Judgment and Decision Making (2026), **vol. 21**:e20 1–28
doi:10.1017/jdm.2026.10043

EMPIRICAL ARTICLE

Adapting to information search costs in sample-based decisions

Linda McCaughey [1,2,3], Johannes Ziegler [4], and Klaus Fiedler[2,†]

[1]Psychology, TU Dresden, Germany; [2]Psychology, Heidelberg University, Germany; [3]Psychology, LMU Munich, Germany and [4]Psychology, Ulm University, Germany

Corresponding author: Linda McCaughey; Email: l.mccaughey@lmu.de

Received: 11 April 2023 UTC; **Revised:** 3 June 2026 UTC; **Accepted:** 16 June 2026 UTC

Keywords: information search; adaptivity; decision-making; metacognition; information costs

Abstract

When searching for information before making a decision, one faces a fundamental trade-off. On the one hand, more information will almost always improve the decision outcome. On the other hand, spending more time, money, or effort on further information search may be more costly than that improvement is beneficial. If decision makers are adaptive, they should use their resources in information searches efficiently. Using a sampling-based decision paradigm with financial costs for each observation, we conducted five experiments to investigate whether and how participants adapt their information amount to relevant aspects of the environment. Experiment 1a ($N = 67$) and 1b ($N = 79$) focused on changes in response to different ratios of information cost and decision payoffs, revealing that participants are capable of adapting, but seem to do so mainly based on a priori planning. Experiment 2 ($N = 131$) assessed the role of a second process, online adjustment, by varying whether participants received feedback. Surprisingly, participants changed their behavior to the same extent irrespective of whether they received feedback. To further improve the opportunities for metacognitive monitoring and control, Experiments 3a and 3b ($N = 160$ and $N = 318$) implemented different feedback types across more decision trials for three different ratios of information cost and decision outcomes, which stayed constant per participant. Whether participants changed their sample size and efficiency across trials depended on that ratio, but not on the type of feedback. This provides crucial insights into how participants adapt and what challenges might befall the necessary metacognitive monitoring, serving as a valuable prompt for future research.

1. Adapting to information costs in sample-based decisions

Searching for information before making a decision is fundamental to improving decision outcomes. The more information one has accumulated in favor of or against the options one is deciding between, the more likely one is to choose the most favorable one. This benefit of information, however, always comes at a cost—be it in terms of money, time, or effort. Hence, its benefits need to be weighed against its costs as part of a superordinate decision: is it worthwhile to keep searching, or is it better to make a decision based on the already acquired information? The answer will depend on the specific decision situation and its characteristics, for example: How consequential is the decision? How difficult is it to acquire more information? How helpful is additional information? Since humans are faced with many

[†] *Sadly, Klaus Fiedler passed away in May 2026, shortly before the final resubmission of this manuscript.*

© The Author(s), 2026. Published by Cambridge University Press on behalf of Society for Judgment and Decision Making and European Association for Decision Making. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-ShareAlike licence (https://creativecommons.org/licenses/by-sa/4.0), which permits re-use, distribution, and reproduction in any medium, provided the same Creative Commons licence is used to distribute the re-used or adapted article and the original article is properly cited.

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

different decision situations that vary along those and other dimensions, it is especially important to be able to flexibly adjust the amount of information one seeks out before making a decision.

Existing findings indicate that humans are, in principle, able to adapt their information search amount to relevant characteristics of decisions; for example, the decision outcome magnitudes (Hau et al., 2008; Pitz et al., 1969). However, it has rarely been examined *how* they are able to do so. We believe that a metacognitive perspective is necessary for such an investigation and will apply just that to explain the empirical results from a sample-based decision paradigm explicitly examining changes of information costs and decision payoffs. For this purpose, we propose two complementary metacognitive categories that govern information search adaptation, a priori *planning*, which happens before encountering a situation, and online *adjustment*, which happens during the experience with the situation.

The foundation for this investigation is naturally provided by previous research on how people conduct their information search. Careful examination of information search behavior started at least as early as the 1950s and 1960s, for example, with Edwards (1965) or Pitz et al. (1969), who suggested the first models of people's information search when drawing samples from one of two Bernoulli distributions with complementary probability parameters to decide which the underlying distribution was (bag and poker chip or urn paradigm). Participants had to pay a certain cost for each draw and received a gain for correct decisions and a deduction for incorrect decisions. Among other aspects, these costs and payoffs were varied across different conditions, demonstrating that participants responded sensibly (e.g., Busemeyer and Rapoport, 1988; Irwin and Smith, 1957; Pitz, 1968). People have also been shown to respond to other characteristics of decisions, for example, the variance among decision options and the similarity of cue validities in attribute-based decisions (Payne et al., 1988), or the distribution of high and low options in the environment in optimal stopping problems (Guan and Lee, 2018).

Reviving sample-based decision paradigms, more recent experiments have examined how people behave when they choose between gambles after acquiring information about the outcome distributions of these gambles by sampling from them at no financial cost (Hertwig et al., 2004). Experiencing the gamble's outcome at no cost before making a final consequential draw is a good proxy for information search and the findings suggest that people display considerable adaptivity (Wulff et al., 2018): They search until they (believe to) have seen all distinct outcomes; (Spektor and Wulff, 2024; Srivastava et al., 2016), they tend to search for longer before making their consequential draw when the stakes are higher (Hau et al., 2008), and they radically curtail information search in presence of a rival who can pre-empt their decision (McCaughey et al., 2025; Phillips et al., 2014). This is a very good indication that humans can select their information search amount guided by the characteristics of the decision, at least under certain circumstances and to some extent.

Most previous research implemented changes of decision characteristics, such as outcomes and their probabilities, exclusively between experimental conditions. Hence, for the majority of examined tasks, the decision characteristics remain stationary over the course of the period under investigation. There are a few exceptions, however, with Rakow and Miler (2009) having looked at gradual changes of outcome probabilities in lotteries from which participants made consequential draws. They concluded that many behaviors that seem irrational in stationary lotteries may be exhibited by participants because they help them adapt to gradual changes in non-stationary environments, which they more commonly encounter. In a very similar task, Navarro et al. (2016) showed that participants changed their search behavior depending on whether they were in an environment with stationary or gradually changing characteristics. While these types of gradually changing environments certainly exist, we are interested in another type of change that may be even more frequent. Moving through life, humans encounter different decision situations fully aware that they are entering a specific situation or even deliberately choosing it. We want to call a specific decision situation like this a 'decision context' to refer to all the characteristics that influence how profitable different amounts of information are. Most prominently, this encompasses the decision payoff, that is, the potential or actual outcomes that a decision results in, and the information costs, but many other characteristics are also entailed by it, such as the number

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

of options and their relation. Hence, the type of decision context change we are interested in is characterized by a sudden concrete change of some characteristics from one stable period to another.

This is precisely what Lee et al. (2014) examined: Rather than investigating a gradually changing context, they implemented one concrete change and a subsequent change back to the initial context characteristics, which were stable otherwise. Their experiment was set up such that participants decided which of two options was the better by sampling attributes. In the initial decision context, very little information sufficed for high accuracy rates. After a covert change to the second context, more information was needed for correct choices, while a second covert change back to the initial context made extensive search no longer necessary. Participants in this experiment were good at responding to the unannounced changes, which they could only learn about through their own choices and the feedback they received. The fact that most reduced their information search amount when less information sufficed was especially intriguing because simple model-free reinforcement learning accounts based on trial-by-trial error updating cannot explain this pattern. Although more sophisticated model-based reinforcement learning models developed since may be able to account for such patterns, we believe a metacognitive framework is most suitable to characterize the processes in a psychologically plausible way (as argued by Chater, 2009).

Lee et al.'s (2014) paradigm focused on covert changes that participants had to detect based on their experience with the task, which is an important aspect of information search. It could not, however, capture a second important aspect of adapting information search, namely, adapting to decision contexts that one already has knowledge about. Since humans typically have experience with many of the decision contexts they face or have other ways of obtaining information about them, they often enter a decision context with knowledge about it, which allows for additional adaptation processes. For example, knowledge of an upcoming change and its likely consequences may allow for mental preparation, for example, by retrieving relevant memories or mentally simulating the situation. The present paper does not aim to examine specific preparatory processes; it focuses instead on their consequences for adaptive behavior, introducing a metacognitive framework that defines two forms of adaptation, a priori planning and online adjustment.

1.1. Specifying metacognitive processes involved in adaptation

While metacognition is often discussed in terms of monitoring and control processes (Nelson and Narens, 1990), having knowledge of an upcoming change and what it entails puts the spotlight on another, more specific metacognitive process: planning. While planning does at least play a secondary role in frameworks of meta-memory (Nelson and Narens, 1990) or self-regulated learning (Panadero, 2017; Schraw and Moshman, 1995), it has not been featured prominently in information search and decision-making literature (except in the related field of model-based reinforcement learning, in which planning is made possible through the model of the world; Dayan and Niv, 2008). In line with the literature discussed above, we define planning as the decision about a future course of action, which often constitutes a strategy for the task at hand. In both learning and memory, an important aspect of this strategy is how much study time to allocate (e.g., Nelson and Narens, 1990). In our paradigm, planning involves the formulation of a strategy, that is, a series of steps to search for information and make a decision. Since participants receive information about the task and all task changes before encountering them, they are able to engage in planning a priori.

As soon as the task is being encountered—be it for the first time or after a change—we assume that processes belonging to the classic monitoring and control categories of metacognition take place. While the control category entails all processes for executing a strategy, be it maintained or adapted, the monitoring category contains interesting further distinctions in Nelson and Narens' classic conceptualization (1990). First and foremost, monitoring involves the observation of various aspects of the task execution and its results at different stages, including, for example, ease-of-learning judgements or judgements of learning in memory tasks. Applying this conceptualization to our paradigm, the explicit performance feedback that is provided is the most obvious and will be the focus of monitoring.

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

Based on these monitoring processes, people evaluate their current strategy as either requiring change or not. In case a change is decided upon, a planning stage that is in principle equivalent to the initial one follows. The result of that plan is then implemented through control processes, after which the monitoring stage starts anew since monitoring and control processes are applied in a continuous loop.

Different levels of monitoring. Our assumption is that monitoring of task aspects and performance can extend to different temporal horizons. At the shortest temporal horizon, monitoring is based only on the current or most recent decision (which is the current or most recent decision trial in the present paradigm) and serves as the primary basis of evaluating the current strategy and adjusting it if necessary. To draw the comparison with reinforcement learning again, this is similar to adjusting parameters in one's model based on feedback rather than changing the model's structure (Gershman and Niv, 2010; Sutton and Barto, 2018). In the cases when monitoring is based on more than the immediate experience, the temporal horizon can be as far as the cognitive and motivational capacities of the decision maker allow for. A natural implication of this is that the temporal horizon chosen for monitoring and evaluation will be influenced by the support and guidance provided in the environment.

The level of monitoring also marks the subtle distinction between the planning stage before starting a task and the planning during task execution. While planning during a task can be based on monitoring at any point along the temporal horizon, from focusing on the most recent decision to integrating extended experience, the planning stage before encountering a task is restricted to abstract information derived from the description of the task, possibly supplemented by memories of similar tasks and mental simulations thereof.

Before and during a task. We want to examine how these processes contribute to adaptation in two different stages of time, before experiencing a new task or task change and while experiencing a new task or task change. While planning can occur in both stages, we will focus on the planning that occurs before experiencing the task and will refer to it as a priori planning. This is pragmatically a very interesting subset of planning processes because people often do have knowledge of decision contexts before they are faced with them. At the same time, a priori planning has received very little attention. The processes that occur during the experience with the task will be termed online adjustment and will encompass monitoring and evaluation processes as the main interest but also planning and control processes as necessary components. We designed a set of experiments based on a sample-based decision paradigm specifically tailored to examine these different processes at work.

1.2. Present paradigm

We used a classic paradigm of sampling from Bernoulli distributions (see Pitz et al., 1969), in which participants can determine the number of draws. After sampling as many binary observations as they wish, they make a rewarded decision about whether the underlying distribution is more likely to generate one or the other event ($p < .5$ or $p > .5$). The information is costly, so participants have to weigh this cost against the benefit of increasing the probability of making a correct decision, which depends on the reward policy. Overt context changes are operationalized as changes in the reward policy, that is, the ratio of information costs and decision payoff in the experiments reported below.

