Doctoral thesis submitted to
the Faculty of Behavioural and Cultural Studies
Heidelberg University
in partial fulfillment of the requirements of the degree of
Doctor of Philosophy (Dr. phil.)
in Psychology

Title of the publication-based thesis
Information Costs in Sample-Based Decisions

presented by
Linda McCaughey

year of submission
2023

Dean: Prof. Dr. Guido Sprenger
Advisors: Prof. Dr. Klaus Fiedler; Prof. Dr. Michaela Wänke

Table of Contents

Abstract

This dissertation examined whether people engage in information search adaptively and what processes contribute to this potential achievement. The four papers report experiments based on sample-based decisions featuring either financial information costs, or temporal information costs implemented in a speed-accuracy trade-off. The findings paint a mixed picture of adaptivity, which I assume to consist of at least two aspects: responding to the decision situation and changes therein through an a-priori planning process, and further improving efficiency through a fine-tuning process based on experience. Participants consistently adapted to explicit changes in the decision situation in a suitable way, mainly in the financial-cost experiment, but also in the temporal-information experiments. However, when it came to further improving efficiency beyond this initial adaptation, participants showed profound difficulties, which were demonstrated in a persistent accuracy bias in the speed-accuracy trade-off experiments and a lack of improvement in the information-cost experiments. Metacognitive monitoring and control are presumably fundamental to making this experience-based improvement possible, but participants could not make use of the assistance provided to metacognitive monitoring in either paradigm. This leads to the conclusion that further improving efficiency to achieve adaptivity is not a straightforward matter. What the exact requirements are and where they are instantiated in the real world to make adaptivity possible are important questions for future research.

Adapting to information search costs in sample-based decisions

Whenever we are faced with a decision, the first thing we want to do is search for information. Collecting information about the options that we have to decide between makes the decision easier and more likely to yield a satisfactory outcome. Despite these benefits, we generally tend not to spend the majority of our time on information search. Presumably, we are also acutely aware of the costs of information search. We might have to spend money on acquiring information, and we most definitely have to spend time on it. That time and money might be better spent on something else – a fact we seem to be aware of. Considering the plethora of decisions that we make each and every day, this is a highly important skill for managing our every-day life. Most of the time, it helps us avoid getting stuck in information search when new information is unlikely to improve decision making just as much as it helps us to avoid making important decisions too. Although everyone can think of a regrettable decision that was made on the basis of too little information or after waiting for too long, the fact that we quickly run into difficulties when trying to come up with further examples seems to indicate that, all in all, we are doing alright. But how good are we really and, more interestingly, how do we achieve this feat?

This dissertation strives to provide new insights on these important questions based on four publications. Implementing time costs of information in a speed-accuracy trade-off, my co-authors and I were able to observe participants' struggle with overcoming their focus on accuracy, reflected in their persistent spending of too much time to collect information before deciding (Fiedler et al., 2021). Despite multiple interventions intended to assist participants in improving, they maintained their accuracy emphasis. We came to the conclusion that the task did not provide enough scaffolding for effective metacognitive monitoring and control, precluding improvement, especially in combination with the task-inherent imbalance in accuracy versus speed emphasis. In the following paper, my co-authors and I reported an experiment that implemented a new approach to helping people improve at the speed-accuracy trade-off task (McCaughey, Prager, & Fiedler, 2023a). Through competition with a virtual rival or observation of a virtual teammate, participants experienced or observed a speedier strategy based on less information, which was intended to aid the evaluation of that strategy as superior to a slower one based on more information. After the interventions, participants did implement a slightly faster strategy than a control group. However, whether this was a result of their successful metacognitive monitoring and control rather than a type of anchoring effect is doubtful, since participants still believed a moderate rather than a fast strategy to be the best approach.

The following publication (McCaughey, Prager, & Fiedler, 2023b) contrasted the experiments using temporal information costs with financial-cost experiments, which have a relevantly different structure: The financial cost of information was much more direct and less abstract than the temporal cost, because it was directly subtracted from the decision payoff in each trial. Accordingly, the costs and benefits of information were not at such an imbalance as in the speed-accuracy trade-off. This means that participants should exhibit a less pronounced or no accuracy bias at all in the financial-cost decisions. Across four different ratios of information cost and payoff, we indeed found that participants acquired both too much information for some ratios as well as too little information for others, without exhibiting a general tendency towards collecting too much information.

In the final paper (McCaughey, Prager, & Fiedler, 2023c), we expanded on these findings by focussing on the adaptation to different information-cost-payoff ratios by examining the role of two potential underlying processes. The results indicated that participants relied on a-priori planning of strategy change in response to the announcement of the new ratio. To a much lesser extent did they use their experience of this ratio and the feedback about their decisions to further fine-tune their strategy during the block, with the presence or absence of feedback making no difference to the extent of strategy changes within blocks. To exclude the possibility that the standard trial-by-trial feedback was too complex to use for metacognitive monitoring, we also provided participants with summarised feedback and a better task structure . However, even across many more trials than previously, participants did not or were not able to use either type of feedback to improve their performance.

The four papers invite the conclusion that humans are adaptive to a certain extent, perhaps even though they struggle to use fine-tuning processes to further approximate the most efficient resource use. Before elaborating on how exactly this conclusion can be derived from the findings in the four papers, a clarification of the concept of adaptivity along with related concepts and literature is in order.

Theoretical Background

In decision making research, adaptivity is most often appealed to in the field of heuristic decision strategies. While the resulting biases were the focus of research at first (e.g. Tversky & Kahneman, 1974), the value of these decisional shortcuts or rules of thumb was soon also acknowledged. Payne, Bettman, and Johnson (1988) showed that people used decision strategies that struck a good balance between the amount of information they used and the accuracy they achieved. The implemented decision strategies were also suited to the respective informational environments. For example, when attributes were differently predictive of the best option, with one attribute particularly predictive of the best option, participants tended to use a decision strategy based on only that attribute. Payne and colleagues termed the adaptation of decision strategies to the characteristics of the decision environment adaptive decision making.

Adaptivity is also closely related to the concepts of bounded rationality and ecological rationality. Herbert Simon (1956) introduced bounded rationality as an alternative to classic normative theories that takes into account the constraints faced by human decision makers. Those constraints apply to cognitive capacity, time capacity, and epistemic possibilities, because we cannot process unlimited amounts of information, we do not have unlimited time for decisions, and we often only have limited information, with some information assumed to be available by normative principles actually impossible to attain in practice. Behaviour can be seen as boundedly rational when it approximates normative principles as best as possible taken into consideration these constraints.

Building on bounded rationality, the concept of ecological rationality emphasises the fit between the behaviour and the environment (Goldstein & Gigerenzer, 2002; Todd & Gigerenzer, 2007). Although introduced as a general concept, its meaning becomes clearest when applied to heuristics. According to Todd and Gigerenzer (2007), the information in certain environments is structured in a way that can be exploited by heuristics. It is ecologically rational to use the heuristic decision strategy that does well given the environmental structure. Each heuristic will do well in some environments and poorly in others, so assessing to what extent participants apply suitable heuristics is a good measure of how ecologically rational they are.

The bigger picture is that human decision makers, who face various constraints, need to be able to find a good balance between expending resources and getting a good outcome. In fact, that is exactly how Lieder and Griffiths (2017) argue that humans learn to apply different heuristics in different environments, by learning whether the outcomes were worth the resource expenditure. This account solves the problem of how people choose which heuristic to rely on in different situations, also called the strategy selection problem (see also Söllner & Bröder, 2015) and suggests a more general principle: That people, through experience with their environments, learn to achieve the best outcome possible with as little resources as possible; they learn to strike a good trade-off between the costs and benefits of information.

To me, this is the core of adaptivity. Adaptive behaviour is behaviour that leads to a moderately good outcome without investing too many resources, and does so by taking into account the relevant features of the environment, that is, the features that influence the outcome. This type of adaptivity is not only important in decision situations in which options are characterised by different attributes described by summarised information or numeric values, since these are not the only types of decisions that we encounter and have to cope with. According to the cognitive-ecological approach (Fiedler & Wänke, 2009) and its precursors (e.g., Hammond et al., 1986), we often have to make decisions and judgements about entities that we cannot observe directly. The only way of doing that is by collecting observations of a proxy as indicators of the underlying variable that is of interest to us. For example, trying to determine whether we will like a movie or not before going to see it at the cinema, we can inquire about it with our family and friends. Deciding whether or not to invite a new colleague to our birthday party, we have to rely on the interactions with them and form an impression of them. These decisions are made based on a sample of observations of the underlying variable. In line with the ecological-cognitive approach, I will call this type of decisions sample-based decisions. Although different variants of sample-based decisions have been studied for decades, a particular type has recently experienced a resurgence, where instead of gambles with described probabilities and outcomes, participants need to experience the outcomes and their probabilities (Hertwig et al., 2004). The researchers engaged with this type of decision rightly argue that we rarely encounter decisions that are like described gambles. Much more often, we do not know the probabilities of the outcomes or even all outcomes, making this a more naturalistic decision type to study. I have focussed on a different, more simple variant of a sample-based decision that has many similarities with the decision studied in the Bayesian tradition of the 1960s (Slovic & Lichtenstein, 1971). In the two-sample version, participants have to decide which of two options, each described by a sample of binary observations drawn from a Bernoulli distribution with a certain probability p of a positive outcome, has a higher value of that probability p. In the one-sample version, participants have to decide whether the p for that sample's distribution is above or below .5. Each observation is represented by a green arrow pointing upwards or a magenta arrow pointing downwards to indicate a positive or negative outcomes, respectively.

