 1. *Screenshot of the planning task. Participants can reveal rewards for a fee to plan a path to one of the target nodes.*

included in the analysis ranged from 18 to 78 years (M = 39.9, SD = 13.6), and 78% of them were female.

2.2. Materials

2.2.1. Planning task

Since it is not possible to observe human planning directly, the underlying cognitive process has to be inferred from people's behavior. To this end, we employed the Mouselab Markov decision process (Mouselab-MDP) paradigm (Callaway et al., 2017; Callaway et al., 2022). Mimicking the challenge that achieving important life goals, such as becoming a doctor, often requires planning multiple steps ahead, the Mouselab-MDP paradigm requires participants to plan a series of actions. Which actions are available in each step depends on which actions the participants choose in the previous steps. As in many challenge real-life decisions, such as choosing a career, which of the initial actions is best primarily depends on the outcomes of later actions. Crucially, the Mouselab-MDP paradigm is designed so that people's information-gathering behavior is highly informative about their planning strategy. In our version of this task, participants were tasked to move a spider from a starting node to one of 18 target nodes. Each possible path consisted of 5 nodes, which contained rewards whose values were initially occluded (Figure 1 shows the task). Participants could reveal the value of a reward by clicking on the corresponding node. A fee of \$1 per click incentivized them to only click on a node when they needed its value for their next planning operation. The participant's sequence of clicks (i.e., which nodes the participant inspected on and in which order they clicked on them) is highly informative about which type of planning strategy they used (e.g., far-sighted vs. near-sighted planning). Once the spider is moved, clicking is no longer possible. The spider uncovers and collects every reward on its way from the start node to the target node. The participants' task is to maximize their game score, which is the sum of the rewards collected by the spider minus the amount spent to uncover the values of the rewards. Rewards are drawn from a Gaussian distribution with mean 0 and standard deviation 1, 2, 4, 8, or 32 for nodes that are 1–5 steps away from the start node, respectively. This means that the rewards near the start nodes vary less than the rewards at the target nodes, making it advantageous to start planning at the target nodes. Participants have to solve 21 instances of this planning task in a row—each time with a different set of hidden rewards.

Table 1. *Reflection prompts that participants in the reflection condition were asked to answer from the first question to the last, grouped by which kind of metacognition they stimulate.*

Level of reflection	Reflection prompt(s)
Description	How many clicks did you make to decide what to do? Where did you click in order to decide what to do? Did you use any particular strategy for selecting your clicks? If so, what was it?
Judgment	How well do you think your current strategy is working?[a]
Evaluation and learning	Why do you think it worked out that way? Based on what you have learned in the last rounds, what tip could you give to a person who performs this task for the first time?
Planning	Based on the previous questions, how many clicks do you plan to do in the next rounds? Based on the previous questions, where do you plan to click in the next rounds? Which strategy do you want to use to select your clicks in the next rounds?

Note: [a] Participants answered this question on a Likert scale with answer choices ranging from 'Very Well' to 'Very Poor'.

2.2.2. Reflection prompts

The reflection prompts started with the instruction 'Please reflect about your planning success in the last 3 rounds by answering a couple of questions'. Participants then received a brief objective feedback on their current performance, displayed for 4 seconds: 'Your average score in the last 3 rounds was X. Your average score in the 3 rounds before that was Y'. Next, they were asked to answer the questions listed in Table 1 in writing, one question at a time. These questions were designed by adapting the reflection script by Wolfbauer et al. (2020) to foster metacognitive reflection about planning in the task shown in Figure 1. The first 3 questions guided the participant to describe their planning strategy. The fourth question asked the participant to judge how well they had planned. The following 2 questions asked the participant to analyze their performance and derive a lesson from it. Next, they were asked to formulate a concrete plan for how to put that lesson into practice in the following trials.

2.2.3. Mundane questions

The control prompts started with a brief objective feedback on their current performance, displayed for 4 seconds: 'Your average score in the last 3 rounds was X. Your average score in the 3 rounds before that was Y'. Next, the participants were asked the following 3 questions: 'Please describe your favorite ¡topic¿?', 'What do you like about it?', and 'What do you dislike about it?' Each prompt addressed a different topic, such as food, books, or sports. The control prompts were designed to match the reflection prompts in their writing and time requirements. The mean response time to reflection prompts was 123 seconds (SD = 111) and to control prompts it was 150 seconds (SD = 97).

2.2.4. Questionnaires

We measured how much each participant liked to think using the Need for Cognition (NFC) Scale (Cacioppo et al., 1984). This scale consists of 18 statements related to the satisfaction one gains from thinking (e.g., 'I prefer complex to simple problems') and asks for the degree of approval on a 5-point scale ranging from 'extremely uncharacteristic of me' to 'extremely characteristic of me'.

In addition, participants in the experimental condition were asked how much effort they had invested into answering the reflection prompts. The answer choices were 'minimal effort', 'some effort', and a 'lot of effort'.

https://doi.org/10.1017/jdm.2023.16 Published online by Cambridge University Press

2.3. Experimental design

Our experiment used a between-subjects design with one experimental group ($N = 128$), one active control group ($N = 65$), and one passive control group ($N = 64$). All groups performed 21 trials of the planning task described above. The 3 groups differed in whether participants were prompted to engage in systematic metacognitive reflection (reflection condition), prompted to answer the mundane questions described above (active control condition), or not prompted to answer any questions during the planning task at all (passive control condition). The rationale of including the active control condition was to control for the amount of time and mental effort participants had to invest to complete the experiment and the time intervals between subsequent decisions. The rationale for including the passive control condition was to examine whether adding reflection prompts has practically relevant benefits that cannot be explained by the absence of the potential negative effect of interrupting participants' learning with mundane questions.

2.4. Procedure

2.4.1. Data collection

After participants gave informed consent and answered the NCS questionnaire, the experiment started with instructions on the planning task. Participants' understanding of the instructions was tested via a quiz comprising 4 basic comprehension questions, for example, asking how to learn the value of a node. If a participant answered one or more questions incorrectly, they had to reread the instructions and retake the quiz until they got all answers right. Participants were then informed about the performance-dependent bonus scheme and completed 21 trials of the planning task. After every third trial except for the last one, the experimental group was prompted to reflect on their decision-making, whereas the active control group was prompted to answer the mundane questions. After completing the 21 trials, participants reported basic demographic information and were asked whether they had participated in an experiment using a similar task before.

2.4.2. Data processing and diagnostics

2.4.2.1. Measuring engagement

We classified participants' responses to the reflection prompts according to the participant's degree of engagement with the question. Each response was categorized as showing either no engagement, low engagement, or high engagement based on the participant's response time and the length and content of their response. A response was categorized as showing no engagement when its content was meaningless (e.g., 'asdf' or '-'). A response was categorized as showing high engagement when it was meaningful and its response time or length were above the fourth quartile of the respective distribution (i.e., longer than 144 seconds or 292 characters). Finally, a response was categorized as showing low engagement when it was meaningful but did not meet the criteria for high engagement or was given too rapidly (i.e., when the response time was below the first quartile of the response time distribution, i.e., 69 seconds). Whether a response was meaningless or meaningful was determined by 2 trained raters. The 2 raters achieved an inter-rater agreement of 99.5% on the first 25% of the responses. Therefore, each of the subsequent responses was categorized by only one rater.