The normative solution to this task is to update one's belief in favor of one or the other decision option by computing the Bayesian odds ratio with each new observation that is sampled (Edwards, 1965; Lee et al., 2014). The ratio quantifies the relative support for one hypothesis over the other (i.e., $p < .5$ vs. $p > .5$) based on the likelihood of the observed sample. Once a critical threshold is exceeded, indicating that the subjective belief favors one alternative strongly enough over the other, a decision is made. The reward-maximizing threshold depends on the information costs and payoff: Higher information costs, for example, imply the need to stop earlier to maximize the total reward (Edwards, 1965). People deviate from this norm in various ways (Busemeyer and Rapoport, 1988; Pitz, 1968), but it is nevertheless useful as a standard of comparison and as a conceptual model to further illustrate the different meta-cognitive adaptation processes.

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

From the perspective of the conceptual Bayesian model, adapting to these changes in a way that maximizes reward means adjusting the odds ratio threshold in accordance with the ratio of information costs to decision payoff. Different threshold settings are known to be implemented by humans to produce behavior to meet environmental demands, such as different instructions in speed versus accuracy conditions (Ratcliff and McKoon, 2008) or to achieve a good reward rate (Bogacz et al., 2010). Hence, they serve as a good example of how metacognitive processes of the two categories could be applied. If the optimal thresholds were known beforehand, they could simply be adjusted as soon as the environment changed. Assuming that the optimal threshold is not known and cannot be derived a priori, metacognitive adaptation processes need to be applied. First, when a change is announced, the decision maker might think about which threshold is sensible and plan to implement it in the new context. This intention might be based on the general (correct) idea that when information becomes cheaper relative to the decision outcome, it makes sense to acquire more of it, while it becoming more expensive should lead one to acquire less. This could lead to the plan to acquire more or less information than previously as part of the (a priori) planning process.

Once one is searching and making decisions in the new context, what we term online adjustment processes are at work. Based on the experience in the new context, one can further adapt the threshold step by step. As Lee et al. (2014) suggest, feedback plays a vital role in this process, and metacognitive monitoring and evaluation are crucial for adjustment to lead to improvement. Evaluation will further determine whether some aspect of the present strategy needs to be changed. Through planning and control processes, a change can be implemented, which will again be monitored and evaluated. The way we conceptualize the task is that the metacognitive level encompasses considering whether the employed strategy and threshold are working well or not for one or multiple trials after a decision is made. There are also theories that treat terminating the search itself as a metacognitive process because one has to evaluate whether one is confident enough or has enough evidence to make a decision (Ackerman et al., 2020). Either way, we assume a superordinate metacognitive level at which the processes we have detailed take place.

To test to what extent a priori planning and online adjustment processes guide adaptation to new decision contexts, we devised three experiments in which the information-cost ratio changed from block to block with prior announcement. The first experiment tested whether people adapted to overt changes in ratio across four blocks, focusing on a priori planning processes. The second experiment increased the number of decision trials and manipulated feedback to examine online adjustment processes across only two changes more closely. Finally, in a third experiment, the focus was exclusively on online adjustment, more specifically on the level of monitoring, with feedback type and ratios varying only between participants to give individuals the opportunity and motivation to implement further adjustment processes.

2. Experiment 1 (a and b)

To establish whether participants were sensitive to changes in the information cost and payoff parameters and to what extent, in Experiment 1a, we varied either the information cost or the payoff across four blocks. The ratio of those determines whether it is worthwhile to acquire larger amounts of information or whether an early decision is more profitable. Since the information amount was the only aspect that participants could control, the ratio should be reflected in participants' sample sizes. Hence, we expected participants to display sensitivity through changes in sample size governed mainly by a priori planning processes, while adjustment processes could only play a smaller role due to the fairly small number of trials per block. While participants' sample size is the most immediate indicator of strategy changes, it can also be reflected in accuracy, albeit afflicted by more extraneous factors, such as practice and attention. Efficiency is the joint result of sample size and accuracy, defined as the sum of the decision payoff (depending on the decision's accuracy) and the total information cost (dependent on the sample size) at the level of the trial. As such, it is an interesting indicator but will again reflect behavioral changes less reliably.

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

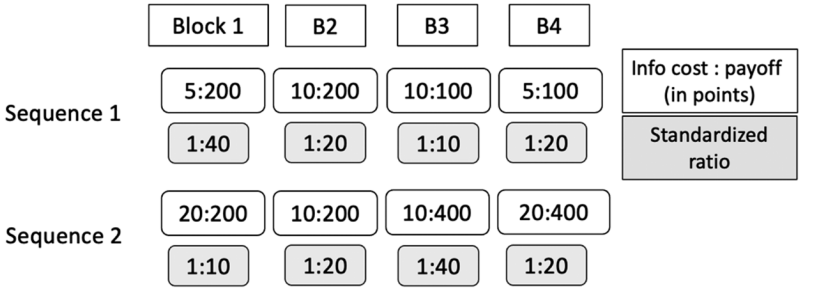


Figure 1. *Illustration of the two sequences of ratios for Blocks 1–4.*
Note: Boxes with a white background show the ratio of information cost and payoff instantiated by the amount of points used in the experiment; boxes with a gray background show the standardized ratio.

While we expected participants to be sensitive to the ratio changes, their absolute efficiency levels may not necessarily be high. The a priori planning process allows for sample size modifications in the direction that seems reasonable based on the announced changes. However, the efficiency after the change depends heavily on the efficiency of the previous behavior. Reliable improvements in absolute efficiency can only be driven by adaptation based on feedback, for which there was probably not enough opportunity in this experiment.

2.1. Methods

2.1.1. Participants

We recruited 78 participants from our local participant pool, of whom 76% were female, 23% male, and one participant preferred not to say. They were aged 19–77, with a mean of 26.3 (*SD* = 11.3), with a vast majority of students (11 of which were psychology students; 6 were not students). Because this was a new paradigm, we recruited the number of participants that could be recruited in the lab-typical recruitment phase of one week. The experiment was part of a one-hour session in our laboratory and was conducted as the second out of four studies in that session. The other three studies pertained to directed forgetting, impression formation, and encoding of information. In addition to either a fixed amount of 5€ or course credit, participants received a variable bonus of, on average, 3€ that ranged from 0€ to 8€ and depended on their performance at the experiment described here. The bonus was determined by converting participants' points to Euros according to a conversion rate that was estimated to yield the reported mean. The session was advertised as yielding 5€ and an additional average bonus payment of 3€, and participants were informed that this bonus payment was dependent on their performance at the experiment described here.

2.1.2. Design, task, and procedure

The experiment had a within-participant design. Each participant faced three different information-cost–payoff ratios realized in four different blocks of 25 decision trials each. This cost–payoff ratio remained unchanged for the 25 trials within a block but was changed between blocks after an announcement of the next block's information cost and payoff amounts. Participants started the first block with ratios 1:40 or 1:10 (cost per information: potential gain or loss), while ratios of the subsequent blocks were created by either doubling or halving one of the ratio components (see Figure 1). To avoid order effects and effects of the concrete number of points that instantiated the ratios, we partially counterbalanced the order of the ratios and the specific number of points that were used to implement them between subjects. Figure 1 shows the order of the ratios and what specific information costs and payoffs were used to instantiate them.

Task. In each trial, participants made a decision about the probability parameter p of a Bernoulli distribution from which they could acquire as many randomly sampled observations as they thought

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

was sensible. The cover story was that they had to decide whether a slot machine tended to yield more gains or more losses overall, where the individual observations were defined as past gains and losses drawn at random from the totality of past recorded gains and losses.

The observations came at a certain cost, which will be referred to as information cost, while the decision resulted in the gain or loss of a certain sum, which will be called the payoff, depending on whether it was correct or incorrect. The screen displayed the trial number and block number out of the total trials and blocks on the top right and the general information cost and decision payoff for the current trial (and block) on the top left. On the left side of the screen, participants saw their own current total balance, which was updated after each trial by subtracting the points incurred for information and adding or subtracting the points from the decision payoff. Participants could acquire information by pressing the space bar, which made an arrow appear at a random location in an 8×5 grid. The arrow was either green and upward-pointing to represent a positive observation or magenta and downward-pointing to represent a negative observation. The maximum sample size possible was 40. As soon as participants wanted to make a decision (or after the maximum sample size was reached), they pressed the upward arrow key or the downward arrow key to indicate their decision that the population was predominantly positive or negative, respectively. Participants' current balance was not updated during the trial. Instead, a separate counter indicating how much they had spent on information in the current trial was updated with each arrow. The probability of a positive observation p of the Bernoulli distribution was drawn at random from the values {.1, .2, .3, .4, .6, .7, .8, .9} for each new trial. Participants received feedback about whether their decision was correct and the resulting decision payoff for that trial in a textbox superimposed over the sampling and decision screen for 1,200 ms. After this, the next decision was loaded, where their updated current total balance was displayed. The experiment was programmed in Python using Pygame.

2.2. Results

2.2.1. Analysis procedure

The reported analyses for all experiments were carried out using R (R Core Team, 2025) and RStudio (Posit team, 2025). Linear mixed models were fitted using the afex package (Singmann et al., 2024) and in accordance with the guidelines authored by Singmann and Kellen (2019), specifically starting with the maximum possible random effects structure and omitting random effects that did not explain variance until the model converged without single fit warnings. Whenever the full random effects structure could not be used, the random slopes that could be included will be reported. A random intercept varying by participant could always be included and will not be reiterated. P-values were determined using the Satterthwaite approximation for degrees of freedom unless otherwise specified. The experimental data are openly available in the Zenodo repository: https://doi.org/10.5281/zenodo.20719463.

2.2.2. Exclusions

Four participants had to be excluded because the experimental program parameters were not set correctly. Three participants were excluded because their accuracy rate did not differ from guessing, as indicated by binomial tests. Four additional participants were excluded because their mean sample size was below 0.75, meaning that they acquired 0 observations in at least ¼ of the trials, which leaves too little data to reliably assess their behavior. None had to be excluded due to slow response times, leaving a total sample size of 67. Sequence 1 had 35 participants and Sequence 2 had 32.

2.2.3. Trade-Off

To check whether this task generally resulted in a trade-off between the costs and benefits of information, we plotted participants' mean accuracy as well as their mean efficiency as a function of their mean sample size (Figure 2). Efficiency is not independent of the ratio, so to equate it across the different ratios for this plot, it was recalculated by scaling all information costs to 10 and all payoffs to

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

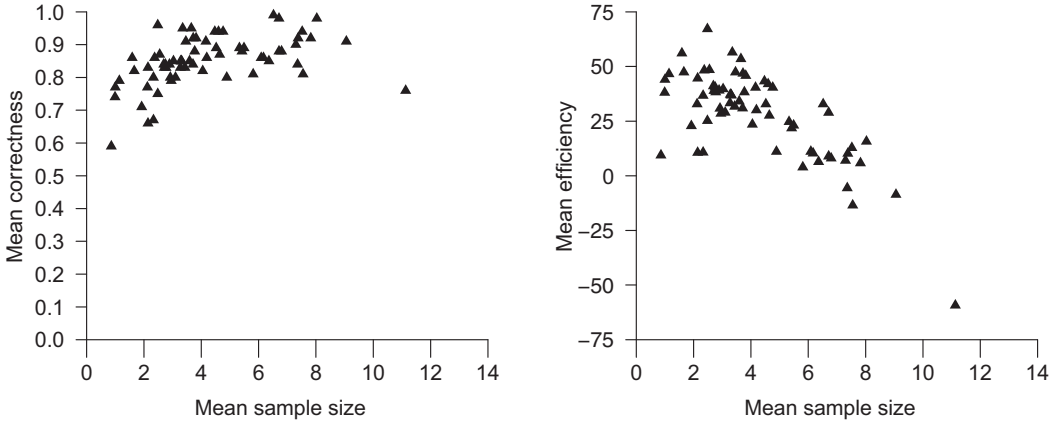


Figure 2. *Participants' mean accuracy and efficiency as a function of sample size.*
Note: Accuracy, efficiency, and sample size were averaged across all blocks per participant, resulting in one value each for each participant.