In this and in most sample-based decisions, participants are directly faced with the decision of how much information they want to acquire before making a decision. In attribute-based decisions, the domain of the typical heuristic decision strategies, most accounts assume that participants choose a particular heuristic for the decision situation and then apply it, which results in a certain information amount acquired as decision basis (most prominently Gigerenzer & Todd, 1999; but see Söllner & Bröder, 2014, or Lee & Cummins, 2004, for alternative views). In sampling-based decisions, people have to decide more directly what amount of evidence or information they want to acquire before making a decision. It is generally assumed that they use some type of accumulation-to-threshold process to tackle such tasks (e.g., Gluth et al., 2012). This means that they collect evidence until they reach a certain threshold that leads to a decision for the option currently favoured by the evidence. The evidence is what the integrated information results in and can take various forms, including Bayesian evidence (Edwards, 1965). Bayesian evidence accumulation up to a threshold is equivalent to the sequential probability ratio test (Wald & Wolfowitz, 1948), shown to be the optimal way of approaching decisions in which the amount of information is self-determined, which is why the optimal strategy used as a comparison standard and detailed in the papers takes this form.

Independently of how well these models describe actual behaviour, participants make a more direct decision about how much evidence or information to collect before deciding in sample-based decisions. As hinted at above, there are also accounts that claim that heuristics are selected according to that very same criterion, that is, by the evidence threshold people want to reach (Söllner & Bröder, 2014). However, heuristic strategies cannot all be transformed into each other at every point, so the evidence threshold might guide strategy selection beforehand, but it does not actually determine the termination of information search – the heuristic decision rule is responsible for that. However, this potentially stronger similarity between attribute-based decisions and sample-based decisions makes it even more pertinent to investigate sample-based decisions, as the insights gained might also be generalisable to other decision types.

Although adaptivity is fundamentally different from classic normative principles, they are nevertheless types of normative principles that play a valuable role in examining adaptivity. The classic normative strategy for attribute-based decisions is what Payne et al. (1988) called the weighted-additive strategy that takes into account all available information. The equivalent for sampling-based decisions would be to look at the entire population instead of just a sample from it, which is absurd. It shows how unrealistic normative principles are if they do not take into account the resources that need to be expended to attain outcomes of a certain quality (e.g., a certain probability of a positive outcome). However, one can also develop optimal strategies that take into account these constraints. In line with the core idea of adaptivity, an optimal strategy is one that maximises the efficiency of resource use by achieving the highest accuracy (or probability of a correct decision) with as little resources as possible. This may sound vague, but as soon as the information cost and the decision payoff are specified, there is a point (or multiple) at which the average payoff minus the information cost is highest. The visual depiction of the accuracy as well as the efficiency people can achieve with a certain sample size (see Figure 1) are an intuitively accessible portrayal of the relation. With the sample size on the x-axis and the mean accuracy of the decision plotted on the y-axis , one can clearly see that the increased sample size causes increasing accuracy in the left panel. However, it is not completely clear at which point the resources (sample size and the costs involved) are used most efficiently. For this, one needs the second scatterplot in the right panel, which displays the efficiency (or balance consisting of the decision payoff minus the information cost) as a function of sample size. This clearly shows which sample size and accuracy combination represents the most efficient use of resources, which is equivalent to payoff maximisation if decision outcome and information costs are framed in monetary terms.[1] Crucially, this kind of

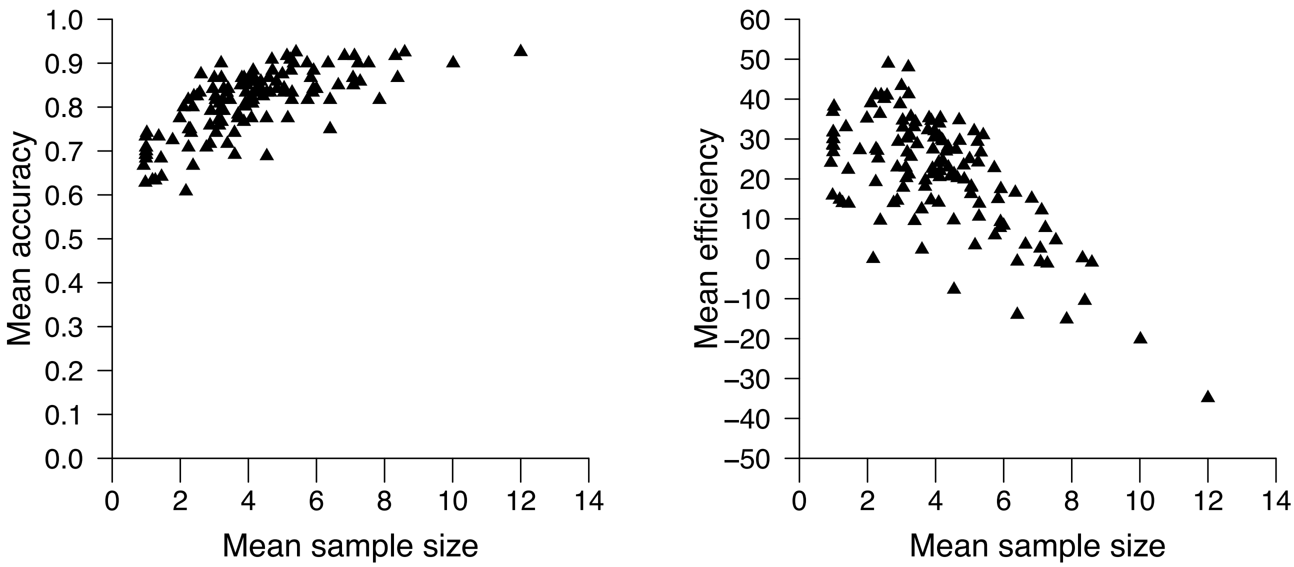

Figure 1. Based on participants' data from Paper 4, Experiment 3: Left panel shows participants' mean accuracy as a function of their mean sample size (aggregated across trials), while the right panel plots participants' mean efficiency as a function of mean sample size

optimal strategy is not always the same, but will change with the environment. Some changes in the environment will affect the right graph: changes in the ratio of information and decision

---

[1] What the graph indicates as best trade-off is, of course, always constrained by the data that is displayed. In this case, the data is actual participant data from Paper 4, Experiment 3, so the "optimum" will only be as good as the best participants. Simulated data can reveal the optimal trade-off.

payoff, for example. Some environment changes, such as the difficulty of a decision, will affect the relation between sample size and accuracy displayed in the left graph, thereby also influencing the relation between sample size and efficiency.

With adaptivity defined as efficient use of resources, there are two aspects one can focus on. First, one can focus on the extent to which people achieve efficient resource use in a given decision situation. A related, but not necessarily identical question is to what extent people succeed in changing their decision strategy in response to a change in the decision situation or environment to maintain efficient resource use across different decision situations.[2] The extent of people's adaptivity is only one of two questions I aimed to investigate. The second question is how people manage to be adaptive, if they are, or what prevents them, if they are not. It is not realistic to assume that, as humans, we just know precisely which amount of resources is most efficient to apply in all decision situations. Hence, for adaptivity to be possible at all, there must be some process of acquiring the knowledge, at least for the decision environments that we face most frequently. This is where the metacognitive perspective becomes relevant. Metacognition in the case of striving for resource efficiency refers to the cognitive processes that occur at the superordinate level (Fiedler, Ackermann & Scarampi, 2019). While the basic cognitive functions are responsible for information search, integration, and arriving at a decision, the metacognitive processes should monitor one's efficiency at the decisions and use control to improve it (Nelson & Narens, 1990). Efficiency is the quality of the decision outcome (e.g., accuracy) put into relation with the resources one needed to achieve it. Monitoring the efficiency of different decision strategies, for example, different evidence thresholds, should allow to identify the most efficient, with control functions being applied to execute it.[3] What form these metacognitive processes take more specifically and whether they are successful was a question investigated as part of this dissertation.