2.4.2.2. Inference of planning strategies and strategy types

The computational microscope (Jain et al., 2022) is a computational method for inferring the planning strategies participants use in the Mouselab-MDP paradigm from their information-gathering behavior. It performs Bayesian inference to determine which of 89 predefined planning strategies is most likely to have generated the sequence of clicks a participant made on a given trial, taking into account which strategy the participants appears to have used in the adjacent trials.[1] The predefined

[1]The prior distribution favors strategy sequences with fewer switches. For the first trial, the prior assigns equal probability to all strategies. For each strategy transition, the prior assigns some probability to the strategy remaining the same. The remainder of

https://doi.org/10.1017/jdm.2023.16 Published online by Cambridge University Press

strategies differ in how much planning they perform, which outcomes they focus on (e.g., immediate outcomes vs. long-term consequences), the order in which different outcomes are inspected (e.g., path-by-path vs. all immediate outcomes first), and in the ways in which the observed outcomes affect whether planning continues (e.g., stop planning upon discovering the highest possible reward) and, if so, how (e.g., if a potential first step yields a positive immediate outcome, then examine its long-term consequences). The set of strategies includes the optimal strategy for this task, which starts by exploring the final outcomes and stops clicking upon finding the maximum value of the reward distribution. The 10 strategies our participants used most frequently are described in Table 2. The computational microscope has been empirically validated on the Mouselab-MDP paradigm; it made accurate inferences[2] and was able to detect the effects of feedback on metacognitive learning (Jain et al., 2022). We employ the computational microscope to identify possible effects of systematic reflection on the temporal evolution of people's decision strategies. Building on He and Lieder (2022b), we grouped the decision strategies our participants used into the 4 types defined in Table 3. Of these 4 strategy types, the far-sighted planning strategies are most adaptive in our task and the no planning strategy is least adaptive.

2.4.2.3. Quantifying the goodness of participants' planning

A key characteristic of the experimental task is that the rewards are drawn from distributions with relatively high variances (see Section 2.2.1 for details). As a consequence, using the very same planning strategy in multiple trials will most likely lead to different scores. This makes the score a relatively noisy measure for the goodness of planning. To overcome this, the computational microscope (Jain et al., 2022) provides an *expected score* for each inferred strategy. The expected score of a strategy is the expected value of the score attained by using the strategy. We estimated it by simulating the application of the strategy to 100,000 different trials of the experimental task and then averaging the attained scores. The expected score can be seen as the expected value of the score distribution of a strategy, whereas the score can be seen as a single draw of it. The expected score thus quantifies the goodness of participant's planning more robustly than the score.

2.4.2.4. Computing outcome measures

In each trial, we recorded the participant's *score* in the planning task, the *planning strategy* inferred from the participant's click behavior, the *expected score* of the planning strategy, and the *type* of the planning strategy. In addition, we measured the *amount of planning* by the number of clicks the participant made prior to choosing a path. In transitions from one trial to the next, we recorded whether the participant's planning strategy changed (*strategy change*), whether the type of the new strategy was different (*strategy type change*), and whether the amount of planning changed (*clicks change*). In addition, we measured the *numerical difference in expected score* and the *numerical difference in the amount of planning* by subtracting the corresponding value of the previous trial from its counterpart for the new trial.

2.4.2.5. Pooling of control conditions

Following our preregistered data analysis plan, we combined the 2 control conditions into a single control condition upon confirming that they did not differ in any meaningful way. To keep the following text simple, we will refer to the pooled data as 'the control condition'. The rationale was to increase the statistical power of the comparisons between the experimental condition and the control condition(s). The passive control condition required 12.6 minutes for the experiment on average. The addition of mundane questions in the active control condition increased this duration to 30.4 minutes to match the duration of the reflection condition (30.2 minutes). An ANCOVA, correcting for NCS and the

the probability is evenly distributed among all alternative strategies. The probability that the strategy remains the same is inferred from the data. For more detail, see Jain et al. (2022).

[2]In simulation studies, the inferred strategy type was correct for 91%–96% of the trials, and the inferred strategy was correct for 76%–88% of the time, depending on how the data were generated. For more information, see Appendix A.3 of Jain et al. (2022).

https://doi.org/10.1017/jdm.2023.16 Published online by Cambridge University Press

Table 2. The 10 most frequently used decision strategies.

Strategy	Used on _% of trials	Used by _% of participants	Used on _% of trials by participants who used this strategy	Strategy type	Expected score
Do not inspect any outcomes (i.e., no planning).	29.8	47.9	62.3	No-planning	0.0
Inspect the final outcomes and stop clicking upon finding the highest possible reward.[a]	13.4	26.8	50.0	Far-sighted	42.9
For each branch, explore one final outcome and its predecessor.	6.1	15.2	40.3	Far-sighted	38.0
Inspect all immediate outcomes and then explore paths that start with a positive immediate outcome.	5.8	15.2	38.5	Far-sighted	25.8
Inspect all final outcomes and then click the predecessor of a positive final outcome and then stop planning.	5.5	11.3	48.4	Far-sighted	22.9
Plan according to Depth First Search and then stop upon finding the highest possible reward.	3.9	10.9	35.9	Other	31.8
Inspect all final outcomes and then inspect the immediate outcomes. If you find the highest possible reward while inspecting final outcomes, then it stops clicking.	3.8	10.5	36.2	Far-sighted	41.0
Inspect the final outcomes. If you find the highest possible reward, then immediately inspect all nodes along the path that leads to the highest possible value, starting from the center of the web. Then stops clicking.	3.3	7.0	46.8	Far-sighted	41.5
Inspect all the final outcomes until you finds the highest possible reward. Then explores the immediate outcome of the path that leads to the highest possible reward.	3.1	9.7	31.2	Far-sighted	42.5
Inspect immediate outcomes. Once you have uncovered a positive immediate outcome, then inspect the final outcomes in such a way that after exploring a final outcome, you next explore its siblings.	2.2	5.1	43.2	Far-sighted	30.7

[a] This is the optimal decision strategy for this task.

https://doi.org/10.1017/jdm.2023.16 Published online by Cambridge University Press

Table 3. *Taxonomy of different types of planning strategies.*

Strategy type	Description	Used on _% of trials	Used by _% of participants	Used on _% of trials by participants who used this strategy type	Expected score
No-planning strategy	This strategy considers no outcome at all.	29.8	47.9	62.3	0.0
Near-sighted strategies	These strategies only consider immediate outcomes.	3.5	18.3	19.4	0.7
Other strategies	These are planning strategies that do not fit in any of the previous categories.	9.4	30.7	30.7	20.9
Far-sighted strategies	These strategies consider final outcomes.	57.2	74.7	76.6	33.9

performance in the first 3 trials, showed that the performance of the 2 control conditions did not significantly differ in terms of score ($F(1, 125) = 0.022, p = .883$) and expected score ($F(1, 125) = 0.026, p = .871$). The average expected score was 18.8 ($SD = 18.1$) in the passive control condition and 22.9 ($SD = 18.6$) in the active control condition.