100 and recalculating the balance on that basis to make all trials equivalent. The accuracy plot in the left panel of Figure 2 clearly illustrates the steep initial increase in accuracy with the first few observations and the flattening increase with additional observations (Pearson correlation coefficient $r = 0.46$, $t(65) = 4.16$, $p < .001$). Combined with the linear information cost increase, this leads to the steep initial increase in efficiency in the right plot and the steep subsequent decrease when the cost of information starts to outweigh the increase in probability that the decision will be correct (resulting in a negative Pearson correlation coefficient $r = -.74$, $t(65) = -8.75$, $p < .001$)

2.2.4. Sample size analysis

To assess the extent to which participants adapted their strategy to the changing ratios, we analyzed the change in sample size both across the different blocks (more specifically, across ratios) and across trials within the blocks. A priori planning processes would be evident when a change in sample size was implemented from the very start of the new block, for example, when ratio or block would show an influence on sample size. Adaptation processes, on the other hand, would be evident in the influence of trial number on sample size. Sample size changes are the most sensitive measure because sample size is directly controlled by participants. We expected participants to increase their sample size if the information cost decreased or the payoff increased and vice versa. We will refer to this as sample size changes in accordance with the ratio change. From participants' perspective, without prior experience with the task and without sufficient experience to know what maximizes efficiency, this would be the most sensible thing to do.

Figure 3 illustrates participants' mean sample sizes for Blocks 1–4, each block subdivided into sections that approximate block thirds, and two separate lines representing the two different sequences of ratios. The graph clearly shows that the change in sample size between the blocks is large, while the changes within the blocks are minimal. Since the different order merely served as a partial counterbalancing, the data were pooled according to the ratio (independent of block and condition) in the analysis to examine the influence of the ratio. The means and standard deviations pooled by ratio are displayed in Table 1. We estimated a linear mixed model with sample size as the dependent variable and ratio, trial, *p*'s deviation from .5, and their interaction as fixed effects and a random intercept, random slopes for ratio, *p*'s deviation[1], and trial, and their two-way interactions as uncorrelated effects

[1]The larger the deviation of the distribution's *p* from .5, the closer it is to .9 and .1, which are *p* parameters that yield the most evidence in favor of a decision option controlling for sample size. Conversely, parameters closer to .5 will tend to contain more mixed samples and less evidence. Including this feature of the trial in the model allows to test whether participants are sensitive to the amount of evidence in the sample by requiring less observations when the distribution's *p* is further from .5.

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

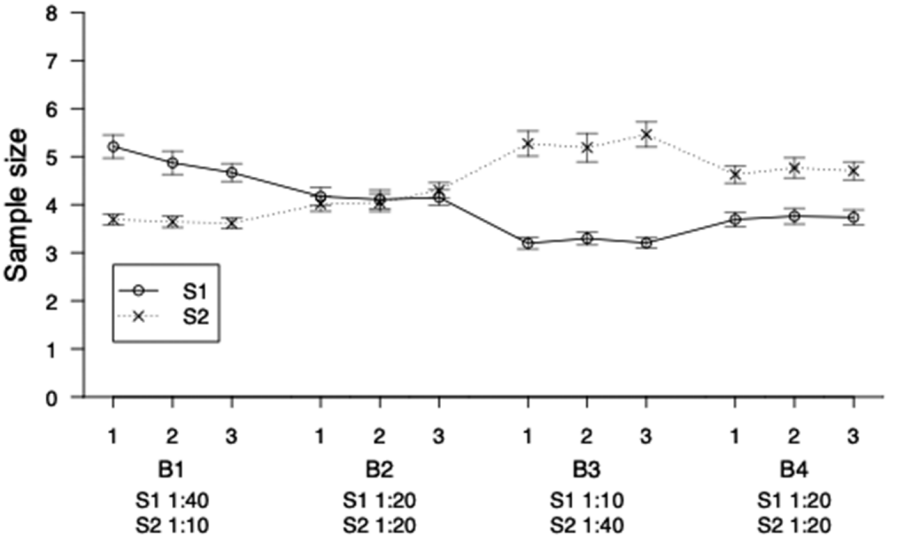


Figure 3. *Sample size across block sections separately for the two ratio sequences.*
Note: Error bars represent the standard error of the mean. Block is abbreviated B. Condition is abbreviated C. The standardized ratio for sample size was aggregated for block sections of 8 or 9 trials and, hence, approximated block thirds.

Table 1. *Descriptive statistics of sample size and accuracy.*

Ratio	*n*(*SD*)	*a*(*SD*)
1:10	3.43 (1.54)	.827 (.085)
1:20	4.16 (2.17)	.855 (.078)
1:40	5.11 (3.10)	.867 (.088)

Note: Mean sample size *n* and accuracy *a*, with standard deviation (*SD*) in brackets, separately for ratios.

varying by participant (ratio coded as −1 for 1:10, 0 for 1:20, and 1 for 1:40, and trial centered on 0, and *p*'s deviation coded as −1.5, −0.5, 0.5, 1.5). The results showed that participants indeed responded to the ratio changes by implementing larger sample sizes when information cost was cheaper, or payoff was higher, with sample size increasing by 0.82 with every coded ratio increase ($b = 0.82$, $SE = 0.12$, $F(1, 65.27) = 43.93$, $p < .001$). The deviation of *p* from .5 also had an influence on sample size, such that the larger *p*'s deviation from .5 and the closer it was to the extremes of 0 and 1, the smaller participants' sample size was ($b = -0.44$, $SE = 0.06$, $F(1, 62.99) = 49.93$, $p < .001$). There was also an interaction between ratio and *p*'s deviation, indicating that the cheaper the information, the larger the difference in sample size for the different levels of *p*'s deviation ($b = -0.21$, $SE = 0.04$, $F(1, 58.42) = 26.83$, $p < .001$). The trial number had no effect, and there were no other significant predictors (see the Supplementary Material for full table).

Hence, participants changed the sample size that they used, on average, to suit the ratio they encountered. The main timing of this change seemed to have been before the start of each block, since we found no evidence of trial influencing the sample size. This is not surprising given the small number of trials and the varying *p*, which makes the decision task quite variable, both of which do not facilitate adjustment processes.

2.2.5. Accuracy analysis

To assess whether the changes in sample size were associated with changes in accuracy as part of the trade-off, we also examined accuracy with a binomial generalized linear mixed model with ratio, trial,

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

Table 2. *Fixed-effect estimates and tests for sample size analysis of Experiment 1b.*

Fixed effect	Estimate	*SE*	*df*	*F*	*p*
(Intercept)	4.17	0.20			
Ratio	**0.42**	**0.07**	**1, 77.73**	**37.07**	**<.001**
Trial	**0.01**	**0.00**	**1, 77.76**	**4.03**	**.048**
***p*'s deviation**	**−0.49**	**0.06**	**1, 78.14**	**63.90**	**<.001**
Ratio × Trial	0.00	0.00	1, 77.66	1.02	0.315
Ratio × *p*'s deviation	**−0.11**	**0.03**	**1, 77.83**	**19.63**	**<.001**
Trial × *p*'s deviation	0.00	0.00	1, 79.73	0.64	.426
Ratio × Trial × *p*'s deviation	0.00	0.00	1, 75.55	1.41	.239

Note: In the column names, *SE* denotes the standard error of the estimate, *df* the degrees of freedom for the *F*-test, *F* the test-statistic value, and *p* the *p* value. Fixed effect labels in bold indicate significant results ($p < .05$).

and *p*'s deviation from .5 as fixed effects and a random intercept and random slopes for trial and *p*'s deviation varying by participant as uncorrelated random effects. P-values were determined using the likelihood ratio test within the afex package. The results showed that each increase in ratio, which made information less expensive, led to an increase in the odds of making a correct decision by about 25% ($b = 0.22$, $SE = 0.06$, $\chi^2(1) = 12.1$, $p < .001$), consistent with the mean accuracy rates of .827, .855, and .867 for ratios 1:10, 1:20 and 1:40. Also, with every .1 increase in distance from $p = .5$, the odds of being correct increased by about 140%. There was also a significant interaction between trial and *p*'s deviation from $p = .5$, indicating that with each trial increase, the positive effect of deviation on accuracy was strengthened by 0.02 ($b = 0.02$, $SE = 0.01$, $\chi^2(1) = 15.72$, $p < .001$). So the difference in accuracy rates for the different deviations was more pronounced for later trials.

2.3. Discussion

All in all, participants were clearly sensitive to the ratios and changed their search behavior in a way that is reasonable based on limited knowledge and experience with the task. This is an important aspect of responding suitably to changes in one's environment or when moving between contexts and belongs to the category of a priori planning. Online adjustment processes did not play a noticeable role, as there was no evidence of sample size adaptation within the blocks. This might be due to the rather small number of trials per block (which is why efficiency was not tested) and will be investigated more closely in the experiments reported below.

3. Experiment 1b

Experiment 1b ($N = 67$) was very similar to Experiment 1a and is therefore reported in the Supplementary Material and only briefly summarized here. An additional ratio (1:80) allowed for the expansion of the conclusion to a wider range of values. The other crucial difference was that the comparison of the ratio was made more difficult through an additional task between the blocks and the fact that both ratio components changed between blocks instead of just one. This was intended to make it harder to derive the appropriate sample size change a priori, before experiencing the block, making more room for potential change during the block. And indeed, a linear mixed model showed that in addition to the main effects of ratio and *p*'s deviation on sample size and accuracy, there was a slight main effect of trial on sample size, however, without ratio interactions (see Table 2 and Supplementary Material for full report). To be able to better investigate the online adjustment processes, we conducted a further experiment better suited to this purpose.

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

4. Experiment 2

To facilitate online adjustment processes, in Experiment 2, we made the task less variable by reducing the possible levels the probability parameter could take on and by increasing the number of trials per block to give participants more time to experience the new ratio and change their behavior in response. For effective assessment of the extent of online adjustment exhibited by participants, we implemented a control group that did not receive any feedback after the first change. We expected participants who received feedback to change their sample size more than participants in the control condition, specifically during Blocks 2 and 3, because the online adjustment processes would be impaired by the lack of feedback. This greater change in sample size should also lead to higher efficiency, as online adaptation based on feedback allows for improvement.

4.1. Methods

4.1.1. Participants

We recruited 148 participants from the study platform Prolific for an online experiment. Since the design included two between-participant factors, we aimed to recruit about twice as many participants as in the previous experiments. They had to speak fluent English and currently reside in the UK to be able to participate to ensure they understood the English instructions. Their mean age was 39.9 (SD = 13.9), ranging from 19 to 78. Of all participants, 69.9% reported their sex as female, 27.7% as male, and 2.7% did not make their data available. About 7.4% indicated they were students, 60.1% were not, and for 31.8%, data on student status were unavailable. The experiment took about 15 minutes and was remunerated with a base remuneration of £1.50 and an additional bonus payment of about £1.25 on average. In the study announcement on Prolific, potential participants were informed that they could earn £1.50 and a bonus payment of up to £2.50.

4.1.2. Design and procedure

The experiment had a 2×2-between-participant design, with feedback (present or absent after Block 1) and ratio change (via payoff or information cost) as orthogonal factors. As an additional within-participant factor, we implemented three blocks with two different information-cost—payoff ratios. All conditions received feedback in the first block to be able to familiarize themselves with the task in the same way. In Blocks 2 and 3, only the feedback conditions continued to receive feedback. As displayed in Figure 4, the sequence of ratios, from 1:10 to 1:40 back to 1:10, was the same for all participants, but the points at which ratio 1:40 was instantiated in the second block differed. The information-cost-change conditions went through a change of information cost to a ratio of 5:200 from 20:200, while the payoff-change-conditions faced a ratio of 20:800 after the initial ratio of 20:200. Instead of being separated by just one doubling or halving of the payoff or information costs as in the previous experiments, these ratios were separated by two halvings or doublings to make the necessity for change even more salient.