With the major role that information costs play for adaptivity, I investigated the questions above using two different types of information cost, financial and temporal. Since time passes relentlessly and cannot be accumulated or 'saved up' for later use, temporal costs of information are often implemented as opportunity costs of time, meaning that the cost of

[2] Of course, if people were always perfectly efficient, they would also respond perfectly to changes in the decision situation. However, as soon as they deviate from the most efficient resource use, the two issues are no longer the same.
[3] My use of strategy is broader than in the heuristic decision domain, where different strategies are characterised by different steps. Here, different strategies will also refer to decision processes that integrate the information the same way but have a different stopping point because one had a higher and the other a lower evidence threshold. While there might be different ways of information integration as well, I ignore that and look only at the resulting mean trade-off between resources and accuracy or decision quality.

spending time on something is the value of the alternative one could be spending it on (Spiller, 2019). This is not always straightforward to operationalise. A common and useful way of implementing a very specific opportunity cost of time is by setting up a speed-accuracy trade-off. For this, the total time of an experimental session needs to be fixed, and participants need to be advised, and ideally incentivised, to work through as many trials as possible with the highest accuracy possible. The task in each trial needs to be such that spending more time on it leads to better outcomes. In the papers described in the following, this was achieved by letting observations, represented by arrows, appear one by one, so participants had to wait longer to receive more information. That way, participants had to trade off the accuracy they achieve on individual trials against the number of trials they could get through – the speed component. The cost of information is the value of the alternative use of time, which is defined quite precisely in this paradigm. The time spent on waiting for additional information could have been spent on working through additional trials instead. If information and time are proportional to each other, assuming one spends a certain amount of time on each decision to acquire a certain amount of information, one could double one's trials by halving the amount of information per decision. Since the relationship between sample size and accuracy is sub-linear even for optimal strategies (when aggregated across trials; see participant data depicted in Figure 1), for most sample sizes larger than 2 or 3 it makes more sense to spend time on increasing one's accuracy for a new decision instead of the current one, because the initial increase is always steeper (assuming moderate presentation time and symmetric payoff among other things).

Financial information costs are easier to set up. Participants simply have to pay for each observation they acquire. In our case, they are told beforehand what the information costs and the payoff from correct and incorrect decisions are and, in most cases, also receive feedback about whether they made a correct or incorrect decision and how that impacted their current payoff along with the information costs. This simplicity also makes it very easy to change the ratio of information costs and payoff. The ratio has an enormous influence on the point at which resources are used most efficiently. If the stakes are high because the payoff is high while the information is relatively cheap, efficient use of resources would mean to collect more information. If the stakes are low, on the other hand, because the payoff is low while information costs are high, then efficient resource use would be reflected in sparse information search.

The first two papers of this dissertation are based on a speed-accuracy trade-off paradigm to investigate adaptivity to temporal information costs. The other two involve financial costs to better be able to investigate participants' adaptation to changes of the environment and to compare the speed-accuracy trade-off with a different task set up that has a relevantly different structure for metacognitive monitoring and control.

First paper: Speed-accuracy trade-offs in sample-based decisions

This first paper aimed to assess participants' adaptivity to temporal information costs in a speed-accuracy trade-off with sample-based decisions. Although it has some characteristics of real-life situations, with sample-based information and time costs, participants are not usually faced with this type of task. Hence, we did not expect participants do be able to adapt to it immediately. Rather, we expected them to be able to improve their resource use to become more efficient through experience with the task. However, the extreme emphasis on accuracy that participants displayed did not improve in any of the four experiments, which were designed to test different forms of assistance. The biggest change was brought about by implementing different payoff structures, which also changed the optimal trade-off substantially. Our focus at the time was not on whether participants would change their behaviour in response to the different optimal trade-offs, but whether that experience would help them understand that they generally tended to wait for too much information and therefore help them come closer to the optimal strategy. Interestingly, participants did change their strategy and the resulting sample size in the right direction. When the payoff consisted of a one-point gain for each correct decision, but no loss for incorrect decisions, participants decreased their sample size. However, relative to the optimal strategy, their efficiency actually worsened. On the other hand, when the payoff consisted of a one-point gain for each correct decision and a three-point loss for each incorrect decision, participants exhibited even larger sample sizes. While the direction of their sample size difference was the same as for the optimal strategy, participants' mean sample size for the previous, symmetric payoff was already large enough for this asymmetric loss-emphasising condition. So by changing in the same direction as the optimal strategy, they also shifted away from the optimal strategy. This demonstrates that there are two sides to adaptivity, which will also play a role in the other papers. On the one hand, participants successfully adapted their sample size in the correct direction in response to changes in the payoff structure, so they understood what the consequences of the payoff change were for the strategy conceptually. On the other hand, they continually overemphasised accuracy and exhibited a sample size much larger than necessary and did not succeed in improving even in the slightest. So while their change in response to the change in payoff structure would have been close to optimal if their sample size was generally smaller, they just do not seem to be able to gain the insight that a more efficient resource use would be achieved by generally smaller sample sizes. From a metacognitive perspective, participants seemed to be having serious difficulties in either metacognitive monitoring, control, or both. Both are necessary to improve one's strategy at the task. Metacognitive monitoring keeps track of one's performance, which can be used to register improvement or worsening, or to compare different strategies to each other. Metacognitive control instructs behaviour and can therefore implement the changes suggested by the results of monitoring. Failure of either of these functions is possible, but they are not equally likely, in my opinion. Since participants have demonstrated that they can implement strategies based on or resulting in a different sample size, I do not find it plausible that participants recognise what an improvement in their strategy would be and simply fail to apply that knowledge. It is more likely that the monitoring function is not supported in a way conducive to a correct evaluation of participants' actually and potentially applied strategies.

One potential problem for participants' monitoring might have been a lack of ability or motivation to fully monitor their performance. I am not suggesting that they were not motivated to engage or perform well at the task. But with a strategy based on a fairly large sample size, the feedback that one receives is positive to a high degree because one has a high accuracy rate. The opportunity cost of spending so much time on individual decisions can easily be ignored, with monitoring focussed completely on the accuracy dimension of the task. And if participants were under the impression that they were doing pretty well at the task, they would have seen no need to try out other strategies and monitor them comparatively. Hence, they would not have had the basis to recognise that faster strategies are more efficient, happily proceeding with their overly accuracy-focussed strategy throughout the whole experiment. As a remedy to this, we identified what we believed to be a promising way of demonstrating the benefits of a faster strategy while also motivating participants to monitor their performance on both dimensions of the trade-off, accuracy and speed. This remedy consisted of participants competing against a virtual rival or observing a virtual teammate at the speed-accuracy trade-off task.

Second Paper: Rivals reloaded – Adapting to sample-based speed-accuracy trade-offs through competitive pressure

In the second paper, my co-authors and I report an experiment that featured an extension of the speed-accuracy trade-off task by either a virtual rival, a virtual teammate, or both. The purpose of these extensions was, in the rival condition, to make participants experience a faster strategy by competing with the rival, who was programmed to pursue a strategy close to optimal. The rival was incorporated into the trade-off task in such a way that the participant and rival saw the same sample unfold, observation by observation, and whoever indicated first that they wanted to stop sampling and make a decision was granted that opportunity while the other player had to watch and received nothing. Instead of competing with the rival themselves, participants in a second condition observed their virtual teammate compete with the virtual rival according to the same principles. In yet another condition, intended to exclude the competitive element, participants observed their virtual teammate go through the standard task, without competing against anybody. Finally, a control condition just played the standard speed-accuracy trade-off task. All conditions went through two blocks of the speed-accuracy trade-off task, with all experimental manipulations placed in the first block, followed by a second block of the standard task. The rival did indeed have the effect that we had hoped for, at least while it was present: Participants reduced their sample size substantially in Block 1 to avoid missing out on decisions too often. This gave them the opportunity to evaluate a strategy that was much faster than what they would have typically engaged in. However, in Block 2, participants' sample sizes were only a little smaller than the control group's and still very far from the optimal sample size. Moreover, they still reported the belief that a moderate strategy is more efficient than a fast strategy. This makes it unlikely that participants actually gained the insight that a fast strategy is the most efficient. They may, however, have realised that a somewhat faster strategy than what they would usually implement is more efficient than a slow strategy. In contrast, the control condition rated the slow and moderate strategies as equally efficient and a fast strategy as less efficient, so perhaps the lack of fast experience led them to maintain the belief that moderate and slow strategies are equally efficient. The causal direction might just as well be reversed, however. Participants may have been led to use smaller samples due to a type of anchoring effect on the strategy they used or saw previously and then reported the strategy they used themselves as most efficient. In any case, they did not arrive at the insight that fast strategies are most efficient.

The metacognitive perspective can offer explanations for why participants did not gain as much insight as we had anticipated. As argued above, monitoring is the likely weak spot in participants' efforts to improve. This might have been particularly difficult for participants competing against the rival. While we assumed that executing a faster strategy oneself would be beneficial for improving because one is experiencing it first-hand and is extremely motivated to perform well and beat the rival, this pressure to compete with the rival may also have made monitoring more difficult. There may only have been limited capacities left to monitor performance. However, participants who observed their teammate compete with the rival should have had enough cognitive resources to monitor the performance because they did not have to make any decisions themselves. Nevertheless, participants did not improve as much as we had hoped.