2.4.3. Data analysis strategy

In the main analysis, we employed linear mixed models (LMMs) for numeric dependent variables and generalized LMMs for binary dependent variables. The exact model formulations and a detailed description of the data analysis are given in our preregistration (Section 2). The complete regression results can be found in Appendix A.3 of the Supplementary Material. Our analysis complies with the best practices outlined in Irwin and McClelland (2001), which means that we describe how we code variables, we include all components of higher-order interactions in our models, and we do not analyze dichotomized variables. We used the following packages from the statistical analysis program R: 'lme4' for fitting, 'lmerTest' to obtain *p*-values, and 'interactions' to resolve interaction effects. We standardized all numeric independent variables before fitting the models to stabilize optimization. We evaluated the models after setting the predictor variable *baseline* one standard deviation below average. Binary variables (reflection, prompt) are coded as dummy codes (0,1). To control for the false discovery rate, we applied the Benjamini–Hochberg procedure to correct the *p*-values for multiple comparisons (Benjamini and Hochberg, 1995). Following the best practices outlined by Spiller et al. (2013), we investigated significant interaction effects using the Johnson–Neyman test. This test determines for which values of a moderator variable the main effect is significant. Further, we have verified that all significant interactions that we reported are crossover interactions if not stated otherwise. Unlike non-crossover interactions, crossover interactions are not removable by monotonic data transformations (Loftus, 1978).

3. Results

As detailed below, our results show that systematic metacognitive reflection boosts people's decision-making competency by fostering fast metacognitive learning. Moreover, we found that the benefit of metacognitive reflection is greatest when people who use poor decision-making strategies reflect on them for the first time.

As illustrated in Figure 2, the expected score of the control condition increased linearly across all trials. By contrast, in the reflection condition, there was a rapid improvement from trials 3 to 7 that

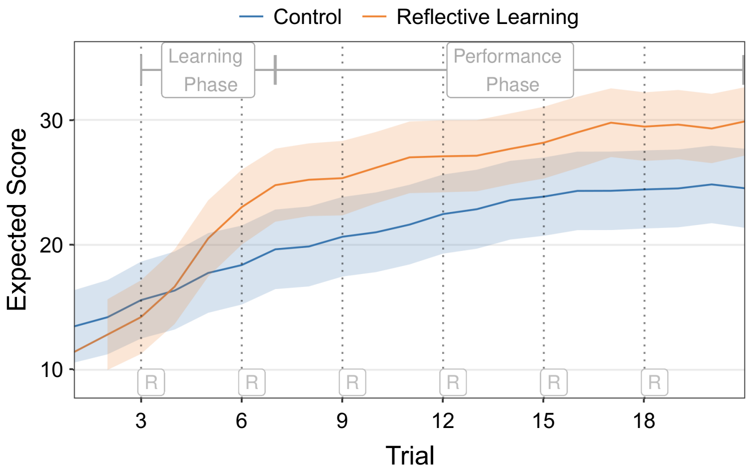

Figure 2. *Expected score as a function of trial number and condition. In the reflection condition after every third trial, a reflection prompt occurred, indicated by the letter R.*

Table 4. *Regression results for the effect of reflection on performance, the amount of planning, and the rate at which they increase with practice (Reflection × Trial no.).*

		Outcome variable					
		Expected score		Score		Clicks	
Trials	Fixed effect	β	p	β	p	β	p
3–7	Reflection	**7.35**	<.001	**8.51**	.015	**2.28**	.001
	Reflection × Trial no.	**4.71**	<.001	2.05	.475	**1.14**	<.001
7–21	Reflection	**12.21**	<.001	**9.96**	<.001	**2.6**	.002
	Reflection × Trial no.	**−0.74**	.021	0.28	.812	**−0.39**	.048

Note: All p-values were corrected for multiple comparisons. Significant predictors are marked in bold. Interaction effects are denoted by a cross (×).

slowed down afterward. According to the LMM summarized in Table 4, the slopes of the expected score differed significantly between those two time periods, producing a kink at trial 7 (z-test for difference in slope: $z = 5.80, p < .001$). To meet the linearity assumption of linear models and to obtain accurate estimates of the slopes, we analyzed trials 3–7 separately from trials 7–21, instead of analyzing trials 3–21 jointly as intended in the preregistration. These 2 phases overlap to include the information about the change from trial 7 to trial 8 in our analysis. Because most of the learning occurred in trials 3–7, we refer to them as the *learning phase*. Relatedly, because people's performance was stable in the subsequent trials, we refer to trials 7–21 as the *performance phase*. We also found that there was no significant difference in performance between the 2 groups prior to the first reflection prompt (i.e., the first 3 trials) ($t(255) = 0.8, p = .42$).

3.1. H1: Systematic metacognitive reflection boosts performance and learning[3]

3.1.1. Improved performance

Considering all trials, the reflection group outperformed the control group in terms of their expected score ($M = 24.5, SD = 17.5$ points vs. $M = 20.9, SD = 18.4$ points) and their score ($M =$

[3]See preregistration: Analyses B1.H1 and B1.H3.

https://doi.org/10.1017/jdm.2023.16 Published online by Cambridge University Press

Table 5. *Regression results concerning the use of different strategy types.*

		Outcome variable							
		No-planning		Near-sighted		Far-sighted		Other	
Trials	Fixed effect	β	p	β	p	β	p	β	p
3–7	Reflection	**−11.2**	.044	26.25	.15	1.5	.549	**3.04**	.044
	Reflection × Baseline	−0.83	.894	−11.1	.319	−0.48	.894	1.42	.561
7–21	Reflection	**−7.66**	.023	3.7	.633	**11.55**	<.001	1.13	.367
	Reflection × Baseline	**−13.0**	<.001	−3.83	.307	**−13.1**	<.001	0.96	.466

Note: All p-values were corrected for multiple comparisons. Significant predictors are marked in bold. Interaction effects are denoted by a cross (×). The predictor *Baseline* is the number of trials in which the participant used the corresponding strategy type in the first 3 trials.

22.8, $SD = 30.5$ points vs. $M = 21.4$, $SD = 32.2$ points). As summarized in Table 4, these differences were significant in both the learning and performance phases.

3.1.2. Accelerated learning

In the learning phase, the reflection group compared to the control group had a steeper learning curve in their expected score ($M = 2.6$, $SD = 4.4$ points/trial vs. $M = 1.0$, $SD = 2.7$ points/trial) and their score ($M = 3.6$, $SD = 10.2$ points/trial vs. $M = 1.9$, $SD = 10.2$ points/trial). This difference in slope was significant for the expected score (Table 4).

3.2. H2: Systematic metacognitive reflection improves how and how much people plan

Having found that systematic metacognitive reflection improves performance, we now investigate which reflection-induced changes in decision-making are responsible for this improvement. We start by examining the cumulative effect of repeatedly reflecting on different decisions. Then, we zoom in on the immediate effect of a single reflection session.

3.2.1. Cumulative effects of reflection on planning[4]

Looking at the cumulative effect of reflection, we found that systematic metacognitive reflection improves both how people plan and how much people plan.

3.2.1.1. Reflection improves how people plan

Considering all trials, the reflection group used the no-planning strategy in 22.8% of trials and the control group in 36.7% of trials. Consistent with the hypothesis that reflection helped people to improve their planning strategies, we found that the reflection group used the no-planning strategy significantly less often (Table 5) than the control group in the learning phase. In the performance phase, the effect of reflection was moderated by how often participants used the no-planning strategy in the first 3 trials, which we will refer to as the *baseline* trials. The Johnson–Neyman test revealed that reflection reduced the use of the no-planning strategy within the group of participants who used this strategy in the majority of the baseline trials.