The blocks consisted of 40 trials to give participants more time to adjust to each new ratio. The task was similar to the other experiments, except that the cover story was concerned with products and whether their reviews were predominantly positive or negative, and the arrows did not appear at a random location of the rectangle but started at the top left and appeared in line, with 8 arrows in each line. A smaller maximum sample size of 16 was chosen since it is already substantially larger than the highly efficient strategies and hence allowed for the assessment of oversampling while at the same time limiting the time participants could spend on individual decisions. In addition, the Bernoulli distributions' p no longer varied but was either .3 or its complement .7, making the task easier because it is less variable. Participants indicated their decision by clicking one of the two buttons below the arrow box on the screen to indicate their chosen decision option. Whenever feedback was provided, the screen following the decision displayed the correctness of the decision, the number of points gained or lost, and the new balance for 1,500 ms before automatically displaying the next screen for the new

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

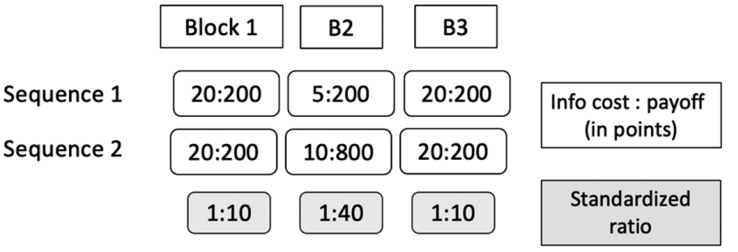


Figure 4. *Illustration of the two ratio sequences for Blocks 1–3 in Experiment 2.*
Note: Boxes with a white background show the ratio of information cost and payoff instantiated by the amount of points used in the experiment; boxes with a gray background show the standardized ratio.

decision. When no feedback was provided, the screen after the decision screen instructed participants to wait for the next trial to start, to which it was automatically redirected after 1,500 ms. In addition, the balance was not displayed when feedback was withheld, because the updated balance would have acted as a feedback substitute.

4.2. Results

4.2.1. Exclusions

Fifteen participants were excluded because their mean sample size was below 0.75, which did not leave enough data to reliably assess their behavior. Two had to be excluded because they took a break longer than 3 minutes, leaving a total sample size of 131 for analysis.

4.2.2. Sample size analysis

Since the ratio changes across the three blocks consisted of an increase and an equivalent decrease back to the original ratio, we decided to run two separate linear mixed models, one comparing the first and second blocks, and another comparing the second and third blocks. To correct for the inflated probability of Type-1 errors, we applied a Bonferroni correction by dividing our alpha level by the number of analyses, resulting in .05/2 = .025, which was used as the alpha level in the analyses below.

As fixed effects, we included ratio, trial, feedback, and type of ratio change (all predictors centered on 0, with the following coding: 1:10 as −1 and 1:40 as 1, feedback absence as −1 and feedback presence as 1, and information-cost change as −1 and payoff change as 1; trial was rescaled by dividing it by 10 to make it more similar in range to the other predictors, facilitating convergence). Random slopes varying by participant were included for ratio, trial, and their interaction, as correlated effects were used.

Figure 5 visualizes the sample size across quarter blocks, aggregated across 10 trials per quarter, separately for conditions, while Table 3 displays the means and *SDs* of sample size for the feedback conditions along with accuracy and efficiency. As is clearly visible in the graph and supported by the analyses detailed in Tables 4 and 5, sample size varied with ratio and trial. Participants tended to use larger sample sizes for the higher ratio in Block 2 compared to both Block 1 and Block 3 (b_{12} = 0.85, b_{32} = 0.49; see Tables 4 and 5 for further details). The higher the trial number, the larger the sample size that participants implemented in all blocks (b_{12} = 0.24, b_{32} = 0.18), clearly indicating that adjustment processes were applied.[2] The influence of the trial number was moderated by the ratio such that the increase in sample size with each trial was higher for the higher ratio, that is, higher in Block 2 compared to both Block 1 and Block 3 (b_{12} = 0.10, b_{32} = 0.16). The sample size change across trials in Block 2 brought participants closer to the sample size of an optimal strategy (see section on Optimal Strategy and Table 6), indicating that the adjustment processes resulted in a change with reasonable

[2]Since trial was rescaled, dividing it by 10, the parameters indicate the change per 10 trials.

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

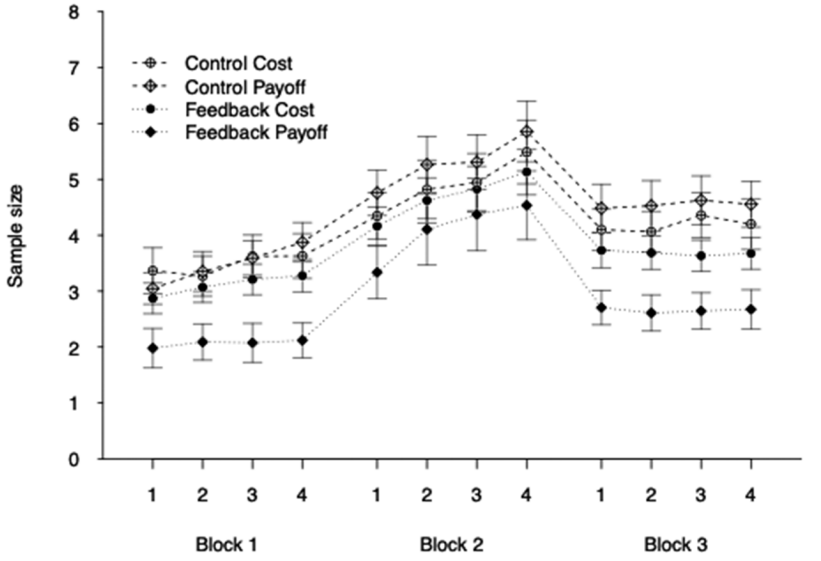


Figure 5. *Sample size displayed by block quarters.*
Note: Each data point is aggregated across one quarter of a block, i.e., 10 trials.

Table 3. *Descriptive statistics by feedback condition and block.*

Feedback condition	Block	Mean *n* (*SD*)	Mean *a* (*SD*)	Mean *Eff* stand. (*SD*)
Control	1	3.47 (1.81)	0.77 (0.11)	2 (1.77)
	2	5.14 (2.78)	0.84 (0.10)	22.23 (6.33)
	3	4.4 (2.4)	0.82 (0.09)	1.97 (2.04)
Feedback	1	2.79 (1.7)	0.75 (0.12)	2.21 (1.9)
	2	4.5 (2.51)	0.81 (0.12)	20.21 (7.72)
	3	3.37 (1.81)	0.79 (0.12)	2.44 (1.75)

Note: n = sample size; *a* = accuracy defined as proportion of correct decisions; mean *Eff.* stand. is the standardized decision payoff minus the standardized information costs, with standardizing to an information cost of 1, resulting in a standardized payoff of 10 for the first and third blocks and of 40 for the middle block.

direction. In Blocks 1 and 3 (ratio 1:10), participants' sample sizes were closer to but larger than those of an optimal strategy to begin with. Hence, although sample sizes should have decreased to maximize gain, a smaller increase than for ratio 1:40 in Block 2 nevertheless shows some sensitivity to the context factors. In the case of Block 3, the Ratio × Trial interaction suggests that the increase of sample size over trials was close to 0, so participants maintained the sample size on average.

None of the other predictors were significant (all $F < 5.06$, all $p > .026$; see Tables 4 and 5 and Supplementary Material for the full table). Of course, no definitive conclusion can be drawn from the absence of a significant feedback effect. The presence of a main effect of trial across the feedback conditions, however, is reliable and surprising. It indicates that online adjustment processes took place regardless of whether participants received feedback or not, highlighting the possibility of changes that are not based on external feedback.

4.2.3. Accuracy analysis

The accuracy results, described in full in the Supplementary Material, mirror the sample-size changes for the main effect of ratio ($b_{12} = 0.22$, $SE = 0.03$, $\chi^2(1) = 39.22$, $p < .001$; $b_{32} = 0.09$, $SE = 0.03$, $\chi^2(1) = 7.83$, $p = .005$) and trial ($b_{12} = 0.07$, $SE = 0.02$, $\chi^2(1) = 10.4$, $p = .001$; $b_{32} = 0.06$, $SE = 0.03$,

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

Table 4. *Mixed model results for sample size for Block 1 and Block 2.*

Fixed effect	Estimate	SE	*df*	*F*	*p*
Ratio	0.85	0.08	1, 126.66	119.23	<.001
Trial	0.24	0.03	1, 127.19	49.02	<.001
Feedback	−0.41	0.19	1, 126.8	4.46	.037
Cost change	−0.15	0.19	1, 126.8	0.61	.438
Ratio × Trial	0.10	0.03	1, 126.39	9.4	.003
Ratio × Feedback	0.05	0.08	1, 126.66	0.41	.523
Trial × Feedback	−0.03	0.03	1, 127.19	0.61	.435
Ratio × Cost change	0.10	0.08	1, 126.66	1.57	.212
Trial × Cost change	0.02	0.03	1, 127.19	0.23	.63
Feedback × Cost change	−0.26	0.19	1, 126.8	1.85	.176

Note: For mixed model of Block 1 and Block 2: Estimate of fixed effect, standard error *SE*, degrees of freedom *df*, *F*-value, and *p*-value for main effects and two-way interactions. Fixed effect labels in bold indicate significant results ($p < .05$). See the Supplementary Material for details of higher-order interactions.

Table 5. *Linear mixed model results for sample size for Block 3 and Block 2.*

Fixed effect	Estimate	*SE*	*df*	*F*	*p*
Ratio	0.49	0.06	1, 126.94	58.43	<.001
Trial	0.18	0.03	1, 126.88	40.58	<.001
Feedback	−0.48	0.21	1, 127	5.06	.026
Cost change	−0.11	0.21	1, 127	0.26	.612
Ratio × Trial	0.16	0.03	1, 127.06	30.11	<.001
Ratio × Feedback	0.12	0.06	1, 126.94	3.58	.061
Trial × Feedback	−0.02	0.03	1, 126.88	0.34	.562
Ratio × Cost change	0.06	0.06	1, 126.94	0.77	.383
Trial × Cost change	0.01	0.03	1, 126.88	0.10	.748
Feedback × Cost change	−0.30	0.21	1, 127	1.98	.161

Note: For mixed model of Block 3 and Block 1: Estimate of fixed effect, standard error *SE*, degrees of freedom *df*, *F*-value, and *p*-value for main effects and two-way interactions. Fixed effect labels in bold indicate significant results ($p < .05$). See the Supplementary Material for details for higher-order interactions.

$\chi^2(1) = 5.49$, $p = .019$), but without their interaction reaching significance. An additional divergence is a significant interaction between ratio and cost-change type for Blocks 2 and 3, such that the positive effect of the 1:40 ratio on accuracy was more pronounced for the conditions that experienced the payoff versus the cost change ($b_{32} = 0.07$, $SE = 0.03$, $\chi^2(1) = 5.07$, $p = .024$). Interestingly, this effect was not found for sample size, so the changes in accuracy may have been based on factors other than the information basis. For Blocks 1 and 2, the interaction of ratio, trial, and cost change was significant ($b_{12} = 0.05$, $SE = 0.02$, $\chi^2(1) = 5.3$, $p = .021$).

4.2.4. Efficiency analysis

To examine whether adjustment led to changes in efficiency (i.e., the number of points participants earned each trial) within the blocks, we fitted a linear mixed model with the same predictors and coding as for the sample size analysis above, including a random slope for the ratio varying by participant. We standardized the balance by recalculating it with an information cost of 1 and a payoff of 10 or 40 to maintain the ratios but control for the specific points the ratio was instantiated with to be able to compare the different change conditions. Since the payoff was still different for different blocks after rescaling, the ratio as a predictor is not of interest, since a difference reflects the payoff amount more than the behavior. The main predictor of interest is trial, since it would indicate an efficiency change

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

Table 6. *Descriptive statistics of the optimal strategy by ratio and block.*

Ratio	Odds ratio thr.	Arrow diff.	Mean sample size (*SD*)	Mean accuracy (*SD*)	Mean Eff (*SD*)	*P*s' mean standardized absolute efficiency	*P*s' mean relative efficiency
1:10	5.44	2	3.45 (2.23)	.846 (.361)	3.47 (7.5)	Block 1 2.10 (1.83)	61.12 (53.3)
						Block 3 2.20 (1.91)	64.06 (55.56)
1:40	69.16	5	8.95 (3.17)	.985 (.120)	29.88 (10.11)	21.23 (7.10)	70.99 (23.74)

Note: Odds ratio thr. is the odds ratio threshold for stopping that maximized the efficiency for that ratio; Arrow diff. is the arrow difference that was equivalent to the mean sample size; correctness and efficiency are the means and standard deviations resulting from those thresholds. Participant efficiency is the absolute mean efficiency per trial (abs.) or the relative mean efficiency compared to the optimal strategy (rel.).

within blocks. However, only the ratio was a significant predictor of efficiency ($b_{23} = 9.60$, $SE_{23} = 0.33$, $F(1,127.01) = 873.15$, $p < .001$), which is a necessary consequence of the different payoff amounts. Importantly, there was no effect of trial on efficiency, indicating that there was no efficiency change across the blocks despite the changes in sample size and accuracy that online adjustment led to (see Supplementary Material for full results).