An additional issue is that there were very big experimental changes from Block 1 to Block 2 – the rival was no longer playing. Whatever participants may have inferred about a fast strategy, they may have only attributed it to the task version with the rival. Once pre-empting the rival was no longer a prerequisite for having an opportunity to attain a gain at all, participants may have inferred that they should change their strategy. Likewise, they may have believed the rival to be following that strategy only because of the competitive situation.

Nevertheless, not even the condition that simply observed their teammate execute a close-to-optimal strategy at the standard, non-competitive task arrived at the insight that a fast strategy was the most efficient. In fact, one could have expected them to imitate the strategy just because that is what they were shown and might have inferred they are supposed to act the same way, even more so since the virtual agent was designated a teammate. Participants nevertheless preferring to pursue a different strategy indicates quite a strong conviction on their part that it is in fact an efficient strategy.

Since monitoring seemed unsuccessful in all of the speed-accuracy trade-off experiments, the cause of the difficulty must be related to general aspects that remained unchanged. A systematic examination of the steps necessary for improvement should help to pinpoint potential causes of monitoring difficulties.

In order to improve one's performance at the trade-off, one needs to increase the efficiency. The efficiency is the number of correct decisions per time unit. It reflects how accuracy and speed are linked in the task and makes different ways of trading off speed and accuracy directly comparable. The structure of the trade-off task is such that the balance, which is the cumulative sum of the decision outcomes, indicates the extent to which the resources are used efficiently. So how can efficiency be improved? Participants need to monitor their accuracy rate and the number of decisions, and put that in relation to the time that has passed. Since they also see their balance, as an alternative to monitoring accuracy and speed and integrating them, they could monitor the balance in relation to the time directly.

The reason participants may have difficulties with this is related to the basic structure of the task and was briefly mentioned above. Although the name may suggest it, speed and accuracy are not equally salient when the task includes trial-by-trial feedback. While feedback about accuracy is the standard way of providing performance feedback, in this task it actually only covers one of two aspects that jointly determine performance, if performance is understood as efficiency. By continually providing feedback about decisions' correctness, participants may be focussing on monitoring only accuracy instead of both accuracy and speed, or balance, directly.

This is further exacerbated by what Hsee and Zhang (2010) call evaluability. According to them, attributes can be inherently more or less difficult to evaluate depending on whether people have an innate reference system to compare the attribute values against. Ambient temperature is one example they give for an inherently evaluable attribute, because we can evaluate how pleasant a temperature is for us without any further information or experience. Many other attributes are harder to evaluate, lower in evaluability, and can only be meaningfully evaluated in comparison with other values on the attribute or after extensive experience. For example, whether a certain car is expensive or not can only be answered in comparison with other car prices. At first, decision correctness seems like a prime example for an attribute that is inherently evaluable: correct decisions are good and incorrect decisions are bad. However, Hsee and Zhang (2010) consider an attribute inherently evaluable only if no learning is necessary to be able to evaluate it. Hence, the case of decision correctness may be more debatable than we thought in Fiedler et al. (2021). What is most relevant about this is not whether decision correctness is more appropriately categorised as inherently evaluable or not. It is the fact that decision correctness seems to fit the category of inherently evaluable so well, although some learning or socialisation is presumably required for this. Our very own task demonstrates that correctness is not always good, because a certain number of incorrect decisions is required for efficiency. This makes its evaluation rather difficult. However, this is a difficulty that is hard to accept. Our socialisation of correct decisions and accuracy as good seems so strong that taking a step back from evaluating it in itself, but rather in relation to the resources required is very difficult. One could call this case misleading quasi-inherent evaluability. The ease with which accuracy is evaluated as good may pose a serious challenge for engaging in accurate monitoring of different aspects of efficiency, since accuracy may attract more attention or be weighed more heavily. The ease may be related to the large predominance of accuracy as sole performance indicator in society, with the majority of school and university exams, competitions of all kinds and job performance indicators all strongly concerned with accuracy, and only paying attention to speed when it tends towards the extremes. For example, no one takes into consideration how long a gymnast trained to achieve a certain quality of performance, the quality of the performance is the only thing that counts. A little less exclusively focussed on accuracy, exams always have a time limit, but within that time limit, only accuracy or quality are graded, without further taking into account time differences. Moreover, how long students prepared for the exam is not taken into account at all.

The quasi-normative goodness of accuracy may not be the only contributor to participants' accuracy focus. From a perspective of weighing benefits and costs of information and striving for a good balance, the benefits of information are extremely clear. More information leads to correct decisions, the consequences of which are emphasised after each trial. The costs one has to accept in return are indirect. They consist of lost opportunities for additional gains, because one will not be able to make as many decisions. However, since one would have to collect less information to be able to make additional decisions, participants may even be uncertain of whether attaining more decisions is a good thing, since they have also experienced the loss that follows an incorrect decision. Probably more pertinently, however, it is easy to neglect or ignore that indirect cost and convince oneself that one is pursuing an efficient strategy, especially since no explicit feedback is given on the cost of information.

Because the different degrees of directness of the benefits and costs are an inherent part of the temporal costs as opportunity costs, designing feedback for the indirect cost of information that is as salient would be very difficult. Hence, we chose a different approach and designed experiments that used financial costs instead of temporal costs.

Third paper: The information cost-benefit trade-off as a sampling problem in information search

This publication links together two types of information cost. It summarises the speed-accuracy trade-off experiments and contrasts them with two new experiments that use financial information costs. Financial costs have the great advantage that they are in the same currency as the gains and losses that constitute the decision outcome, which I refer to as payoff. This means that the two can simply be summed to form the total outcome for that trial, which is equivalent to the efficiency. Since financial time costs can be directly subtracted, there is no need for a total session time in which participants have to attain as many decisions as they think is reasonable. The experimental structure can simply be a series of trials. Accordingly, there should no longer be an imbalance between the costs and benefits of information. Both the advantages of information, a higher probability of receiving the gain, and the cost of information, the points deducted for each observation, are experienced directly for each trial. If anything, the information cost may be more salient, because it is deterministic, while the benefit of the information is only an increase in probability of getting the gain, but will also result in a loss at least some of the time. The directness of the information cost also means that simple, intuitively accessible feedback can be given on it, further balancing the perceived weight of costs and benefits of information.

If this imbalance was the main driver of participants' accuracy focus in the speed-accuracy trade-off, then there should be less of an accuracy focus in a paradigm that eliminated that imbalance. The experiments assessed this across multiple different amounts of information costs that were grouped into blocks. For the other advantage of financial costs is that they are very simple to change in a way that participants can follow and experience. Since the ratio of information costs to payoff determines how much information is most efficient to use, this is a valuable aspect to vary. Across four different ratios, the experimental results indeed showed that there was no general accuracy focus. However, the lack of an imbalance is not the only reason for those findings. What is also evident is that whether participants can be said to have acquired too much or too little information depends very much on the optimal strategy (see Figure 15.4 in Paper 3). The optimal strategy changes a lot more across the four ratios than participants' strategy, so the difference between them is mainly driven by the optimal strategy. Because our experiments included a whole range of ratios, we were alerted to this fact, making us cautious about the inferences we draw. If we had chosen only one or two of those ratios, we might have felt confident enough to draw very different conclusions. Looking at only the lower two ratios one would come to the opposite, seemingly well-supported conclusion as looking at the higher two ratios.

Looking at the whole range of ratios, there was no overall accuracy bias. In fact, the only condition in which participants sampled substantially more than the optimal strategy was for the ratio in which the optimal sample size amounted to a single observation. Although we do not want to understate this, it does seem understandable that participants had greater trouble approximating this optimal strategy just because of its extremity. Supporting this assumption is that the extent to which participants sampled too little for the ratio with the cheapest information (or highest payoff) is even larger. If they were generally prone to acquire too much information, one would expect them to find it easier to sample more when the ratio required it for the sake of efficiency. However, the adjustment was similarly small, resulting in a sample size much smaller than optimal. A factor that differs between sampling too much and sampling too little, however, is the additional cost of time that participants were facing. Getting through the experiment more quickly would have allowed them to spend time on something else. This was counteracted to some extent by the incentivisation, but it was still an asymmetrical influence. When sampling too much, participants accepted a higher information cost than would have been most efficient and, in addition, accepted a higher time cost than would have been necessary. When sampling less than the optimal strategy, at least some of the deviation can be explained by participants taking into account the time costs that the optimal strategy could not take into account.

Hence, the accuracy focus seems to be related to the structural emphasis on accuracy in the speed-accuracy trade-off in two ways. First, it may heavily contribute to the initial accuracy focus that participants exhibited consistently. Second, and more importantly for adaptivity, it may also be a good explanation for why participants did not succeed in improving their efficiency. Their focus on accuracy and the easier accessibility of accuracy-related information made it impossible to engage in balanced monitoring of both dimensions and efficiency as their combination. Since the structural imbalance is not present in the financial-cost paradigm, adaptivity and potential underlying processes should be easier to investigate in that paradigm. This was the focus of the final paper.