For far-sighted strategies, we observed the opposite pattern. Considering all trials, the reflection group used far-sighted planning strategies in 60.3% of trials and the control group in 54.2% of trials. Again, the effect of reflection in the performance phase was moderated by the initial use of far-sighted strategies (Table 5). Using the Johnson–Neyman test, we found that reflection increased the use of far-sighted strategies within the group of participants who did not use far-sighted strategies in the baseline trials.

[4]See preregistration: Analyses B1.H6, C1.H1, B1.H1, and B1.H3.

https://doi.org/10.1017/jdm.2023.16 Published online by Cambridge University Press

Table 6. *Regression results for the effect of reflection on the frequency and magnitude of changes in performance and the amount of planning from one trial to the next.*

	Outcome variable									
	Change of						Magnitude of change of			
	Strategy		Strategy type		Clicks		Expected score		No. of clicks	
Fixed effect	β	p	β	p	β	p	β	p	β	p
Reflection	**0.98**	<.001	**0.69**	.013	**1.13**	<.001	**2.69**	<.001	−0.1	.721
Reflection prompt	0.63	.088	0.85	.074	**0.67**	.006	0.75	.278	**1.31**	.006
Reflection prompt × Previous expected score	−0.17	.451	−0.26	.451	−0.28	.257	**−0.8**	.045	−0.26	.436

Note: All p-values were corrected for multiple comparisons. Significant predictors are marked in bold. *Previous expected score* denotes the expected score of the strategy used in the previous trial. Interaction effects are denoted by a cross (×).

These differences might arise partly because reflection helps people overcome the no-planning strategy. Consistent with this hypothesis, an exploratory follow-up analysis showed that participants who did not plan at all in the baseline trials learned to use far-sighted strategies significantly more often when they were in the reflection condition than when they were in the control condition (β = 3.39, $p = .032$).

Considering all trials, the control condition used near-sighted strategies in 2.9% of trials and undefined strategies in 6.2% of trials. The reflective learning group used near-sighted strategies in 4.2% of trials and undefined strategies in 12.7% of trials. We found that the reflective learning group used undefined strategies significantly more often than the control group in the learning phase (Table 5).

3.2.1.2. Reflection fosters strategy exploration

Supporting the interpretation that these improvements were achieved through the exploration of alternative decision strategies, we found that reflection caused participants to change their decision strategy more often. That is, compared with the control group, the reflection group more frequently changed their decision strategy (1.8 times vs. 2.5 times), and the type of their decision strategy (0.7 times vs. 1.1 times). Both differences were statistically significant (Table 6). Accordingly, the average improvement per trial, in terms of the expected score, was significantly higher in the reflection condition than in the control condition (0.92 points per transition vs. 0.55 points per transition; Table 6).

3.2.1.3. Reflection improves how much people plan

The reflection group performed significantly more clicks on each trial than the control group (M = 6.3, SD = 6.8 clicks vs. M = 4.7, SD = 6.1 clicks; Table 4). This is an improvement because the strategy with the highest expected score performs 14.9 clicks on average. The reflection group also learned significantly faster to increase their number of clicks in the learning phase (M = 1.0, SD = 1.9 clicks/trial vs. M = 0.4, SD = 1.3 clicks/trial; Table 4). In addition, the reflection condition changed their number of clicks significantly more often than the control condition (13.4 times vs. 11.2 times; Table 6).

3.2.2. Immediate effects of reflection on planning[5]

We then took a closer look at what happens immediately after a person engages in systematic metacognitive reflection, and how that differs from the changes that occurred in the control condition and in the trials of the reflection condition where no prompt was provided.

[5]See preregistration: Analysis C1.H1.

https://doi.org/10.1017/jdm.2023.16 Published online by Cambridge University Press

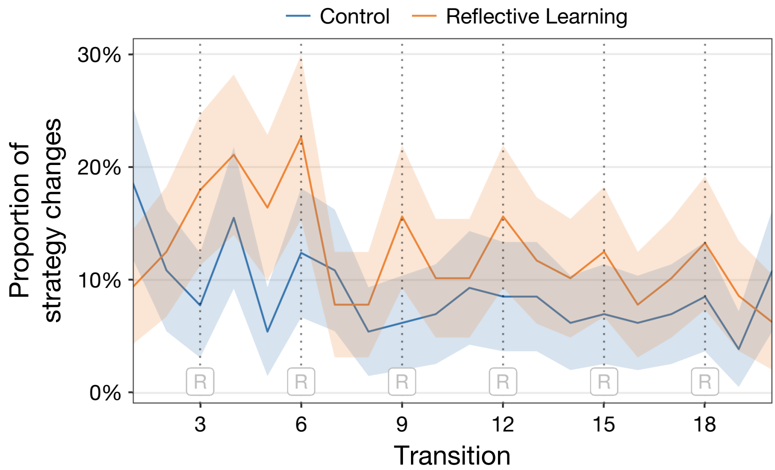

Figure 3. *The proportion of performed strategy changes as a function of transition number and condition. Participants in the reflection condition reflected on their planning success in every third transition (R).*

3.2.2.1. Reflection fosters strategy exploration on the next trial

In the reflection condition, participants changed their planning strategy more often immediately after a reflection prompt than in transitions without a reflection prompt (16.3% vs. 10.7%). The relation between reflection prompts and the number of strategy changes is illustrated in Figure 3. Overall, reflection prompts led to a not statistically significant increase in the proportion of participants who switched to a different planning strategy in the next trial (Table 6).

3.2.2.2. Reflection leads to increased planning on the next trial

Reflection had an immediate effect on how much people planned on the very next trial. That is, the reflection group changed their number of clicks significantly more often after transitions with reflection prompts than after transitions without reflection prompts (72.5% vs. 64.4%; Table 6). In addition, the average change in the number of clicks was significantly larger after transitions with reflection prompts than after transitions without reflection prompts (+1.2 clicks vs. −0.2 clicks; Table 6).

3.3. H3: Reflection is most beneficial when the decision was made poorly

Having found that systematic metacognitive reflection is beneficial for people *on average*, we now examine whether and, if so, how these benefits depend on the person and the situation. Concretely, we examine if the benefit of reflection depends on the person's initial planning skills, their NFC, which strategy they used, and the outcome of their decision.

3.3.1. People who plan poorly benefit more

3.3.1.1. Reflection is especially helpful for low performers[6]

We found that the average expected score in the first 3 trials moderated the effect of reflection on expected score and score (Table 7). The moderation implies that the effectiveness of reflection decreased with increasing baseline performance, which is consistent with the fact that people who start out with good planning strategies have less room for improvement. The Johnson–Neyman procedure revealed significant main effects of reflection on expected score and score for participants who were in

[6] See preregistration: Analyses B1.H4 and B3.

https://doi.org/10.1017/jdm.2023.16 Published online by Cambridge University Press

Table 7. *People's performance in the baseline trials moderates the effect of reflection on their performance and their amount of planning in the subsequent trials.*

		Outcome variable					
		Expected score		Score		Clicks	
Trials	Fixed effect	β	p	β	p	β	p
3–7	Reflection × Baseline	**−3.84**	.001	**−7.81**	.001	−0.41	.416
	Reflection × Baseline × Trial no.	**−2.48**	<.001	0.32	.845	**−0.49**	.033
7–21	Reflection × Baseline	**−6.52**	<.001	**−6.89**	<.001	−0.32	.606
	Reflection × Baseline × Trial no.	**0.6**	.015	−0.1	.904	**0.3**	.039

Note: All p-values were corrected for multiple comparisons. Significant predictors are marked in bold. The baseline value was given by the average value of the corresponding outcome variable in the first 3 trials. Interaction effects are denoted by a cross (×).