4.2.5. Optimal strategy and comparison

We specified an optimal strategy based on the odds ratio of the posterior probabilities of the hypotheses that .3 versus .7 were the true *p* parameters of the Bernoulli distribution as a stopping criterion. We determined the ratio that maximized the efficiency through simulations of the information acquisition and decision process with the same constraints as the experimental task (regarding cost per sample, payoff, maximum sample size of 16 for this experiment, and 32 for Exp. 3) with odds ratio thresholds {1, 2, 5, 10, 20, 50, 100, 200, 400}, with one representing equal evidence for both hypotheses, and each step beyond capturing the odds ratio change brought about by an additional difference in upward and downward pointing arrows with 100,000 runs per threshold. We aggregated the mean sample size, arrow difference, payoff, and actual achieved odds ratio to determine the odds ratio threshold that maximized the efficiency (see Table 6).

Compared to the maximum efficiency possible, participants fell short, achieving only 60–70% of the maximum. Although their efficiency scores did not change within the blocks, their sample size and accuracy over the course of Block 2 did increase, showing that the prerequisites for an efficiency change through online adjustment processes were present but did not quite translate into an efficiency increase. In Block 3, on the other hand, the lack of efficiency change across the trials was a reflection of stable sample size and accuracy. Although sample size should have ideally decreased to lead to an increase in efficiency, sample sizes were fairly close to those of an optimal strategy from the start. So, although participants fell short of the maximum efficiency, in contrast to Block 2, they did not increase their sample size over the course of the block. So, they did not just execute the same behavior, showing some adaptivity to the new context factors and successfully avoiding a decrease in efficiency.

4.3. Discussion

The results show that participants adapt to changes in the cost–payoff ratio not just by changing their sample size at the start of each block via a priori planning processes, but also by changing it within the blocks through online adjustment processes. The adjustment was also reflected in accuracy changes but did not manifest in efficiency changes. Nevertheless, in the case of the less expensive ratio (1:40

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

in Block 2), the adjustments within the block brought participants closer to the sample size of the optimal strategy. For the more expensive ratio (1:10, Blocks 1 and 3), participants tended to have sample sizes larger than the most efficient strategy. Here, the online adjustment seems to have been misguided initially, leading to a slight increase in sample size away from more efficient sample sizes in Block 1. In Block 3, at least, the sample sizes were stable, albeit still not moving in the direction of the sample sizes exhibited by the optimal strategy.

Surprisingly, there was no evidence that conditions with feedback exhibited a different within-block change than those conditions without feedback. All participants changed their sample size across trials approximately the same way. Feedback provides the necessary information to evaluate one's performance to be able to improve it as part of online adjustment processes. In the conditions without feedback, the within-block change occurred without this information. Instead of processing actual feedback, participants in the no-feedback conditions may have been mentally simulating the feedback and the implications of the change in ratio and further adjusting the sample size based on the experience of the simulation. In fact, Elwin et al. (2007) showed that participants with and without feedback on decisions performed equally well, leading them to assume that participants used a type of self-generated feedback that replaced actual feedback when it was not available. In addition, it may take a few trials of experiencing the task and monitoring the planned strategy to fully implement and specify it.

Since the conditions that did receive feedback had a much better informational basis to improve their efficiency, it is nevertheless surprising that they did not improve more than the no-feedback conditions. A reason for this may lie in the temporal horizon of monitoring. Since feedback was provided after every trial, the simplest way of using it was to consider mainly the most recent trial. However, such processing of feedback is not very informative for performance in a task with probabilistic elements in which a suitable trade-off needs to be achieved between the cost of information and its benefit for accuracy. To be able to evaluate a particular strategy, one would need to monitor with a wider horizon, aggregating feedback (about accuracy, info costs, and gains) across multiple trials to be able to compare accuracy rates instead of individual instances of correct and incorrect decisions. We designed a final experiment to investigate how different feedback structures support different monitoring horizons for online adjustment processes as evidenced by performance improvement.

5. Experiments 3a and 3b

To examine whether different monitoring horizons could lead to differences in performance improvements, we implemented two different feedback types in a new structure: each block comprised 10 consecutive decision trials, with the ratio of information costs and payoff staying constant across all blocks. In Experiment 3a, this ratio was 1:20; in Experiment 3b, two ratios, 1:10 and 1:40, were implemented as additional between-participant factors. In both experiments, the first of two feedback-type conditions, the trial-feedback condition, received the same trial-by-trial feedback as realized in Experiments 1 and 2 about the different aspects of each decision within a block, which we assumed would facilitate a monitoring horizon too short to lead to improvements. The block-feedback condition, by contrast, received only aggregated feedback provided at the end of each 10-trial block, summarizing all aspects across those trials. This feedback type facilitated a longer monitoring horizon and reduced cognitive strain, since participants did not have to integrate their accuracy, information costs, and gains or losses across multiple trials themselves. This is also supported by Balzer et al. (1989), who found that in multiple-cue decision-making tasks, the trial-by-trial feedback is unstable because the task is probabilistic. Providing aggregated feedback increases performance improvements, unless participants struggle with aggregating the information on their own (Newell et al., 2009). The structure and blocked feedback not only facilitate the evaluation of a strategy in the block-feedback condition, but they also offer clear units to try out different strategies and easily compare their performance based on the summarized feedback for precisely that unit. Since a priori planning is not inherently different from planning that happens during a decision context, online planning, the new structure also made

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

room for more focused online planning phases because it did not have to be done concurrently with other task demands. In combination with the longer monitoring horizon, this was expected to facilitate performance improvement.

Hence, participants were provided with excellent conditions for improvement through online adjustment. We hypothesized that participants who received block feedback would change their sample size across blocks more than those who received trial-by-trial feedback in the direction of sample sizes that are more efficient.

5.1. Methods

5.1.1. Participants

For Experiment 3a, a total of 176 participants completed the study via Prolific. To be able to access the study on Prolific, participants had to be fluent in English and live in the UK. Their age ranged from 20 to 79, with a mean of 35.7 (*SD* = 11.3); 36.9% reported their sex as female, 60.2% as male, and one participant preferred not to indicate. Twelve percent reported being students, 6% did not make any indication.[3] The study took about 25 minutes and was remunerated with a fixed rate of £2.50 and an additional bonus that amounted to £1.25, on average, and ranged from 0 to £2.50. Potential participants were informed that they would earn £2.50 and an additional bonus of up to £2.50 depending on their performance in the announcement of the study on Prolific.

For Experiment 3b, we recruited 370 participants via Prolific with the same prerequisites and study details, including duration and remuneration. Participants' ages ranged from 18 to 75 years with a mean of 39 (*SD* = 11.94); 35.7% reported their sex as female, 64% as male, and 0.3% preferred not to say. Full-time employees made up 52.7% of the sample, while 16.7% worked part-time and 6% were unemployed.

5.1.2. Design and procedure

Feedback type was the main between-participant factor, with participants in the block-feedback condition only receiving summarized feedback at the end of each block and participants in the trial-feedback condition only receiving feedback after each trial. Compared to the previous experiments, the trials were separated into more but shorter blocks of 10 trials each in both conditions so the new block feedback could be offered often enough in the block-feedback condition. The block feedback informed participants about their average performance at the end of each block: their mean accuracy, the mean number of arrows and the mean number of points spent on them, and the mean amount of payoff from the decision outcomes. In addition, the sum of the payoffs and information spending for all 10 trials was displayed as total block balance. The trial-by-trial feedback consisted of the equivalent feedback components for the current trial shown after every trial and appeared one by one at a rate of 750 ms to facilitate processing: whether the decision was correct and the payoff (number of points gained or lost), the points spent on the information (without reiterating the number of arrows), the resulting total payoff for that trial, and the updated cumulative balance. After an additional 1,500 ms, the next trial started automatically. Instead of the trial feedback, the block-feedback condition saw a screen asking them to wait for the next trial to start, which happened automatically after 2,000 ms.[4] The summarized feedback at the end of each block was displayed until participants clicked a button on the screen to continue. Instead, the trial-feedback condition saw a screen informing them that the block had ended and what their final balance was for the block. Participants worked through a total of 11 blocks in Experiment 3a[5] and a total of 10 blocks in Experiment 3b. The basic task was exactly as for Experiment 2, except that

[3] Four participants' demographic data were missing, because their Prolific IDs were not saved and the demographic data provided by Prolific could not be matched in those cases.

[4] This was not quite equivalent to the break between trials in the trial-feedback condition. However, a break any longer increased the risk of boredom and inattentiveness too much.

[5] Due to an incorrect variable specification in the experimental program, the trial-by trial feedback condition actually worked through 12 blocks instead of 11. The last block was discarded to achieve equal data for the two conditions.

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

the maximum sample size per decision was 32. Experiment 3b additionally had the between-participant factor of ratio (1:10 vs. 1:40), resulting in a 2×2 between-participant design, while Experiment 3a used only one ratio (1:20).

At the end of the experiment, participants filled in a brief questionnaire that asked whether they changed their strategy over the course of the experiment (binary yes/no answer with option to comment if 'yes' was selected) and how much they incorporated the different components of the feedback. For each of the components—decision accuracy, decision outcomes (gains and losses), points spent on reviews, and the total balance—there was a separate question worded 'To what extent did you take into account the feedback about X for your strategy?' with the component replacing X in the question. Participants answered each question by clicking on a continuous slider anchored on 'not at all' and 'very much' at the ends.

Experiment 3b asked five additional questions on a second page, which were all answered on a sliding scale with ends labeled 'not at all' and 'very much.' The questions concerned confidence ('To what extent were you searching for additional information in order to gain more confidence?' and 'Overall, to what extent did searching for additional information increase your confidence by the desired amount?') and motivation ('How motivated were you to improve your performance to receive a higher bonus payment?' 'To what extent did you have the impression that it was possible to improve your performance in order to receive a higher bonus payment?' and 'How engaging did you find the task?' with the additional remark that participants' answers to this question will not be used for evaluating their participation for payment purposes, asking for a truthful answer).

5.2. Results

5.2.1. Exclusions

In Experiment 3a, 8 participants were excluded because they either took a break of more than 3 minutes or three or more breaks of one minute or longer, as based on the decision response time and time before requesting the first observation. The latter is of particular importance because the previous feedback page redirected to the sampling page automatically. Hence, it was very likely that those who missed the start of the trial also did not pay attention to the preceding feedback. Another eight participants had to be excluded because their mean sample size was below .75. This means that they had not drawn a single sample for at least three quarters of the trials, which did not yield enough data to analyze reliably. The resulting sample consisted of 160 participants.

For Experiment 3b, 18 participants were excluded based on response time in the same way as for Experiment 3a. Another 34 participants were excluded based on sample size: Either their mean sample size was below .75 or a sample size of 0 was chosen for more than 33% of trials, in which cases the data are too sparse to analyze reliably. After exclusions, 318 individual datasets remained.

5.2.2. Sample size analysis

First, we wanted to assess whether participants changed their sample size across blocks and within blocks and whether this was influenced by the type of feedback they received (and for Experiment 3b, by the ratio). Trial-by-trial feedback should facilitate short-horizon monitoring and therefore support changes within blocks, while block feedback should facilitate long-horizon monitoring and therefore also facilitate trying out a stable strategy for the duration of a block. Accordingly, to analyze the sample size data, which are depicted in Figure 6, we estimated one linear mixed model per experiment with block, trial, and feedback condition (and ratio for Exp. 3b) as fixed effects (feedback condition coded −1 for trial feedback and +1 for block feedback, block and trial centered on 0 and rescaled by dividing by 10 to make the value ranges of all variables more similar and facilitate convergence of the models). For both Experiments 3a and 3b, the maximum random effects structure that converged included the trial, the block, and the interaction of the trial and the block, all varying by participant as uncorrelated effects.