Fourth paper: Adapting to information search costs in sample-based decisions

The fourth paper extends the first two financial-cost experiments by another two that were designed with a stronger focus on the processes that might underlie adaptivity. It also elaborates on aspects of the first two financial-cost experiments that were set aside in the third paper. For the different ratios did not only serve the purpose of assessing participants' accuracy focus. They were also ideally suited to investigating the extent to which they adapted to a change in the factors that influence which the most efficient strategy is. This is of great relevance, since different decision situations are often characterised by different stakes and different accessibility of information, that is, by different payoffs and information costs, which determine how much information is most efficient to acquire. Being able to adjust to different situations is a core element of adaptivity.

The first experiment had participants face four blocks, each with a different ratio. The ratio changes were implemented either through a change in the information cost or a change in the payoff to make the next block easily comparable to the last. The results showed that participants were indeed sensitive to those changes and adapted their sample size, which was also reflected in accuracy changes. The question that arises is how participants achieved this adaptation. Since the blocks were easy to compare to each other directly, it is likely that the adaptation was based on an a-priori plan of strategy change in response to the announcement of the new cost ratio to a large extent. A second process that should be at work is the direct experience of the ratio during the decisions, and the adjustments that are made in response to that experience mediated by the feedback. For example, if one has just worked through a block in which each observation cost 10 points and the payoff was 200 points (ratio 1:20), and moves to a block in which the payoff changes to 100 (ratio 1:10), one should quickly realise that pursuing the same strategy as before leads to a larger proportion of the gain being spent on information cost, or in case of an incorrect decision, that the total loss that one faces consists of information costs to quite a large extent. We have called this process fine-tuning of one's strategy because the changes are likely to be more gradual, with people perhaps trying out slightly different strategies, monitoring them, and evaluating them to improve their strategy.

In order to assess to what extent the first a-priori planning process is at work, the second experiment also featured changes between four blocks. This time, however, both components of the ratio changed at the same time and participants additionally had to work on a distractor task during the break between the blocks. These measures were intended to make the comparison between the block ratios more difficult. This increased difficulty was indeed reflected in the data. Participants were still sensitive to the ratios, but a smaller proportion of participants managed to change their sample size in the same direction as the change in the optimal strategy. This indicates that the adaptation did happen based on a-priori planning of strategy change to a large extent, while the fine-tuning contributed less to adaptation.

This is in line with the comparison to optimality: while generally sensitive to the ratio, participants did not adapt enough quantitatively. This might be precisely because of how they adapt. The right direction of change is easiest to infer from the comparison of blocks. Although more difficult, this was still possible in the second experiment. To improve efficiency beyond that initial change, one would have to rely on one's experience and feedback as part of the fine-tuning process. To further examine to what extent participants engaged in fine-tuning, we conducted a third experiment in which we varied whether participants got feedback or not. The experiment consisted of three blocks, of which the first and third had the same ratio. This means that the second change was in fact a change back to a ratio that participants had already experienced. To make sure everyone had a chance to use feedback to improve their initial strategy, all conditions received feedback during the first block. Only the feedback conditions continued to receive feedback in the second and third block.

The findings were surprising in two respects. First, the feedback did not make a difference. Participants did not have significantly different sample sizes depending on whether they received feedback or not. This was very unexpected, since we assumed that the feedback was essential to participants' adaptation to the different ratios via the fine-tuning process described above. And while we suspected that this fine-tuning was not the main way through which participants adapted, we nevertheless assumed that the feedback had to be helping.

Considering this, what is perhaps even more surprising is that participants did increase their sample size over the course of the second block, coming a little closer to the optimal strategy. However, this increase was not influenced by feedback – participants who received feedback displayed the same increase, on average, as participants who did not. This is puzzling. Perhaps the condition without feedback simulated the information costs based on the announcement and the accuracy based on their previous experience with some upward adjustments. As Elwin et al. (2007) showed, under certain circumstances participants can do just as well without feedback as participants with.

Alternatively, perhaps neither of the conditions used feedback and the sample size increase was just a delayed effect of the a-priori planning of change in response to the announcement of the new ratio. Perhaps participants had trouble monitoring their performance despite the more balanced structure and, accordingly, more balanced feedback. Although the feedback was more suited to monitoring than in the speed-accuracy trade-off experiments, there might still have been some unsurmountable challenges for participants. First, the feedback was still provided on a trial-by-trial basis, which means that it would been subject to a certain fluctuation because decisions could only be correct or incorrect. Hence, participants still needed to average the provided information over a few trials to get a good estimate of their current performance. For this to actually yield a reliable estimate, however, participants also needed to consistently pursue the same strategy over a defined period of time and bring the performance estimate in connection with that strategy. However, this would only yield a single evaluation, which would not be very evaluable on its own. One would need to compare it with other strategies. For example, if one were to pursue a strategy based on evidence accumulation to a certain threshold, one would have to try out different threshold settings that would result in higher or lower mean sample size, to be able to recognise which of the strategies yields the best efficiency. This means that participants needed to deliberately try out different strategies, average the feedback over the time periods they were trying them out, keep those evaluations and the strategies from which they resulted in mind, and then compare those evaluations. Clearly, there are still a few steps that could have gone wrong even if the feedback was more balanced than in the speed-accuracy trade-offs.

To test whether a different type of feedback combined with a different trial partitioning would help participants fine-tune their performance, we ran a final experiment. Instead of testing different ratios, we focussed on just a single ratio across all trials. Those trials were partitioned into blocks of ten trials each. This was intended as opportunity for participants to try out different trade-offs in segments that were clearly separated from each other. One condition received the usual trial-by-trial feedback. The other condition, however, received only a summarised feedback at the end of each block, that is, after ten trials. This should have made it very easy to compare the different blocks to each other. In addition, the lack of trial-by-trial feedback should have made it easier to focus on pursuing a consistent strategy and evaluate it at the end.

However, the results of the experiment showed no difference between the types of feedback. Moreover, there was no indication of a sample size change at all, not to mention improvement of efficiency. No matter whether participants received trial-by-trial or blocked feedback, they did not change their sample size. Hence, although the blocks and the blocked feedback would have provided the ideal ground for trying out and evaluating different strategies in order to identify at least slightly better strategies, participants did not seem to even try out different strategies.

One might argue that the number of trials did not suffice for participants to improve. However, if they were to improve with more time and more trials, there should already be some indication of that in the first 100 trials, even if it did not yet come to fruition. But if they tended not to try out different strategies in ten blocks of ten trials each, it is hard to imagine that they should suddenly start doing so with more than 100 trials.

The paper concluded with a mixed message on adaptivity. While participants adapted fairly readily to changing characteristics of decision situations to a certain, small extent, it was not possible to bring about improvements beyond that, which would have brought them closer to using their resources efficiently. It is important to remember here that the optimal strategy is not one that is ignorant of information costs as most classic normative principles are. The optimal strategy for the speed-accuracy trade-off and the information-cost-benefit trade-off are both specified such that they use the resources most efficiently, with efficiency being the hallmark of adaptivity. So participants not being able to further approximate this standard means that they fail to improve the efficiency of their resource use. Understanding adaptivity in terms of efficient resource use, the implications of these findings for human adaptivity need to be discussed.

General Discussion

The conclusion that can be drawn from these four papers is that people are adaptive to a certain extent. They plan changes that lead to behaviour that is more efficient in new decision situations. However, these changes are rather small, with large deviations between their and the optimal sample size. Since we do not expect people to analytically derive the optimal trade-off in advance, this rather small initial change can be fully consistent with adaptivity. However, one would expect to see at least some sign of improvement in efficiency based on experience and feedback, especially after presumed challenges had been reduced. The fact that we did not observe improvement in our experiments is a more reliable indication that adaptivity might be impaired. Since the experiments used fairly artificial tasks, which were intended to emulate only a few aspects of naturalistic decision situations, they do not allow for the conclusion that participants are not adaptive, because participants may very well display adaptive behaviour in real-life situations that they are familiar with. This is a clear limitation of the set of papers and experiments presented. However, assessing the adaptivity of people's behaviour in the real world was not the aim of these papers. They were designed specifically to assess adaptivity of behaviour in an unfamiliar artificial task, which had the potential to reveal the processes through which participants adapt to decision situations they are not familiar with. This is the strength of the experiments that represents the flipside of the limitation just described. If participants had improved their efficiency under some but not all conditions provided by us, we would have gained valuable insights that could not be provided by naturalistic decision situations. Since participants only adapted to changes and did not improve their efficiency, we could mainly infer that the prerequisites for further fine-tuning were not met, which is still an insight that naturalistic decision situations would not have made possible.