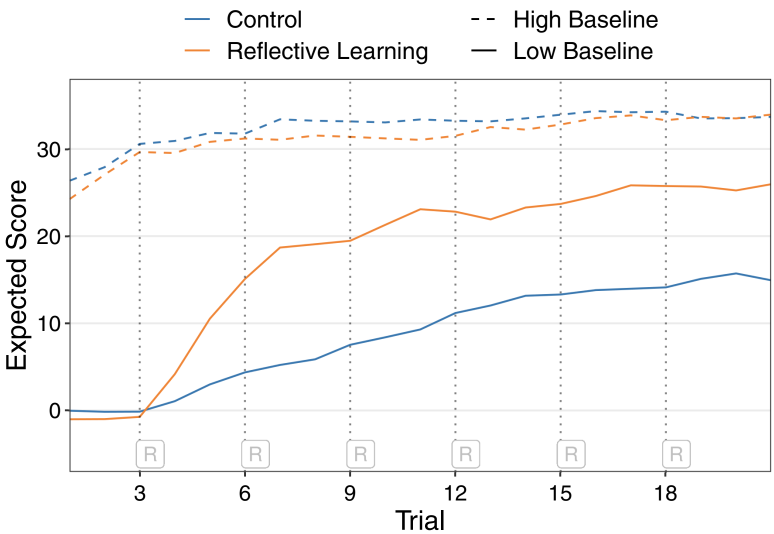

Figure 4. *Expected score as a function of trial number and condition and baseline performance. Participants in the reflection condition reflected on their planning success after every third trial (R).*

the lower half of the baseline performance range (i.e., who had an average expected score in the first 3 trials that was below median).

Accordingly, in trials 3–7, the learning curves were steeper for participants with lower baseline performance as illustrated in Figure 4. We found that reflection increased the rates at which the expected score and the number of clicks increased with the trial number even more strongly for low performers (Table 7). In trials 7–21, the effect of baseline performance on the learning rate was significantly smaller in the reflection condition than in the control condition (Table 7). As illustrated in Figure 4, this suggests that low performers in the control group still had more room for improvement left after the first 7 trials than low performers in the experimental group.

3.3.1.2. Reflection is especially helpful for bad decision-makers

To examine how the effect of reflection depends on the participant's initial planning strategy, we used the type of the strategy participants used in the first 3 trials as a moderator. Using this analysis, we found that the effect of reflection was greater for people who started with the no-planning strategy. They experienced greater increases in score ($\beta = 5.64, p = .01$), expected score ($\beta = 4.93, p = .01$), and

https://doi.org/10.1017/jdm.2023.16 Published online by Cambridge University Press

number of clicks ($\beta = 1.56, p = .019$) in the reflection group than in the control group. In addition, we found that participants who started with near-sighted decision strategies experienced a greater increase in score ($\beta = 4.63, p = .037$) in the reflective group than in the control group.

3.3.2. Reflecting on bad planning is more beneficial than reflecting on good planning[7]

We now turn to the question of how the immediate effect of reflection depends on how the participant planned on the previous trial and the outcome of their decision.

3.3.2.1. The quality of planning on the previous trial moderates the effect of reflection

We investigated whether the expected score of the decision strategy used in the trial preceding a reflection prompt alters the immediate effects of the reflection prompt. We found that the beneficial effect of reflection on the expected score of the strategy used on the next trial was larger when the expected score on the preceding trial was lower (Table 6).

3.3.2.2. Effect of the previous score

We investigated whether the planning success (score) the participant experienced in the trial preceding a reflection prompt altered the immediate effects of the reflection prompt. We found that with decreasing planning success, participants were more likely to change their strategy type ($\beta = -0.69$, $p = .039$) and to increase their expected score ($\beta = -1.22, p = .001$) immediately after a reflection prompt. This is again consistent with the interpretation that reflecting on poor planning is especially beneficial.

3.3.2.3. Reflecting on bad strategies

In an exploratory analysis, we further found that the addition of a reflection prompt led to significantly larger improvements in expected score ($\beta = 8.29, p < .001$) when the participant had used a near-sighted strategy on the previous trial. Further, we found that the number of clicks changed significantly more often after a reflection prompt, when the participant had used the no-planning strategy on the previous trial ($\beta = 0.93, p = .044$).

In Appendix A.1 of the Supplementary Material, we show that our findings are robust to the influence of the no-planning strategy and that reflection also helps people who are already planning to switch to more adaptive planning strategies.

3.4. H4: Repeated reflection on the same kind of decision has diminishing returns[8]

To test our fourth hypothesis, we investigated whether the effect of reflection changed from the first reflection period (i.e., trials 4–6), to the second reflection period (i.e., trials 7–9), to the third reflection period (i.e., trials 10–12). We limited the analysis to these 3 periods as they capture most of the learning dynamics, as illustrated in Figure 2.

3.4.1. The frequency of change appears to decrease over time

On average, 18.5%, 12.8%, and 12.0% of the participants in the reflection group changed their strategy in the first, second, and third reflection periods, respectively. By contrast, only 9.6%, 9.6%, and 7.5% of the participants in the control group changed their strategy in the first, second, and third reflection periods, respectively. The decrease was not significantly larger in the reflection condition than in the control group ($\beta = -0.17, p = .24$). The 2 groups significantly differed in the first and third reflection periods ($\chi^2(1, 257) = 12.0, p = .001; \chi^2(1, 257) = 3.92, p = .048$) but not in the second period ($\chi^2(1, 257) = 1.68, p = .2$).

[7] See preregistration: Analyses C1.H2, C3, and EA1.
[8] See preregistration: Analysis D1.

https://doi.org/10.1017/jdm.2023.16 Published online by Cambridge University Press

3.4.2. The amount of improvement decreases over time

We found a significant decrease in the effect of reflection on how rapidly the expected score increased from the 3rd trial to the 12th trial ($\beta = -0.87, p = .003$). The first reflection prompt was the most effective. On average, the reflection group improved by 2.9, 0.8, and 0.6 points/trial in the first, second, and third reflection periods, respectively. By contrast, the control group improved by only 0.9, 0.8, and 0.6 points/trial in the first, second, and third reflection periods, respectively. The 2 groups significantly differed in the first reflection period ($W = 7092, p = .021$) but not in the following ones (all $p > .68$).

3.4.3. Engagement decrease only slightly over time

One possible reason why improvement decreased over time could be that participants became increasingly more disengaged over the course of the experiment. Indeed, we found that the level of engagement slightly decreased over time ($r(644) = -.09, p = .029$). Although statistically significant, the absolute change in the proportion of engaged responses was minimal. Comparing the first 3 reflection prompts to the last 3 reflection prompts, the proportion of responses categorized as showing no-engagement increased from 0.8% to 1.8%, the proportion of responses categorized as showing low engagement increased from 63.3% to 64.9%, and the proportion of responses categorized as showing high engagement decreased from 35.9% to 33.3%. In general, we found that our participants' engagement was encouragingly high. In the reflection condition, 55.5% of participants reported having invested a lot of effort into answering the prompts, and 44.5% of participants reported having invested some effort into answering them. The median response length to a reflection prompt was 188 characters ($M = 236$), and the median response time was 97 seconds ($M = 123$).