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

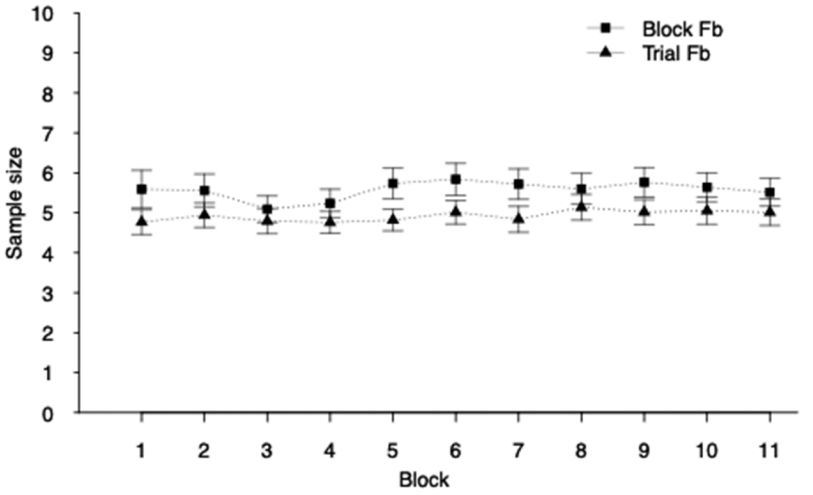


Figure 6. *Experiment 3a: Sample size across blocks separately for conditions.*
Note: Error bars indicate standard error of the mean. Very light symbols depict individual participants' values according to the legend.

Table 7. *Means and standard deviation for sample size and accuracy separately for block feedback and trial feedback.*

Ratio	Feedback condition	*n* (*SD*)	*a* (*SD*)	*Eff (SD)* 1st half	*Eff (SD)* 2nd half
1:10	Block	4.65 (4.49)	0.82 (0.06)	15.82 (89.99)	18.78 (86.42)
	Trial	3.55 (1.45)	0.80 (0.08)	25.91 (80.47)	23.11 (80.79)
1:20	Block	5.58 (3.25)	.850 (.081)	81.67 (147.26)	86.66 (140.85)
	Trial	4.92 (2.33)	.838 (.085)	88.05 (146.25)	83.96 (149.23)
1:40	Block	7.24 (5.74)	0.87 (0.07)	217.85 (276.03)	229.39 (260.48)
	Trial	5.02 (2.31)	0.85 (0.07)	214.96 (298.97)	237.45 (274.32)

Note: Means with SD in brackets for sample size *n*, accuracy *a*, and efficiency *Eff* with data for ratios 1:20 stemming from Exp. 3a and data for ratios 1:10 and 1:40 stemming from Exp. 3b.

For Experiment 3a (ratio 1:20), the analyses showed that only trial had an influence on sample size, such that participants sampled more in later trials within a block, with an average increase of about a quarter of an observation per 10 trials ($b = 0.26$, $SE = 0.08$; $F(1, 158.04) = 9.88$, $p = .002$). Neither the block nor the feedback condition, nor any of the interactions, had an influence on sample size (all $F < 2.02$, all $p > .157$; see the Supplementary Material for full table). This was not in line with our expectations. The feedback condition not making a difference means that there is no evidence that participants implemented different sample sizes depending on the type of feedback, and no evidence that they differed in how they changed their sample sizes across the trials and blocks. In fact, it seems that participants just did not change their sample size over the course of the task. Instead, there was only a change over the course of the trials within the blocks. Since one condition did not receive any feedback during the block, but only at the end, this is surprising. The descriptive statistics for the two conditions are displayed in Table 7.

For Experiment 3b, Figure 7 and Table 8 illustrate the main effect of ratio, with participants with cheaper information (ratio 1:40, white symbols in Figure 7, coded 1 in analysis) displaying sample sizes larger than the mean by about 1. There is also a main effect of the feedback condition, with participants in the block-feedback condition displaying sample sizes about 0.8 larger than the mean of the feedback conditions (black symbols in Figure 7, coded 1 in analysis). Further, block and trial also significantly

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

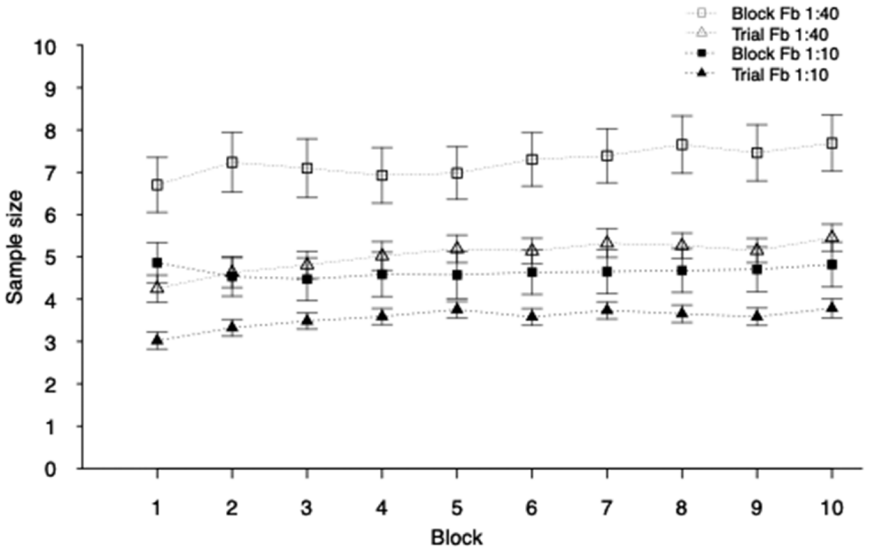


Figure 7. *Experiment 3b: Sample size across blocks separately for conditions and ratios.*

Table 8. *Linear mixed model results for sample size for Experiment 3b.*

Fixed effect	Estimate	SE	*df*	*F*	*p*
Intercept	5.12	0.23			
Ratio	**1.02**	**0.23**	**1, 307**	**20.06**	**<.001**
Feedback condition	**0.83**	**0.23**	**1, 307**	**13.37**	**<.001**
Block	**0.67**	**0.16**	**1, 306.99**	**17.35**	**<.001**
Trial	**0.26**	**0.05**	**1, 306.78**	**24.89**	**<.001**
Ratio × Fb condition	0.28	0.23	1, 307	1.53	0.217
Ratio × Block	**0.32**	**0.16**	**1, 306.99**	**3.97**	**0.047**
Condition × Block	−0.16	0.16	1, 306.99	1.05	0.306
Ratio × Trial	**0.15**	**0.05**	**1, 306.78**	**8.17**	**0.005**
Fb condition × Trial	−0.05	0.05	1, 306.78	1.11	0.293
Block × Trial	0.03	0.2	1, 307.12	0.02	0.88

Note: In the column names, *SE* denotes the standard error of the estimate, *df* the degrees of freedom for the F-test, *F* the test statistic value, and *p* the *p* value. Fixed effect labels in bold indicate significant results ($p < .05$).

influenced the sample size as main effects. Across the 10 blocks, the sample size increased by 0.67, on average. Across the 10 trials within each block, the sample size increased by 0.26 on average.

The change in sample size across blocks and trials needs to be examined further to assess the processes belonging to the fine-tuning category. Indeed, block and trial were the only fixed effects involved in interactions, with the ratio influencing both the effect of block and the effect of trial on sample size. To further examine the interaction, we fit the equivalent model to the two ratios separately (with the intercept varying by participants and random slopes for trial and block and their interaction varying by participant). This showed that the effect of block on sample size was not significant for ratio 1:10 ($b = 0.35$, $SE = 0.22$ $F(1, 158.99) = 2.60$, $p = .109$). However, the effect of trial on sample size was ($b = 0.11$, $SE = 0.05$, $F(1, 159.31) = 4.77$, $p = .030$), meaning that they increased their sample size very slightly by 1/10 per 10 trials. (Condition was also a significant predictor in line with it being a main effect without a significant interaction in the full model, $b = 0.55$, $SE = 0.27$, $F(1, 159) = 4.15$ $p = .043$).

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

The follow-up test for ratio 1:40 showed that the mean effects of block and trial on sample size were significant ($b = 0.99$, $SE = 0.24$, $F(1, 148.01) = 17.24$, $p < .001$ and $b = 0.41$, $SE = 0.09$, $F(1, 147.67) = 19.22$, $p < .001$), with sample size increasing by about one per 10 blocks and by about 0.4 per 10 trials. (Here, too, the main effect of condition was significant, $b = 1.11$, $SE = 0.37$, $F(1, 148) = 8.96$, $p = .003$, in line with the significant main effect in the full model.) There were no other significant predictors in the full model (all Fs < 2.3, all $p > .131$; see the Supplementary Material for complete tables), in particular no significant interactions with the feedback condition, meaning that the different types of feedback did not lead to differences in sample size or to differing patterns of the interactions. The interaction between block and trial was not significant, meaning that the increase of sample size with trial number did not differ systematically across the different blocks. All in all, participants with a ratio of 1:40 increased their sample size consistently, while participants with a ratio of 1:10 only displayed an increase across trials, which did not translate into an increase across blocks—similar to the ratio of 1:20 in Experiment 3a.

Accuracy. We fitted a generalized mixed linear model to the accuracy data with the same predictors as for the sample size analysis. The highest number of random effects that still allowed the model to converge for Experiment 3a was a random slope for block, varying by participant as correlated effects. There was a main effect of block, such that accuracy increased across blocks ($b = 0.19$, $SE = 0.08$, $\chi^2(1) = 5.77$, $p = .016$). Interpreted as odds, this means that accuracy increased by about 21% per 10 blocks. There was also a significant interaction of block and trial, which is easiest to interpret as block number moderating the change of accuracy across trials, despite the main effect of trial not being significant ($b = 0.12$, $SE = 0.07$, $\chi^2(1) = 2.68$, $p = .102$). For early blocks, the increase in accuracy across trials is very strong, but this increase fades toward the middle of the task and even reverses to a decrease in accuracy across trials in late blocks ($b = -0.48$, $SE = 0.23$, $\chi^2(1) = 4.07$, $p = .044$). There were no other significant effects (see the Supplementary Material for full table).

For Experiment 3b, there was also a main effect of block ($b = 0.26$, $SE = 0.06$, $\chi^2(1) = 16.72$, $p < .001$), indicating that accuracy increased by 29.7% for 10 blocks. The feedback condition had a small main effect ($b = 0.09$, $SE = 0.03$, $\chi^2(1) = 7.84$, $p = .005$), but it did not moderate the development of accuracy across blocks ($b = 0.01$, $SE = 0.06$, $\chi^2(1) = 0.03$, $p = .874$). As would be expected, the ratio of 1:40 had a higher accuracy than the ratio 1:10 ($b = 0.2$, $SE = 0.03$, $\chi^2(1) = 38.08$, $p < .001$). In line with the higher sample size increase for ratio 1:40, there was an interaction between ratio and block such that the accuracy increased more for ratio 1:40 than ratio 1:10 ($b = 0.19$, $SE = 0.06$, $\chi^2(1) = 8.52$, $p = .004$). No other predictors were significant (see the Supplementary Material for full table).

Efficiency. To examine differences in efficiency, we fitted a mixed model with condition, block, trial, and their interactions as fixed effects for Experiment 3a and the additional factor of ratio and its interactions for Experiment 3b. Block and trial were not scaled for these models. The maximum random effects structure for Experiment 3a included only the intercept varying by participant. For Experiment 3b, block, trial, and their interaction were included as uncorrelated random effects varying by participant.