All in all, this may mean that participants need more assistance to arrive at adaptivity – or that they are not as adaptive as expected. However, it may help to take a step back and return to the notion of adaptivity to discover a different interpretation. The assumption of adaptivity is that people do fairly well without using many resources. In the findings presented here, this may not apply to the basic task level, since participants used too few or too many resources and did not improve the efficiency of their resource use. However, in the information-cost experiments, especially considering the balance participants earned, one must admit that they did fairly well and achieved a balance that was not much worse than the optimal balance. So perhaps we are dealing with a different kind of efficiency and adaptivity here. So far, I have been detailing the efficiency of resource use within the decision task. However, since we designed this task as a new decision environment which participants are not familiar with, they are faced with two different tasks at two different levels. There is the level of the basic task, for which one needs to use information efficiently to maximise one's reward. But there is also the superordinate task of how many resources one should spend on improving one's strategy at the basic task, which – as previously detailed – involves trying out different strategies, monitoring, evaluating, and comparing them. This is a superordinate exploration-exploitation trade-off. Normally, exploration-exploitation trade-offs are investigated with regard to such basic tasks as well. One needs to collect information about different options on the one hand, but also consider at which point just consistently benefiting from the currently best option is more profitable than further exploring a mixture of different options (see Mehlhorn et al, 2015 for an integrative review). Similarly, Harris et al. (2020) looked at how features of the environments motivated exploration versus exploitation of two options. In contrast, what might be implied by our experimental paradigm is an exploration-exploitation trade-off at a superordinate level. Instead of exploring or exploiting different options, one needs to explore or exploit different strategies for solving a decision task as efficiently as possible. On the one hand, further exploration brings the benefit of potentially finding a better strategy than the one currently pursued. On the other, exploration will also lead to worse strategies at least some of the time and in addition requires effortful strategy planning, monitoring and evaluation. Exploitation, in contrast, leads to a consistent performance or efficiency at the current level with no need for very much metacognitive effort. Therefore, one needs to trade-off the benefits of exploration against the benefits of exploitation with regard to strategy at a level superordinate to the basic task level.

What cannot be left unmentioned in this context is the distribution of efficiency that different strategies yield. There may be tasks where performing slightly differently to the optimal strategy implies a large drop in efficiency. For other tasks, trading off the costs and benefits differently than the optimal strategy may hardly decrease efficiency at all. The exploration-exploitation trade-off will also depend on what this distribution looks like. If the efficiency distribution of strategies or trade-offs is fairly flat, a slightly different strategy will not result in a much lower efficiency. In these cases, spending further resources on the improvement of one's strategy will probably not be outweighed by the benefits they result in.

If we look at the comparison of the participants' actual and the optimal balance in the financial-cost experiments, it is definitely the case that despite large deviations from the optimal strategy in terms of sample size, the efficiency (represented by the balance) is not much worse than optimal. Perhaps participants knew or assumed that their strategy is already fairly efficient and did not think it worthwhile to invest into exploration. However, this did not seem to be the case for the speed-accuracy trade-off experiments, where participants achieved a much lower payoff than optimal. Here, the structural imbalance between speed and accuracy and its reflection in feedback may come back into play. Above, I offered the speculation that participants in the speed-accuracy trade-off may be led to believe that their strategy is already pretty good because they focus too much on accuracy and disregard the speed component.

Perhaps this effect actually occurs at a superordinate level, with the imbalance leading participants to believe not that their strategy is good enough in itself, but rather, that the effort and other costs it would take to improve it are not worth the benefits one could gain from the better strategy.

If this interpretation turns out to be correct, people might be aware of resource use and its efficiency in ways not previously considered. Taking into account this superordinate exploration-exploitation trade-off implies a whole new perspective that, to the best of my knowledge, is not represented in the literature so far. Although Lieder's and Griffiths' (2017) work is focussed on resource efficiency and proposes this as principle underlying people's adaptive selection of heuristic strategies, they stay firmly rooted in the basic-task level of efficiency and do not consider the efficiency at the level of strategy exploration versus exploitation. Similarly, Khodadadi et al.(2014) introduced a model that describes how an agent may learn to set two different thresholds in an evidence-accumulation strategy for tasks of two difficulty levels. But the question of whether or to what extent investing the resources into learning these thresholds is worth the benefits is not considered. This potential new layer of adaptivity can lead to a deeper understanding of what processes are involved in achieving adaptivity over time and perhaps even explain why it may be adaptive to approach some tasks with less-than-maximal efficiency. Naturally, this perspective also brings new questions. How might people be able to estimate how much better a strategy can get? Perhaps, the environment is structured in a way that we can use a simple heuristic for this or perhaps people can infer efficiency distributions of different strategies based on the initial efficiency increases or decreases.

It is, of course, also possible that people are not adaptive at this level and that the findings in the financial-cost paradigms just happened to look like a fairly adaptive trade-off between exploration and exploitation. Which of these scenarios is more likely cannot be determined based on the data presented here and goes beyond the scope of this dissertation. Investigating this question systematically would be a fascinating and extremely important avenue for future research. A suitable approach would include both experiments using more naturalistic situations to assess participants' current adaptivity, but also further experiments based on precisely controllable artificial tasks, which allow to systematically investigate participants' behaviour in new situations. Ideally, a concerted approach would analyse naturalistic situations and our experiences with them, and systematically investigate them with artificial tasks by emulating specific aspects and assessing people's behaviour to discover the prerequisites of adaptive behaviour, both in terms of trading off costs and benefits of information in decision situations, but also in terms of trading off the costs and benefits of further improving a strategy, which is a much more complex matter. It would eventually have to include a perspective across the lifespan, addressing questions such as whether, with experience, we can learn the distribution of possible performance levels for different situations to be able to estimate the benefits of further performance improvement. Being able to estimate those benefits as well as the costs of further improvement, including both the effort and the relative loss through exploration, is necessary to be able to trade them off against each other. Alternatively, perhaps this research might show that instead of learning performance distributions for different situations, we use indicators for how much more there is to gain and at what cost it would come as a type of higher-order heuristic, which will often work well, but sometimes be misleading.

Advancing our knowledge in this way would be a huge step toward fully understanding human adaptivity, replacing the vague appeals to evolutionary selection processes somehow resulting in our being adapted to the environment we live in – most of the time. A theory of how adaptivity in information search and decision making develops would give us a deeper understanding of when it is reasonable for adaptive behaviour to fail and when not; it would also provide a good basis for developing effective assistance to the development of adaptive behaviour as well as direct assistance to decision making where we know even adaptive behaviour is prone to fail, for example, because cues that work most of the time are misleading in particular situations.

Conclusion

The summary and discussion of the four papers constituting this dissertation show that human adaptivity in information search and decision making is a complex but fascinating and important research area. Humans definitely show adaptive behaviour to some extent, with different processes presumably contributing to it. A-priori plans of strategy changes in response to announced changes in the decision tasks were the most important contributor in the paradigms reported in these papers. However, fine-tuning of strategy based on feedback and experience to further approach efficiency is bound to play a role as well. Many potential hurdles for the fine-tuning were investigated and discussed, but none of the interventions were successful in assisting the process enough for it to lead to an improvement in efficiency. This could be interpreted in one of two ways. One could assume that despite our efforts, we have not yet found the right conditions for people to be able to successfully engage in fine-tuning. One could take into consideration the superordinate trade-off between strategy exploration and exploitation to arrive at the conclusion that people are adaptive in ways not previously considered – but also seem to struggle with it in new ways. Both these interpretations open up fascinating and worthwhile avenues for future research; for the extent to which people are adaptive can only be settled with further investigation.

Further pursuing this matter would deepen our understanding of adaptivity in information search and decision making, including its prerequisites and underlying processes. This would ultimately allow us to more fully understand how adaptive behaviour change occurs, but also how it develops in the first place. It should also form an ideal basis both for assisting the development of adaptive behaviour and for assisting specific decision situations under circumstances where we will then know that adaptive behaviour tends to fail – precisely because it is adaptive.

References

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. 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

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

Fiedler, K., Ackerman, R., & Scarampi, C. (2019). Metacognition: Monitoring and controlling one's own knowledge, reasoning and decisions. In R. J. Sternberg & J. Funke (Eds.). *The Psychology of Human Thought: An Introduction* (pp. 89-111). Heidelberg University Publishing.

Fiedler, K., & Wänke, M. (2009). The cognitive-ecological approach to rationality in social psychology. *Social Cognition*, *27*(5), 699–732. https://doi.org/10.1521/soco.2009.27.5.699

Gluth, S., Rieskamp, J., & Buchel, C. (2012). Deciding When to Decide: Time-Variant Sequential Sampling Models Explain the Emergence of Value-Based Decisions in the Human Brain. *Journal of Neuroscience*, *32*(31), 10686–10698. https://doi.org/10.1523/JNEUROSCI.0727-12.2012

Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: the recognition heuristic. *Psychological Review*, *109*(1), 75-90.