3.5. H5: Need for cognition does not moderate the effect of systematic metacognitive reflection[9]

We did *not* find that NFC moderates the effect of reflection on the expected score ($\beta = 0.63, p = .737$). Nor did we find that NFC moderated the effect of reflection on how quickly the expected score increased with practice ($\beta = 0.25, p = .555$). This indicates that the reflection prompts were effective regardless of the strength of the participant's NFC. Furthermore, we found that NFC was uncorrelated with the number of clicks that participants performed in the first 3 trials ($r(255) = .02, p = .802$) and with the baseline performance ($r(255) = .09, p = .132$).

Based on Anseel et al. (2009), we had hypothesized that people with higher NFC would engage more deeply in the reflection, and that higher engagement would yield more learning. Contrary to our prediction, we found that higher NFC was not correlated with higher engagement in answering the reflection questions ($r(126) = .07, p = .422$). Moreover, whether participants answered the reflection prompts with high engagement or not did not affect the frequency of strategy changes (18.3% vs. 15.2%; $\beta = 0.18, p = .551$) or the increase in the number of clicks ($M = 1.4, SD = 6.9$ vs. $M = 1.1, SD = 6.0$; $\beta = 0.58, p = .449$). However, for participants' expected score, we found a significant interaction of high engagement with previous performance ($\beta = -1.52, p = .043$). The Johnson–Neyman procedure revealed that high engagement, compared with no engagement and low engagement, promoted significantly larger improvements in expected score for participants whose previous performance was below the 30th percentile. Thus, when people reflect on poor planning, higher engagement leads to a larger improvement in their expected score.

Further, we have conducted an exploratory analysis of the effect of self-evaluation skills on reflection benefits which we report in Appendix A.2 of the Supplementary Material.

4. Discussion

This article introduced a new approach to improving human decision-making, namely prompting people to engage in systematic metacognitive reflection. Unlike the forms of systematic reflection that have

[9]See preregistration: Analysis B1.H5.

https://doi.org/10.1017/jdm.2023.16 Published online by Cambridge University Press

been studied previously (Ellis et al., 2014), our Socratic questions direct people's attention to the mental processes that generated their decisions. It does so in a very granular, step-by-step manner that walks people through the metacognitive operations necessary to discover far-sighted decision strategies. Moreover, using a new process-tracing method to measure the reflection-induced changes in people's planning strategies allowed us to rigorously characterize the effects of systematic (metacognitive) reflection in more detail than was previously possible.

We found that systematic metacognitive reflection led to rapid improvements in decision-making. Participants who were prompted to systematically reflect on their decision-making changed their decision strategy more often, improved faster, performed better, and adopted more adaptive, far-sighted decision strategies than participants who practiced decision-making without systematic metacognitive reflection. Moreover, we found that the benefits of metacognitive reflection were not limited to helping people overcome the no-planning strategy. Reflection also helped people who were already planning to switch to more adaptive planning strategies (Appendix A.1 of the Supplementary Material).

In our experiment, reflecting on instances of poor decision-making was especially beneficial. Systematic metacognitive reflection helped many of our participants to overcome poor decision strategies by learning to plan more and becoming more far-sighted. In other words, systematic metacognitive reflection was most effective for the people who most needed to improve. This is good news, given that prior work suggested that reflection only benefits people who already have considerable domain knowledge (Kirschner et al., 2006; Renner et al., 2020). Whether this difference is due to the generality of the decision strategies, we asked our participants to reflect on, or due to the simplicity of our planning task remains to be seen.

Unlike Anseel et al. (2009), we did not find that people with a lower NFC reflected less. This might be because our reflection prompts were more numerous and more detailed than those by Anseel et al. (2009). We required our participants to provide written responses to 9 Socratic questions that (i) explicitly directed their attention to the most relevant aspects of their decision strategy and (ii) guided them through the process of metacognitive learning, from describing how they reached their decisions to planning how to implement the lesson they learned from the outcomes of their decisions (Table 1). This detailed guidance might have succeeded to engage people with a low NFC in crucial steps of reflection that they might have skipped if we had asked only 4 questions, as Anseel et al. (2009) did.

In our experiment, reflecting once was enough to significantly improve decision-making. This suggests that our reflection prompts might enable people to rapidly improve their decision-making with minimal effort. In our experiment, subsequent reflection sessions led to increasingly less improvement. Therefore, future work should investigate how enhanced metacognitive learning can be sustained over an extended period of time.

The main limitation of the present work is that it investigated metacognitive reflection in a single artificial task and used Socratic questions that are somewhat specific to that task. Therefore, future work should investigate whether our findings generalize to other, more naturalistic, and less structured scenarios. For instance, our task is an instance of the decisions from description paradigm, but many real-world decisions have to be made from experience (Hertwig and Wulff, 2022). This makes investigating whether, when, and how metacognitive reflection can improve decisions from experience an important task for future research. Ultimately, the real question is whether the principles of systematic metacognitive reflection instantiated by our reflection prompts are also effective in the real world. If this were true, then the principles of effective metacognitive reflection identified in this article could be applied to develop interventions that help people and organizations leverage their real-life experience to learn how to make better decisions. By guiding the user to reflect on how they arrived at their best and their worst decisions, such chatbots could help people gain valuable insights into what might be the most effective ways to make different types of decisions and how to avoid catastrophic mistakes.

Scaffolding metacognitive reflection could also be combined with other approaches to boosting decision-making, such as providing decision-makers with descriptions or simulated experience

https://doi.org/10.1017/jdm.2023.16 Published online by Cambridge University Press

(Hertwig and Wulff, 2022). Recent findings suggest that, without further assistance, investors may learn very little from simulated decisions (Hertwig and Wulff, 2022; Lejarraga et al., 2022). Our results suggest that adding scaffolded reflection to the simulation-based training could be helpful. Other research has found that the benefits of describing the risks associated with a decision are also limited by how people use the described information (Hertwig and Wulff, 2022). Based on our findings, prompting people to reflect on how they used the information provided by such descriptions in their decisions might be able to improve it, especially if such reflection is embedded in a series of (simulated) decisions with observed outcomes. At this point, these ideas are mere speculation. This makes testing whether metacognitive reflection can boost the efficacy of boosting decision-making with descriptions or simulated experience an important direction for future research.

While we found that a particular form of systematic metacognitive reflection was effective, many more questions remain. For instance, since we only tested a single series of reflection prompts, it remains unclear how much each step in the reflection process contributes to the effect of our intervention. Moreover, systematic reflection can be performed in many ways. Therefore, testing the effectiveness of our reflection prompts against alternative reflection prompts that foster other forms of reflection, such as comparison, explanation, and counterfactual reasoning, is an important avenue for future research. Furthermore, future research should investigate to which extent the effectiveness of reflection depends on directing people's attention to specific crucial aspects of their decision-making. More fundamentally, the cognitive mechanisms through which reflection boosts metacognitive learning are still unknown. Future research on these mechanisms will lay a more solid foundation for designing interventions that boost metacognitive learning.