Although no significant main effects emerged for Experiment 3a (ratio 1:20), we found two interesting interactions. First, the interaction between the feedback condition and block ($b = 0.69$, $SE = 0.35$, $F(1, 17433.04) = 3.96$, $p = .046$) indicates that the block-feedback condition increased their efficiency across the blocks more than the trial-feedback condition did, whose change across blocks was even closer to 0 than the non-significant main effect of block ($b = 0.56$, $SE = 0.35$, $F(1, 17433.04) = 2.62$, $p = .105$). With each block, the block-feedback condition is estimated to have increased their efficiency by $0.56 + 0.69 = 1.25$, on average. This effect was not accompanied by a significant change in sample size, so participants improved—very slightly—without implementing a change in strategy that affected their mean sample size. It is plausibly related to the increase in accuracy identified above, although that was the same for the two feedback conditions, so it cannot explain the entirety of the effect. Similarly, despite the absence of main effects of block or trial, the effect of the trial number on efficiency was moderated by the block number. In early blocks, the efficiency increased across trials,

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

Table 9. *Linear mixed model results for efficiency for Experiment 3b.*

Fixed effects	Estimate	*SE*	*Df*	*F*	*p*
(Intercept)	122.92	2.21			
Feedback condition (Fb cond)	−2.44	2.21	1, 307.18	1.23	.269
Block	**1.58**	**0.43**	**1, 307.37**	**13.47**	**<.001**
Trial	0.64	0.41	1, 307.34	2.44	.119
Ratio	**102**	**2.21**	**1, 307.18**	**2,133.39**	**<.001**
Fb cond × Block	−0.17	0.43	1, 307.37	0.15	.697
Fb cond × Trial	−0.53	0.41	1, 307.34	1.64	.202
Block × Trial	−0.06	0.14	1, 307.19	0.16	.687
Fb cond × Ratio	1.16	2.21	1, 307.18	0.27	.601
Block × Ratio	**1.72**	**0.43**	**1, 307.37**	**15.89**	**<.001**
Trial × Ratio	**0.81**	**0.41**	**1, 307.34**	**3.9**	**.049**
Fb cond × Block × Trial	0.01	0.14	1, 307.19	0	.944
Fb cond × Block × Ratio	−0.65	0.43	1, 307.37	2.29	.131
Fb cond × Trial × Ratio	−0.73	0.41	1, 307.34	3.16	.076
Block × Trial × Ratio	0.02	0.14	1, 307.19	0.01	.905
Fb cond × Block × Trial × Ratio	0.11	0.14	1, 307.19	0.6	.439

Note: In the column names, *SE* denotes the standard error of the estimate, *df* the degrees of freedom for the F-test, *F* the test statistic value, and *p* the *p* value. Fixed effect labels in bold indicate significant results ($p < .05$).

but this increase disappeared or reversed in later blocks ($b = -0.26$, $SE = 0.12$, $F(1, 17433.04) = 4.49$, $p = .034$).

For Experiment 3b (ratios 1:10 and 1:40), the results, displayed in Table 9, show that block and ratio had significant main effects on efficiency, with a mean increase of 1.58 efficiency units per block and ratio 1:10 (coded −1) having an efficiency about 200 points lower than ratio 1:40 (coded 1). The ratio moderated the efficiency increase across the blocks, as indicated by the significant interaction. For ratio 1:40 (coded 1), the increase across the blocks was about 3.3, while for ratio 1:10, it amounted to close to 0. There is also a significant interaction of trial and ratio, likewise indicating that efficiency for ratio 1:40 may have been increasing across trials within blocks, while it is close to 0 for ratio 1:10.

Following up on these simple effects, we fit separate mixed-effect models for each ratio. For the ratio 1:10 model, we could not maintain the same random effects structure because it did not converge. Instead, only a random intercept and a random slope for the block, both varying by participant, could be included as uncorrelated effects. However, the estimates of the predictors were consistent with the full model. For ratio 1:40, block increased efficiency by 3.31 points ($b = 3.31$, $SE = 0.81$, $F(1,148.06) = 15.23$, $p < .001$). The trial's estimated effect of 1.46 ($b = 1.46$, $SE = 0.82$, $F(1, 147.97) = 3.13$, $p = .079$) was not significant. For a ratio of 1:10, the effect of block on efficiency was not significant ($b = -0.14$, $SE = 0.27$, $F(1, 159.22) = 0.25$, $p = .616$), and neither were any of the other predictors (see the Supplementary Material for full table).

Optimal strategy and comparison. The optimal strategy was determined the same way as for Experiment 2 except that a maximum sample size of 32 was implemented. Table 10 displays the maximum efficiency, the sample size and accuracy it resulted in, and the arrow difference and odds ratio underlying it (see the Supplementary Material for tables of all simulated ratios). Achieving almost 80%, participants' efficiency was high compared to the maximum possible for ratios 1:20 and 1:40. For ratio 1:10, however, they only reached about 60%, a substantially lower amount.

Questionnaire. In Experiment 3a, 60.56% of participants in the trial-feedback condition indicated that they changed their strategy during the experiment, while 41.57% did in the block-feedback condition. This difference was significant, as a test for equality of proportions showed, $\chi^2(1) = 4.96$, $p = .026$. In Experiment 3b, 48.61% in the trial-feedback and 50.30% in the block-feedback condition

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

Table 10. *Descriptive statistics for optimal strategy by ratio.*

Ratio	Odds ratio thr.	Arrow diff.	Mean sample size (*SD*)	Mean correct (*SD*)	Mean optimal efficiency (*SD*)	Mean participants' efficiency abs. (*SD*)	Mean participants' efficiency rel.
1:10	5.44	2	3.45 (2.23)	.846 (.361)	3.47 (7.56)	2.07 (8.48)	59.67%
1:20	12.70	3	6.37 (4.17)	.927 (.261)	10.70 (11.23)	8.50 (2.41)	79.45%
1:40	29.64	4	9.14 (5.50)	.967 (.178)	28.26 (15.23)	22.48 (27.70)	79.56%

Note: Odds ratio thr. is the odds ratio threshold for stopping that maximized the efficiency for that ratio; Arrow diff. is the arrow difference that was equivalent to the mean sample size; correctness and efficiency are the means and standard deviations resulting from those thresholds. Participant efficiency is the absolute mean efficiency per trial (abs.) or the relative mean efficiency compared to the optimal (rel.).

Table 11. *Descriptive and inferential statistics for questionnaire results of Experiments 3a and 3b.*

Feedback aspects	Trial-feedback Mean (*SD*)	Block feedback Mean (*SD*)	*df*	*t*	*p*	*Cohen's d*
Experiment 3a						
Correctness	733.5 (246.5)	700.0 (270.6)	158	0.81	0.42	−0.129
Outcomes	758.5 (198.6)	653.9 (268.4)	158	2.74	0.007	−0.436
Points spent on reviews	627.8 (289.6)	561.4 (303.1)	158	1.4	0.162	−0.223
Balance	667.1 (284.5)	590.0 (307.3)	158	1.63	0.105	−0.259
Experiment 3b						
Correctness	704.8 (242.8)	759.8 (231.1)	309	−2.04	0.042	−0.232
Outcomes	666.5 (258.3)	709.4 (252.9)	309	−1.48	0.141	−0.168
Points spent on reviews	654.0 (261.1)	591.6 (299.4)	309	1.94	0.053	0.222
Balance	599.7 (296.1)	656.1 (296.2)	309	−1.67	0.095	−0.19

Note: Participants' mean ratings of how much they incorporated the feedback about different aspects into their strategy (scale 0–1,000) and standard deviation (*SD*) in brackets, test statistics *t* of two-sample t-tests between the two conditions, *p*-value indicated by *p*, and effect size by Cohen's *d*. Dependent variable labels in bold indicate significant results ($p < .05$).

indicated that they changed their strategy, which was not a significant difference, $\chi^2(1) = -.14$, $p = .855$. The mean ratings of how much participants incorporated the feedback components into their strategy, displayed in Table 11, only differed significantly between conditions for the feedback about decision outcomes in Experiment 3a and for the feedback about correctness in Experiment 3b. The trend in Experiment 3a toward the trial-feedback condition, indicating to have paid more attention to feedback overall, did not appear in Experiment 3b, where the only significant difference was in the opposite direction. Across both experiments, participants indicated values above the midpoints, meaning that they reported having incorporated the feedback into their strategy to a considerable extent with only minor differences between the conditions. Hence, the lack of sample-size change in Experiment 3a, for ratio 1:20, was presumably not just the result of participants not paying attention to the feedback.

The additional questions posed at the end of Experiment 3b show that participants were motivated to achieve a higher bonus payment, thought it was feasible, and found the task engaging (see Table 12). The two conditions' ratings were in a very similar range descriptively, which is important to know because trial-by-trial feedback may have made the task seem more feasible or made it more engaging. Two questions concerning the role of confidence were included to inform future studies and are not of interest here.

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

Table 12. *Descriptive statistics of additional questions in Experiment 3b.*

Ratio	Feedback condition	Motivation to increase bonus (*SD*)	Feasibility of increasing bonus (*SD*)	Task engaging (*SD*)	Confidence as aim (*SD*)	Confidence aim achieved (*SD*)
1:10	Block	854.7 (190.1)	734.4 (228.8)	856.3 (194.2)	626.4 (254.2)	650.7 (244.1)
	Trial	858.3 (185.9)	737.1 (240.8)	845.0 (164.5)	628.3 (230.5)	652.6 (248.3)
1:40	Block	832.5 (197.2)	726.4 (263.1)	855.4 (157.7)	661.7 (260.6)	649.7 (257.0)
	Trial	881.2 (142.6)	766.3 (243.6)	849.9 (183.9)	626.6 (274.8)	653.6 (271.5)

Note: Participants' mean ratings answering the additional questions described above with *SD* in brackets (scale 0–1,000).

5.3. Discussion

Whether participants changed their sample size across the blocks of the experiment depended on the cost–payoff ratio that was implemented. For ratio 1:40 in Experiment 3b, participants did increase their sample size across the blocks (and trials), which also led to an efficiency increase, clearly demonstrating online adjustment processes. However, for the other two ratios, 1:10 in Experiment 3b and 1:20 in Experiment 3a, participants did not change their sample size across the blocks. In Experiment 3a, with a ratio of 1:20, participants did not change their sample size across the blocks but merely across trials within blocks. Across all three cost–payoff ratios, the feedback type did not have an influence on participants' sample size development. It did have an effect on the development of efficiency across the blocks in the case of Experiment 3a, ratio 1:20, where the block-feedback condition exhibited a higher increase in efficiency across blocks than the trial-feedback condition. Nevertheless, the block feedback in combination with the block structure was not more successful at facilitating the evaluation of strategies based on longer-horizon monitoring than the trial-by-trial feedback. Also, instead of varying their sample size across blocks but not within blocks to systematically try out different strategies and compare them, participants more often increased their sample size within the blocks, which speaks against systematic exploration of strategies in different blocks.

We expected that 100 (or 110) trials of a fairly simple task would provide participants with ample opportunity to explore and improve. We assumed that the ten-trial blocks would provide the necessary structure and incentive to try out different strategies, with the block feedback ideally suited for comparing the blocks and strategies, especially since participants did indicate having taken the feedback into consideration. However, evaluating different strategies comparatively even with a more suitable structure is still a cognitively challenging task. In addition, to be able to evaluate different strategies, one needs to implement different strategies in the first place. Since participants did not vary their sample size systematically across blocks, there may not have been much to monitor and evaluate with the help of feedback.

One plausible reason for the lack of exploration and monitoring is offered by the efficiency comparison, which shows that participants' efficiency is quite high. This may very well explain the lack of exploration and improvement because the signal to initiate a monitoring and control process may have been missing. Participants may have been satisfied with the efficiency they perceived to have and may not have felt the need to try out different strategies to find an even better one.

6. General discussion

We investigated two stages of processes through which people adapt their information search to changes in decision contexts. The first, a priori planning, is characterized as an intentional plan to implement behavior made before a transparent change in the context, allowing for an immediate behavior adaptation in line with the new context. We found support for processes of this stage in

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

Experiments 1a and 1b, with participants changing their sample size between blocks in response to the announced changes. The results of Experiment 2 likewise showed support in the form of block-wise changes. It also offered intriguing insight into the second stage of adaptation processes, namely, online adjustment, which encompasses all the processes that occur within the same context. We focused on monitoring and evaluation as parts of online adjustment, with online planning playing a role in Experiments 3a and b. Experiment 2 showed that processes based on externally provided feedback are indeed not the only possible processes, since participants changed their sample size across the blocks independently of whether they received feedback or not. It is plausible to assume that participants without feedback relied on their own mental simulation of the feedback to further adapt to the new context, similar to what Elwin et al. (2007) documented. Consequently, monitoring is applied not only to external information but also to internal signals (Nelson and Narens, 1990) and thus occurs beyond situations in which feedback is available.

This is an important extension for metacognitive theories, since it contributes another layer of structure through phases characterized by different metacognitive processes. The a priori planning phase has been neglected because it can only be investigated with transparent announced changes (or through recurring contexts where decision makers can apply prior experience), which is hardly ever the case in metacognitive research. More broadly, it builds our understanding of humans' adaptivity to the many different demands of their environment.