Hammond, K. R., Stewart, T. R., Brehmer, B., & Steinmann, D. O. (1986). Social Judgment Theory. In H. R. Arkes & K. R. Hammond (Eds.), *Judgment and decision making: An interdisciplinary reader.* (pp. 56–76). New York, NY: Cambridge University Press.

Harris, C., Fiedler, K., Marien, H., & Custers, R. (2020). Biased preferences through exploitation: How initial biases are consolidated in reward-rich environments. *Journal of Experimental Psychology: General*, *149*(10), 1855–1877. https://doi.org/10.1037/xge0000754

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

Khodadadi, A., Fakhari, P., & Busemeyer, J. R. (2014). Learning to maximize reward rate: A model based on semi-Markov decision processes. *Frontiers in Neuroscience*, *8*. https://doi.org/10.3389/fnins.2014.00101

Lee, M. D., & Cummins, T. D. R. (2004). Evidence Accumulation in Decision Making: Comparing 'Take the Best' and 'Rational' Models. *Psychonomic Bulletin & Review*, *11*(2), 343–352. https://link.springer.com/article/10.3758/BF03196581

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., M. Huys, Q. J., & 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., Prager, J. & Fiedler, K. (2023a). Rivals reloaded – Adapting sample-based speed-accuracy trade-offs through competitive pressure. Under review at *Journal of Experimental Psychology: Learning, Memory, and Cognition.*

McCaughey, L., Prager, J. & Fiedler, K. (2023b). The information cost-benefit trade-off as a sampling problem in information search . In Fiedler, K., Juslin, P., & Denrell, J. (Eds.). *Sampling in Judgment and Decision Making*. Cambridge University Press. *In press.*

McCaughey, L., Prager, J. & Fiedler, K. (2023b). Adapting to information search costs in sample-based decisions. Under review at *Judgment and Decision Making.*

Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., Hausmann, D., Fiedler, K., & Gonzalez, C. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, *2*(3), 191–215. https://doi.org/10.1037/dec0000033

Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. *Psychology of Learning and Motivation*, *26*, 125–173.

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.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review, 63,* 129 –138. http://dx.doi.org/10.1037/ h0042769

Slovic, P., & Lichtenstein, S. (1971). Comparison of Bayesian and regression approaches to the study of information processing in judgment. *Organizational Behavior and Human Performance*, *6*(6), 649–744. https://doi.org/10.1016/0030-5073(71)90033-X

Söllner, A., & Bröder, A. (2015). Toolbox or Adjustable Spanner ? A Critical Comparison of Two Metaphors for Adaptive Decision Making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *42*(2), 215–237. https://doi.org/10.1037/xlm0000162

Spiller, S. A. (2019). Opportunity cost neglect and consideration in the domain of time. *Current Opinion in Psychology, 26,* 98–102. http://dx .doi.org/10.1016/j.copsyc.2018.10.001

Todd, P. M., & Gigerenzer, G. (2007). Environments that make us smart: Ecological rationality. *Current Directions in Psychological Science*, *16*(3), 167–171. https://doi.org/10.1111/j.1467-8721.2007.00497.x

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: heuristics and biases. *Science, 185*, 1124-1130. https://doi.org/10.1126/science.185.4157.1124

Wald, A., & Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, *326*, 339.

APPENDIX 1: First paper

Speed-Accuracy Trade-Offs in Sample-Based Decisions

Klaus Fiedler, Linda McCaughey, Johannes Prager, Jürgen Eichberger, and Knut Schnell
Heidelberg University

Success on many tasks depends on a trade-off between speed and accuracy. In a novel variant, a speed-accuracy trade-off with sample-based decisions in which both speed and accuracy jointly depend on (self-truncated) sample size, we found strong accuracy biases. On every trial of a sequential investment game, participants chose between 2 investment funds based on binary samples of the funds' past outcomes. Participants could stop sampling and decide whenever they felt sufficiently informed. Total payoff was the product of choice accuracy and number of choices completed within the available time (speed). Participants' failure to understand the dominance of speed over accuracy—that speed decreases more than accuracy improves with increasing sample size—led to dramatic oversampling. Our research aimed to examine to what extent metacognitive functions of monitoring and control could correct for the accuracy bias. Experiments 1a through 1c demonstrated similarly strong accuracy biases and payoff losses in psychology and economics students, depressed, and control patients. In Experiments 2 through 4, the accuracy bias persisted despite several manipulations (feedback, sample limit, choice difficulty, payoff, sampling truncation as default) that underlined the speed advantage, reflecting a conspicuous metacognitive deficit. Even when participants faced no risk of losing on incorrect trials but could still win on correct trials (Experiment 3) and when sampling was contingent on the active solicitation of every new element (Experiment 4), participants continued to sample too much and failed to overcome the accuracy bias. The final discussion focuses on psychological reasons and possible remedies for the metacognitive deficit in trade-off regulation.

Keywords: rationality, speed-accuracy trade-off, sample-based decisions, metacognitive myopia

A pervasive question in modern information societies is this: How much information is enough? A rationalist answer says that extended search and information costs must be worthwhile. Even though gathering larger samples of relevant information serves to improve the accuracy of judgments and the likelihood of correct decisions in accordance with the law of large numbers (Bernoulli, 1713) and the wisdom of crowds (Surowiecki, 2004), benefits in accuracy must be pitted against the costs of increasingly large samples and decision time. Adaptive decision making then amounts to solving a speed-accuracy trade-off.

Speed-Accuracy Trade-Offs

Speed-accuracy trade-offs are ubiquitous in real life (Beilock, Bertenthal, Hoerger, & Carr, 2008; Pachella & Pew, 1968; Ghisletta, Joly-Burra, Aichele, Lindenberger, & Schmiedek, 2018). Whether we make consumer decisions, partner choices, select university courses, or decide on holiday destinations, we truncate information search at some point when we feel we know enough to go ahead and make a decision. We try to gather sufficient information to decide correctly, but also try not to lose time and resources gathering unnecessary information. Speed-accuracy trade-offs are of particular importance for research on bounded rationality in societal, political, ecological, and economic problem contexts, involving such fundamental issues as information cost and opportunity cost (Payne, Bettman, & Luce, 1996; Spiller, 2019), satisficing decision strategies (Simon, 1955), and cost-benefit analysis (Swets, Dawes, & Monahan, 2000).

Yet, the greatest part of experimental research on speed-accuracy trade-offs does not relate to such higher order, deliberative decision tasks but uses more elementary task settings involving simple perceptual decisions (Simen, Vlasov, & Papadakis, 2016). Rather than studying complex decision processes in politics, marketing, organizational behavior, risk assessment, or health-related problem solving, most of these experiments with perceptual-motor tasks pursue more fundamental, theoretical, or methodological goals. For instance, perceptual speed-accuracy trade-offs are commonly investigated using drift diffusion models, which have led to important insights into the process of perceptual regulation (Bogacz, Hu, Holmes, & Cohen, 2010; Ratcliff, Smith, & McKoon, 2015). Beyond the investigation of mere perceptual tasks, researchers also examined the apparent gap in optimality between perceptual and conceptual decisions, leading to a comparison of the two types in a speed-accuracy trade-off framework (Jarvstad, Hahn, Warren, & Rushton, 2014; Jarvstad, Rushton, Warren, & Hahn, 2012).

Speed-Accuracy Trade-Offs in Sample-Based Choices

In the present investigation, we propagate a novel trade-off paradigm that uses a sample-based choice task, with self-determined sample size as a conjoint scale for measuring speed and accuracy. While the correctness of choices increases with the size of a random sample of the choice options' previous value, increasing sample size costs time and thereby reduces the number of choices completed within a limited period of time. In the context of this task, solving a speed-accuracy trade-off amounts to drawing samples large enough to warrant a high correctness rate but also making do with samples small enough to maximize the number of completed choices, thereby jointly maximizing the expected payoff per choice as well as the number of choices made in total across the entire series.

This novel set-up resembles the inductive-statistical inference tasks that have been the focus of most research on biases and shortcomings in judgment and decision making (JDM) in the tradition of Tversky and Kahneman (1974); Gigerenzer & Gaissmaier, 2011, or Einhorn and Hogarth (1981). However, the speed-accuracy trade-off explicitly built into the payoff scheme of our sequential decision task sets it apart from other JDM approaches to understanding the dependence of decision quality on sample size that do not implement explicit costs (Erev, Ert, Plonsky, Cohen, & Cohen, 2017; Fiedler & Kareev, 2006, 2011; Kareev, 2000).

Tversky and Kahneman's (1971) famous research on people's naïve reliance on small samples, as if they believe in a "law of small numbers," was neither based on self-determined sample size nor on an explicit trade-off between accuracy and number of judgments. More recent work on decision by experience has led to the conclusion that people tend to draw too small samples (Hertwig & Pleskac, 2010); others came to highlight the adaptive value of small samples (Bell, Raiffa, & Tversky, 1988; Hertwig & Pleskac, 2008; Kareev, 2000) or intuitive decisions (Dijksterhuis & Nordgren, 2006; Gigerenzer & Brighton, 2009). Yet, none of these approaches to studying the relation between decision quality and information quantity involved an explicit trade-off task built into the payoff scheme of a sequential choice task, or any explicit payoff scheme at all, which is crucial for comparisons with a normative standard.