Overall, our findings suggest that systematic metacognitive reflection on how decisions were reached is a promising approach to improving human decision-making. This makes developing (digital) tools and interventions for fostering systematic metacognitive reflection in the real world (e.g., reflective learning chatbots) an important direction for future research. Fostering systematic metacognitive reflection could thereby become a valuable complement to conventional forms of systematic reflection that are already being applied to foster learning and improvement within organizations (Ellis et al., 2014; Wood Daudelin, 1996). Integrating these 2 forms of reflection and comparing their individual and combined effects is an interesting direction for future research. Moreover, systematic metacognitive reflection could also be highly beneficial to individuals seeking to improve their own thinking and decision-making. Furthermore, prompts for fostering systematic metacognitive reflection could also be used to develop educational interventions for helping students learn how to make better study choices (Azevedo, 2005; Bannert et al., 2009) and better decisions in real life (Heller et al., 2016; Hertwig and Grüne-Yanoff, 2017; Ryan and Ryan, 2013; Wilson and Jan, 1993).

Supplementary material. The supplementary material for this article can be found at https://doi.org/10.1017/jdm.2023.16.

Data availability statement. The code for the experiment and the statistical analysis is available under https://github.com/RationalityEnhancementGroup/ReflectiveLearning.

Funding statement. This work was supported by the MPI-IS Grassroots Program under Grant GR1040.

Competing interest. The authors declare none.

References

Ackerman, R., & Thompson, V. A. (2017). Meta-reasoning: Monitoring and control of thinking and reasoning. *Trends in Cognitive Sciences*, *21*(8), 607–617. https://doi.org/10.1016/j.tics.2017.05.004

Anseel, F., Lievens, F., & Schollaert, E. (2009). Reflection as a strategy to enhance task performance after feedback. *Organizational Behavior and Human Decision Processes*, *110*(1), 23–35. https://doi.org/10.1016/j.obhdp.2009.05.003

Azevedo, R. (2005). Computer environments as metacognitive tools for enhancing learning. *Educational Psychologist*, *40*(4), 193–197. https://doi.org/10.1207/s15326985ep4004_1

Azevedo, R., Landis, R. S., Feyzi-Behnagh, R., Duffy, M., Trevors, G., Harley, J. M., Bouchet, F., Burlison, J., Taub, M., Pacampara, N., Yeasin, M., Rahman, A. K. M. M., Tanveer, M. I., & Hossain, G. (2012). The effectiveness of pedagogical agents' prompting and feedback in facilitating co-adapted learning with metatutor. In S. A. Cerri, W. J. Clancey, G. Papadourakis, & K. Panourgia (Eds.), *Intelligent tutoring systems* (pp. 212–221). Springer: Berlin–Heidelberg.

Bannert, M., Hildebrand, M., & Mengelkamp, C. (2009). Effects of a metacognitive support device in learning environments. *Computers in Human Behavior*, *25*(4), 829–835. https://doi.org/10.1016/j.chb.2008.07.002

Bannert, M., & Reimann, P. (2012). Supporting self-regulated hypermedia learning through prompts. *Instructional Science*, *40*(1), 193–211. https://doi.org/10.1007/s11251-011-9167-4

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, *57*(1), 289–300. https://doi.org/10.1111/j.2517-6161.1995.tb02031.x

Bruine de Bruin, W., Parker, A. M., & Fischhoff, B. (2007). Individual differences in adult decision-making competence. *Journal of Personality and Social Psychology*, *92*, 938–956. https://doi.org/10.1037/0022-3514.92.5.938

Cacioppo, J. T., Petty, R. E., & Kao, C. F. (1984). The efficient assessment of need for cognition. *Journal of Personality Assessment*, *48*(3), 306–307. https://doi.org/10.1207/s15327752jpa4803_13

Callaway, F., Lieder, F., Krueger, P. M., & Griffiths, T. L. (2017). Mouselab-MDP: A new paradigm for tracing how people plan. In *The 3rd multidisciplinary conference on reinforcement learning and decision making*. https://doi.org/10.31219/osf.io/7wcya

Callaway, F., van Opheusden, B., Gul, S., Das, P., Krueger, P. M., Griffiths, T. L., & Lieder, F. (2022). Rational use of cognitive resources in human planning. *Nature Human Behaviour*, *6*(8), 1112–1125. https://doi.org/10.1038/s41562-022-01332-8

Ellis, S., Carette, B., Anseel, F., & Lievens, F. (2014). Systematic reflection: Implications for learning from failures and successes. *Current Directions in Psychological Science*, *23*(1), 67–72. https://doi.org/10.1177/0963721413504106

Ellis, S., & Davidi, I. (2005). After-event reviews: Drawing lessons from successful and failed experience. *Journal of Applied Psychology*, *90*(5), 857–871. https://doi.org/10.1037/0021-9010.90.5.857

Ellis, S., Mendel, R., & Nir, M. (2006). Learning from successful and failed experience: The moderating role of kind of after-event review. *Journal of Applied Psychology*, *91*(3), 669–680. https://doi.org/10.1037/0021-9010.91.3.669

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, *112*(4), 912–931. https://doi.org/10.1037/0033-295X.112.4.912

Ford, J., Schmitt, N., Schechtman, S. L., Hults, B. M., & Doherty, M. L. (1989). Process tracing methods: Contributions, problems, and neglected research questions. *Organizational Behavior and Human Decision Processes*, *43*(1), 75–117. https://doi.org/10.1016/0749-5978(89)90059-9

Goda, G. S., Levy, M., Manchester, C. F., Sojourner, A., & Tasoff, J. (2019). Predicting retirement savings using survey measures of exponential-growth bias and present bias. *Economic Inquiry*, *57*(3), 1636–1658. https://doi.org/10.1111/ecin.12792

Griffiths, T. L., Callaway, F., Chang, M. B., Grant, E., Krueger, P. M., & Lieder, F. (2019). Doing more with less: Meta-reasoning and meta-learning in humans and machines. *Current Opinion in Behavioral Sciences*, *29*, 24–30. https://doi.org/10.1016/j.cobeha.2019.01.005

He, R., Jain, Y. R., & Lieder, F. (2021). Measuring and modelling how people learn how to plan and how people adapt their planning strategies to the structure of the environment. In *International conference on cognitive modeling*. https://is.mpg.de/uploadsfile/attachment/attachment/671/20210720ICCMsubmissionfinal.pdf

He, R., & Lieder, F. (2022a). Learning-induced changes in people's planning strategies. Preprint. https://doi.org/10.13140/RG.2.2.12721.66407

He, R., & Lieder, F. (2022b). Where do adaptive planning strategies come from? Preprint. https://doi.org/10.13140/RG.2.2.28966.60487

Heller, S. B., Shah, A. K., Guryan, J., Ludwig, J., Mullainathan, S., & Pollack, H. A. (2016). Thinking, Fast and Slow? Some Field Experiments to Reduce Crime and Dropout in Chicago. *The Quarterly Journal of Economics*, *132*(1), 1–54. https://doi.org/10.1093/qje/qjw033

Hertwig, R., & Grüne-Yanoff, T. (2017). Nudging and boosting: Steering or empowering good decisions. *Perspectives on Psychological Science*, *12*(6), 973–986. https://doi.org/10.1177/1745691617702496

Hertwig, R., & Wulff, D. U. (2022). A description–experience framework of the psychology of risk. *Perspectives on Psychological Science*, *17*(3), 631–651.

Hilliger, I., De Laet, T., Henríquez, V., Guerra, J., Ortiz-Rojas, M., Zuñiga, M. Á., Baier, J., & Pérez-Sanagustín, M. (2020). For learners, with learners: Identifying indicators for an academic advising dashboard for students. In C. Alario-Hoyos, M. J. Rodríguez-Triana, M. Scheffel, I. Arnedillo-Sánchez, & S. M. Dennerlein (Eds.), *Addressing global challenges and quality education* (pp. 117–130). Cham: Springer International Publishing.