Certain aspects of the findings also raise questions. External performance feedback should allow monitoring to lead to improvement more easily, especially where the blocked structure facilitated a longer monitoring horizon (Balzer et al., 1989). This was not the case in Experiments 3a and b. Likely candidate explanations for the lack of improved performance despite blocked feedback are the persisting cognitive difficulty of evaluating strategies comparatively and, more importantly, the lack of different strategies to be compared. One combination of metacognitive monitoring and control functions that is not often mentioned but would be required for improvement here is for the control function to implement different strategies (according to an exploration plan) so the monitoring can be informative at all. Although we expected the short-block structure to naturally invite this kind of exploration with facilitating feedback provided at the end, participants still needed to feel the need to try out different variations and monitor them closely. The high efficiency that they achieved may have been satisfactory enough not to elicit that need. Past research (Fiedler et al., 2021) showed that participants omit exploration of strategies (in a speed-accuracy decision context) even when their efficiency is substantially suboptimal, potentially because their efficiency seems subjectively satisfactory. Further investigating to what extent people can estimate the potential increase in efficiency and the effort to achieve it would be very valuable to further understand the processes that help people engage with their environment suitably.

An indication of how much better participants could perform or how others are performing might be enough to prompt systematic exploration and evaluation. This would be similar but also crucially different to the types of feedback that Evans and Brown (2017) implemented. Their elaborate feedback not only provided participants with the information of how much better they could be but also advised them directly how they needed to change their strategy to achieve it. Hence, no exploration, monitoring, or evaluation was necessary before a better strategy could be executed. To study the metacognitive processes involved as part of online adjustment, it would be more interesting to prompt participants to explore themselves and identify the minimum assistance required for them to improve efficiency.

There may also be other reasons for participants not to attain even higher levels of efficiency. Perhaps people expect dynamic environments and were applying strategies accordingly (Navarro et al., 2016) or might have been searching for patterns where there are none (Cohen and Teodorescu, 2022; Plonsky et al., 2015) instead of accepting the stochastic nature of the task and focusing on improving their efficiency. Even though we tried to exclude these possibilities by clarifying the nature of the task in the instructions, participants may not have read those carefully enough or may not have understood the implications—or may just be displaying behavior that is adaptive in the real, dynamic world full of patterns instead of the lab regardless. The reason we implemented a somewhat artificial task that was not

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

intended to represent the environments' informational structure (Dhami et al., 2004) is that we are not merely interested in determining whether people are adaptive in close-to-real-world decision situations. We wanted to investigate processes that allow people to become adaptive by confronting them with an unfamiliar task that only emulated selected real-world features, such as the sampling of observations. Hence, the artificial task structure is both a strength and limitation of the present experiments.

Ideally, these questions should be pursued further with both a naturalistic and an 'artificial' approach. It would be of great value to specify the assistance necessary for people to successfully adapt through online adjustment processes. The flipside of this question is to what extent those conditions are actually encountered in the real world for people to become adaptive, for which a naturalistic approach would be better suited. Lieder and Griffiths (2017) found support for a model of learning heuristic decision strategies that assumes that people can learn the cost–benefit trade-offs of different heuristics in different environments. This indicates that people are indeed capable of monitoring and evaluating different strategies along with their costs and benefits to select the most efficient under some circumstances. Discovering how exactly they do so and what those circumstances are would greatly advance our understanding of adaptivity in information search and decision-making. The complementary knowledge gained about the environmental structure might even allow us to understand when and why people fail to become adaptive or why a seemingly flawed response may be adaptive after all when taking into account the resources that would need to be expended for improving it (e.g., Lieder et al., 2018).

7. Conclusion

Overall, the experiments demonstrated that participants respond to changes in decision contexts via a priori planning processes and, to some extent, online adjustment processes for further efficiency improvements. They also showed that the conditions for further improvement of efficiency are not straightforward due to the complex nature of the metacognitive processes required. Further investigating the different metacognitive processes involved would allow us to better understand human information search and decision-making and how it develops to be adaptive. This might eventually enable us to assist not only decision-making directly—where even an adaptive response is bound to fall short—but also the development of adaptive decision-making for environments that may not naturally offer the required conditions.

Supplementary material. The supplementary material for this article can be found at http://doi.org/10.1017/jdm.2026.10043.

Funding statement. This work was supported by the Deutsche Forschungsgemeinschaft (DFG) under Grant FI 294/30-1 and by an internal grant from TU Dresden.

Competing interest. The authors declare none.

References

Ackerman, R., Douven, I., Elqayam, S., & Teodorescu, K. (2020). Satisficing, meta-reasoning, and the rationality of further deliberation. In S. Elqayam, I. Douven, J. St, B. T. Evans, & N. Cruz (Eds.), *Logic and uncertainty in the human mind: A tribute to David E. Over* (1st ed., pp. 10–26). Routledge: Abingdon.

Balzer, W. K., Doherty, M. E., & O'Connor, R. (1989). Effects of cognitive feedback on performance. *Psychological Bulletin*, *106*(3), 410–433. https://doi.org/10.1037/0033-2909.106.3.410.

Bogacz, R., Hu, P. T., Holmes, P. J., & Cohen, J. D. (2010). Do humans produce the speed-accuracy trade-off that maximizes reward rate? *The Quarterly Journal of Experimental Psychology*, *63*(5), 863–891. https://doi.org/10.1080/17470210903091643.

Busemeyer, J. R., & Rapoport, A. (1988). Psychological models of deferred decision-making. *Journal of Mathematical Psychology*, *32*(2), 91–134. https://doi.org/10.1016/0022-2496(88)90042-9.

Chater, N. (2009). Rational and mechanistic perspectives on reinforcement learning. *Cognition*, *113*(3), 350–364. https://doi.org/10.1016/j.cognition.2008.06.014.

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

Cohen, D., & Teodorescu, K. (2022). On the effect of perceived patterns in decisions from sampling. *Decision*, *9*(1), 21–42. https://doi.org/10.1037/dec0000159.

Dayan, P., & Niv, Y. (2008). Reinforcement learning: The good, the bad and the ugly. *Current Opinion in Neurobiology*, *18*(2), 185–196. https://doi.org/10.1016/j.conb.2008.08.003.

Dhami, M. K., Hertwig, R., & Hoffrage, U. (2004). The role of representative design in an ecological approach to cognition. *Psychological Bulletin*, *130*(6), 959–988. https://doi.org/10.1037/0033-2909.130.6.959.

Edwards, W. (1965). Optimal strategies for seeking information: Models for statistics, choice reaction times, and human information processing. *Journal of Mathematical Psychology*, *2*(2), 312–329. https://doi.org/10.1016/0022-2496(65)90007-6.

Elwin, E., Juslin, P., Olsson, H., & Enkvist, T. (2007). Constructivist coding: Learning from selective feedback. *Psychological Science*, *18*(2), 105–110. https://doi.org/10.1111/j.1467-9280.2007.01856.x.

Evans, N. J., & Brown, S. D. (2017). People adopt optimal policies in simple decision-making, after practice and guidance. *Psychonomic Bulletin & Review*, *24*(2), 597–606. https://doi.org/10.3758/s13423-016-1135-1.

Fiedler, K., McCaughey, L., Prager, J., Eichberger, J., & Schnell, K. (2021). Speed-accuracy trade-offs in sample-based decisions. *Journal of Experimental Psychology: General*, *150*(6), 1203–1224. https://doi.org/10.1037/xge0000986.

Gershman, S. J., & Niv, Y. (2010). Learning latent structure: Carving nature at its joints. *Current Opinion in Neurobiology*, *20*(2), 251–256. https://doi.org/10.1016/j.conb.2010.02.008.

Guan, M., & Lee, M. D. (2018). The effect of goals and environments on human performance in optimal stopping problems. *Decision*, *5*(4), 339–361. https://doi.org/10.1037/dec0000081.

Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description–experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making*, *21*(5), 493–518. https://doi.org/10.1002/bdm.598.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decision from experience and the effect of rare events in risky choice. *Psychological Science*, *15*(8), 534–539. https://doi.org/10.1111/j.0956-7976.2004.00715.x.

Irwin, F. W., & Smith, W. A. S. (1957). Value, cost, and information as determiners of decision. *Journal of Experimental Psychology*, *54*(3), 229–232. https://doi.org/10.1037/h0049137.

Lee, M. D., Newell, B. R., & Vandekerckhove, J. (2014). Modeling the adaptation of search termination in human decision making. *Decision*, *1*(4), 223–251. https://doi.org/10.1037/dec0000019.

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

Lieder, F., Griffiths, T. L., Huys, Q. J. M., & Goodman, N. D. (2018). The anchoring bias reflects rational use of cognitive resources. *Psychonomic Bulletin & Review*, *25*(1), 322–349. https://doi.org/10.3758/s13423-017-1286-8.

McCaughey, L., Ziegler, J., & Fiedler, K. (2025). Rivals reloaded: Adapting to sample-based speed–accuracy trade-offs through competitive pressure. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *51*(7), 1034–1047. https://doi.org/10.1037/xlm0001408.

Navarro, D. J., Newell, B. R., & Schulze, C. (2016). Learning and choosing in an uncertain world: An investigation of the explore-exploit dilemma in static and dynamic environments. *Cognitive Psychology*, *85*, 43–77. https://doi.org/10.1016/j.cogpsych.2016.01.001.

Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. *Psychology of Learning and Motivation*, *26*, 125–173. https://doi.org/10.1016/s0079-7421(08)60053-5.

Newell, B. R., Weston, N. J., Tunney, R. J., & Shanks, D. R. (2009). The effectiveness of feedback in multiple-cue probability learning. *The Quarterly Journal of Experimental Psychology*, *62*(5), 890–908. https://doi.org/10.1080/17470210802351411.

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.

Phillips, N. D., Hertwig, R., Kareev, Y., & Avrahami, J. (2014). Rivals in the dark: How competition influences search in decisions under uncertainty. *Cognition*, *133*(1), 104–119. https://doi.org/10.1016/j.cognition.2014.06.006.

Pitz, G. F. (1968). Information seeking when available information is limited. *Journal of Experimental Psychology*, *76*(1), 25–34. https://doi.org/10.1037/h0025302.

Pitz, G. F., Reinhold, H., & Scott Geller, E. (1969). Strategies of information seeking in deferred decision making. *Organizational Behavior and Human Performance*, *4*(1), 1–19. https://doi.org/10.1016/0030-5073(69)90028-2.

Plonsky, O., Teodorescu, K., & Erev, I. (2015). Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychological Review*, *122*(4). https://doi.org/10.1037/a0039413.

Posit team. (2025). *RStudio: Integrated development environment for R*. Posit Software, PBC. http://www.posit.co/.

R Core Team. (2025). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. https://www.R-project.org/.

Rakow, T., & Miler, K. (2009). Doomed to repeat the successes of the past: History is best forgotten for repeated choices with nonstationary payoffs. *Memory & Cognition*, *37*(7), 985–1000. https://doi.org/10.3758/MC.37.7.985.

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*(4), 873–922. https://doi.org/10.1162/neco.2008.12-06-420.

Schraw, G., & Moshman, D. (1995). Metacognitive theories. *Educational Psychology Review*, *7*(4), 351–371. https://doi.org/10.1007/BF02212307.

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.

Singmann, H., & Kellen, D. (2019). An introduction to mixed models for experimental psychology. In D. H. Spieler & E. Schumacher (Eds.), *New methods in cognitive psychology* (pp. 4–31). Psychology Press: New York.

Singmann, H., Bolker, B., Westfall, J., Aust, F., & Ben-Shachar, M. (2024). *afex: Analysis of factorial experiments*. R package version 1.4-1. https://doi.org/10.32614/CRAN.package.afex.

Spektor, M. S., & Wulff, D. U. (2024). Predecisional information search adaptively reduces three types of uncertainty. *Proceedings of the National Academy of Sciences of the United States of America*, *121*(47). https://doi.org/10.1073/pnas.2311714121.

Srivastava, N., Muller-Trede, J., Schrater, P., & Vul, E. (2016). Modeling sampling duration in decisions from experience. In *Proceedings of the annual meeting of the cognitive science society*, *38*. https://escholarship.org/uc/item/6db3h1j0.

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). The MIT Press.

Wulff, D. U., Mergenthaler-Canseco, M., & Hertwig, R. (2018). A meta-analytic review of two modes of learning and the description-experience gap. *Psychological Bulletin*, *144*(2), 140–176. https://doi.org/10.1037/bul0000115.

Cite this article: McCaughey, L., Ziegler, J., and Fiedler, K. (2026). Adapting to information search costs in sample-based decisions. *Judgment and Decision Making*, e20. https://doi.org/10.1017/jdm.2026.10043

Downloaded from https://www.cambridge.org/core. 16 Nov 2026 at 22:01:49, subject to the Cambridge Core terms of use.