Irwin, J. R., & McClelland, G. H. (2001). Misleading heuristics and moderated multiple regression models. *Journal of Marketing Research*, *38*(1), 100–109. https://doi.org/10.1509/jmkr.38.1.100.18835

Jain, Y. R., Callaway, F., Griffiths, T. L., Dayan, P., He, R., Krueger, P. M., & Lieder, F. (2022). A computational process-tracing method for measuring people's planning strategies and how they change over time. *Behavior Research Methods*. https://doi.org/10.3758/s13428-022-01789-5

Jain, Y. R., Gupta, S., Rakesh, V., Dayan, P., Callaway, F., & Lieder, F. (2019). *How do people learn how to plan?* In *2019 conference on cognitive computational neuroscience* (pp. 826–829). https://doi.org/10.32470/CCN.2019.1313-0

Johnson, E. J., Camerer, C., Sen, S., & Rymon, T. (2002). Detecting failures of backward induction: Monitoring information search in sequential bargaining. *Journal of Economic Theory*, *104*(1), 16–47. https://doi.org/10.1006/jeth.2001.2850

https://doi.org/10.1017/jdm.2023.16 Published online by Cambridge University Press

Johnson, E. J., Schulte-Mecklenbeck, M., & Willemsen, M. C. (2008). Process models deserve process data: Comment on Brandstätter, Gigerenzer, and Hertwig (2006). *Psychological Review*, *115*, 263–272. https://doi.org/10.1037/0033-295X.115.1.263

Kinnier, R. T., & Metha, A. T. (1989). Regrets and priorities at three stages of life. *Counseling and Values*, *33*(3), 182–193. https://doi.org/10.1002/j.2161-007X.1989.tb00761.x

Kirschner, P., Sweller, J., & Clark, R. E. (2006). Why unguided learning does not work: An analysis of the failure of discovery learning, problem-based learning, experiential learning and inquiry-based learning. *Educational Psychologist*, *41*(2), 75–86. https://www.asec.purdue.edu/LCT/HBCU/documents/AnalysisoftheFailureofDiscoveryPBLExperientialInquiryLearning.pdf

Kocielnik, R., Xiao, L., Avrahami, D., & Hsieh, G. (2018). Reflection companion: A conversational system for engaging users in reflection on physical activity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, *2*(2), 1–26. https://doi.org/10.1145/3214273

Krueger, P. M., Lieder, F., & Griffiths, T. (2017). Enhancing metacognitive reinforcement learning using reward structures and feedback. In *Proceedings of the 39th annual conference of the cognitive science society*. https://doi.org/10.13140/RG.2.2.33793.58729

Larrick, R. P. (2004). Debiasing. In D. J. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (pp. 316–338). Malden, MA , Blackwell Publishing Ltd. https://doi.org/10.1002/9780470752937.ch16

Lejarraga, T., Ranganathan, K., & Wulff, D. U. (2022). Can simulated experience be harnessed to help people make investment decisions? PsyArXiv Preprint. https://doi.org/10.31234/osf.io/cw8f4

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

Loftus, G. R. (1978). On interpretation of interactions. *Memory & Cognition*, *6*(3), 312–319. https://doi.org/10.3758/BF03197461

Pammer-Schindler, V., & Prilla, M. (2021). The reflection object: An activity-theory informed concept for designing for reflection. *Interacting with Computers*, *33*(3), 295–310. https://doi.org/10.1093/iwc/iwab027

Panadero, E. (2017). A review of self-regulated learning: Six models and four directions for research. *Frontiers in Psychology*, *8*, 422. https://doi.org/10.3389/fpsyg.2017.00422

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*(3), 534–552. https://doi.org/10.1037/0278-7393.14.3.534

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge: Cambridge University Press.

Renner, B., Wesiak, G., Pammer-Schindler, V., Prilla, M., Müller, L., Morosini, D., Mora, S., Faltin, N., & Cress, U. (2020). Computer-supported reflective learning: How apps can foster reflection at work. *Behaviour & Information Technology*, *39*(2), 167–187. https://doi.org/10.1080/0144929X.2019.1595726

Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, *135*(2), 207–236. https://doi.org/10.1037/0096-3445.135.2.207

Ryan, M., & Ryan, M. (2013). Theorising a model for teaching and assessing reflective learning in higher education. *Higher Education Research & Development*, *32*(2), 244–257. https://doi.org/10.1080/07294360.2012.661704

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

Siegler, R., & Jenkins, E. A. (2014). *How children discover new strategies* (1st ed.). New York: Psychology Press. https://doi.org/10.4324/9781315807744

Siegler, R. S. (1999). Strategic development. *Trends in Cognitive Sciences*, *3*(11), 430–435. https://doi.org/10.1016/S1364-6613(99)01372-8

Simon, H. A., & Newell, A. (1971). Human problem solving. *American Psychologist*, *26*(2), 145–159. https://doi.org/10.1037/h0030806

Sitkin, S. B. (1992). Learning through failure: The strategy of small losses. *Research in Organizational Behavior*, *14*, 231–266. https://scholars.duke.edu/display/pub913886

Spiller, S. A., Fitzsimons, G. J., John G., Lynch, J., & Mcclelland, G. H. (2013). Spotlights, floodlights, and the magic number zero: Simple effects tests in moderated regression. *Journal of Marketing Research*, *50*(2), 277–288. https://doi.org/10.1509/jmr.12.0420

Veenman, M. V. J., Van Hout-Wolters, B. H. A. M., & Afflerbach, P. (2006). Metacognition and learning: Conceptual and methodological considerations. *Metacognition and Learning*, *1*(1), 3–14. https://doi.org/10.1007/s11409-006-6893-0

Wang, Y., & Sloan, F. A. (2018). Present bias and health. *Journal of Risk and Uncertainty*, *57*(2), 177–198. https://doi.org/10.1007/s11166-018-9289-z

Willemsen, M. C., & Johnson, E. J. (2011). Visiting the decision factory: Observing cognition with mouselabweb and other information acquisition methods. In M. Schulte-Mecklenbeck, A. Kuehberger, & J. G. Johnson (Eds.), *A handbook of process tracing methods for decision research* (pp. 21–42). New York: Psychology Press.

Wilson, J., & Jan, L. W. (1993). *Thinking for themselves: Developing strategies for reflective learning*. Portsmouth, NH, Heinemann.

Wolfbauer, I., Pammer-Schindler, V., & Rosé, C. (2020). Rebo junior: Analysis of dialogue structure quality for a reflection guidance chatbot. In *EC-TEL impact paper proceedings 2020: 15th European conference on technology enhanced learning*. https://par.nsf.gov/biblio/10174707

Wood Daudelin, M. (1996). Learning from experience through reflection. *Organizational Dynamics*, *24*(3), 36–48. https://doi.org/10.1016/S0090-2616(96)90004-2

Zimmerman, B. J. (1990). Self-regulated learning and academic achievement: An overview. *Educational Psychologist*, *25*(1), 3–17. https://doi.org/10.1207/s15326985ep2501_2

Cite this article: Becker, F., Wirzberger, M., Pammer-Schindler, V., Srinivas, S., & Lieder, F. (2023). Systematic metacognitive reflection helps people discover far-sighted decision strategies: A process-tracing experiment. *Judgment and Decision Making,* e15. https://doi.org/10.1017/jdm.2023.16