And although the research on expanded judgment and optional stopping tasks (Fried & Peterson, 1969; Irwin & Smith, 1957; Rapoport & Tversky, 1970) can be seen as precursor to our paradigm because they did implement information cost in sample-based decisions, none used a speed-accuracy trade-off to do so. In more general terms, the paradigms used to investigate the theory of melioration (Herrnstein & Prelec, 1991) exhibit interesting similarities with our paradigm. Melioration tasks are constructed to comprise a discrepancy between the strategy at the individual trial level and the strategy at the level of the overall task. However, the elementary task involves neither sampling nor a speed-accuracy trade-off, although the two phenomena may belong to a similar overarching class of paradigms.

The paucity, or almost total absence, of speed-accuracy trade-offs in JDM research is somewhat surprising, given the key role attributed to information costs and efficiency. A distinct novel feature of the trade-off task we are propagating consists in the so far unrecognized normative constraints imposed on the relative weight of speed and accuracy in sample-based choices. The statistical functions relating sample size to speed and accuracy differ systematically, such that the drawbacks of drastically reduced speed resulting from increasing samples size clearly outweighs and overrides the much smaller benefits of increasing accuracy. Consequently, under a wide range of task conditions (choice difficulty; payoff schemes; dynamic presentation format), it is necessary to sacrifice accuracy and increase speed to maximize the total payoff in sample-based choices. The extent to which speed dominates accuracy (under the conditions of the present task) is so counter-intuitive and striking that it is likely to result in a strong accuracy bias and speed neglect.

In the next sections, we first outline the sequential choice task and the normative constraints it imposes on the payoffs. We then describe a computer simulation of the predictable accuracy bias. Finally, the remainder of this article is devoted to a series of experiments on a genuinely metacognitive question: Is it possible to correct, or even to overcome, the accuracy bias and speed neglect over multiple trials, through monitoring and control in the light of feedback and decision aids that ought to sensitize participants for the dominance of speed over accuracy? Our research approach presupposes that metacognitive functions are of primary importance for the regulation of speed and accuracy. Diagnosing and correcting the expected accuracy bias in a secondary metacognitive stage strikes us as more feasible than improving primary cognitive capacities of learning, memory, and reasoning to prevent biases in the first place. Yet, prior work on metacognitive myopia (Fiedler, 2000, 2012) and meta-reasoning (Ackerman & Thompson, 2017) suggests that monitoring and control of one's own memories and inferences may often be ineffective and that metacognitive deficits constitute a major source of irrational behavior.

Dominance of Speed Over Accuracy

Consider the binary choice task used in the experiments below. On each trial of an extended series, participants choose between two options, A and B, based on a sample of positive or negative outcomes for A and B, drawn at random from two universes with experimentally controlled positivity rates of p_A and p_B. Decisions must rely on inferences from finite samples about which of two latent winning rates (p_A or p_B) is superior. To regulate the trade-off between speed and accuracy, participants themselves determine when to stop sampling. By drawing larger samples, they can decrease sampling error and increase the mean choice accuracy. Increasing sample size n serves to increase the reliability of sample estimates and, hence, the correctness rate c and the average payoff a per choice. However, increasing n takes time, reducing the total number k of choices completed in each time period. The efficiency, or total payoff across the entire series, amounts to $Eff = a \cdot k$, that is, the product of expected accuracy a per choice times the number k of choices completed assuming unit payoff of +1 experimental currency unit (ECU) for correct and -1 ECU for incorrect choices.[4]

---

[4] Thus, the expected payoff is the weighted sum $a = c \times (+1 \text{ ECU}) + (1 - c) \times (-1 \text{ ECU})$, that is, the positive and negative unit payoffs multiplied by the respective rates of correct (c) and incorrect ($1 – c$) responses. For instance, correctness rates of $c = .8$ or $c = .9$ transform into average payoffs of $a = .6$ or $a = .8$, respectively.

The same increase in n that lets the a-component in total payoff $Eff = k \cdot a$ grow lets the k component shrink. However, crucially, statistical sampling theory implies a stronger impact of n on speed (k) than accuracy (a): With increasing sample size n, the number k of choices decreases proportionally (i.e., taking twice as much time for doubling n reduces k to $k/2$), but a increases at a much weaker (and negatively accelerated) slope. Therefore, k decreases more rapidly than a increases. To maximize the final payoff, *Eff*, it is thus necessary to make do with a small sample size n, sacrificing the ideal of optimal accuracy to gain a higher number of trials k.

Simulation

Let us illustrate this point in a simulation. On each simulation trial, the algorithm samples n binary observations from each of two choice options, A and B, which differ in their winning probabilities p(A) and p(B), respectively. Manipulation of the differential winning rate $\Delta_p = p(A) - p(B)$ across trials yields easy (Δ_p = .70–.30), intermediate Δ_p = .60–.40), and difficult Δ_p = .55–.45) choice tasks, respectively. The algorithm simply chooses the option with the higher sampled winning rate p^*; it chooses A if the sample contingency is positive, $\Delta_p = p^*(A) - p^*(B) > 0$, and B if it is negative ($\Delta_p < 0$). The vertical axis in Figure 1 shows the proportion of correct choices (across N = 625,000 simulated trials) for fixed sample sizes n.

The positive slope of all three curves reflects increasing correctness rates with increasing n, consistent with Bernoulli's (1713) law of large numbers. Moreover, correctness rates are highest for easy choice tasks (triangles), lowest for difficult tasks (squares), with intermediate tasks (circles) falling in between. Less obvious but central to the trade-off is that across the entire range of n, drawing large samples is hardly worthwhile. At medium difficulty of Δ_p = .20 = .60–.40, doubling n = 6 to n = 12 raises the accuracy rate from .75 to .83. Losing twice as much time and completed choices produces only about 10% higher accuracy. The same is true in the bottom region of sample size; increasing n = 2 to n = 4 increases the correctness rate from .65 to .71. Again, doubled information cost comes along with an accuracy gain of roughly 10%. In either case, the costs of lost time and opportunities (reduced k) resulting from increasing n very quickly become too high compared with a modest gain in accuracy. For choice tasks of any difficulty Δ_p, accuracy is a negatively accelerated function of n (see Figure 1).

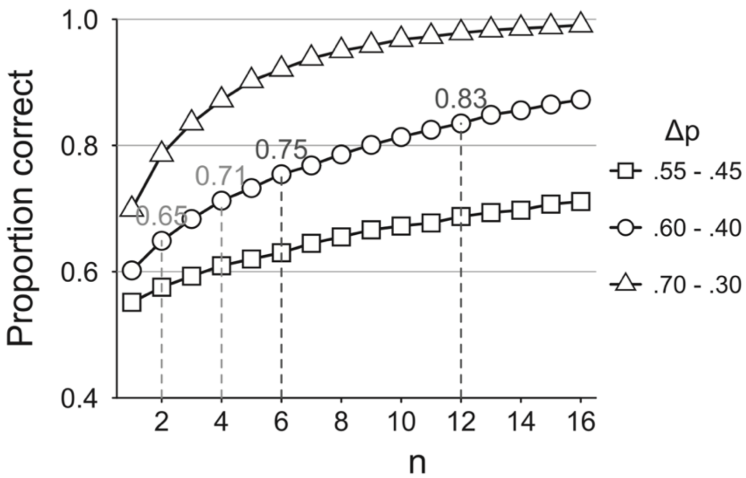

Figure 1. Mean proportion correct as a function of sample size n and choice difficulty Δ_p = .70–.30 (easy), Δ_p = .60–.40 (intermediate), and Δ_p = .55–.45 (difficult). Two pairs of data points indicate accuracy gains from doubling n (from 2 to 4, light gray, or from 6 to 12, dark gray).

Benchmark Strategy

In sample-based choices, the normative-statistical constraints come to interact with several complicating factors that render the mathematical determination of a global optimum impossible. In reality, optimal sample size is a complex function of decision makers' prior assumptions about the superiority of options, their sense of the diagnosticity of a growing sample, the distribution of choice difficulties, the samples' reliability, and assumptions about dynamically changing environments. Despite this insurmountable uncertainty, we can use as surrogate for a mathematical optimum a simple algorithm based on Edwards' (1965) posterior odds ratio test. Assuming that there is no a priori preference for either option, the algorithm simplifies the binary choice to an easily applicable "tie-breaker" strategy. The tie breaker stops sampling and makes a choice as soon as sampled information favors one option over the other to a minimal degree. It then chooses the option with the higher sample winning rate. From the simulation results in Figure 1, we see that accuracy shows the strongest increase from 0 to the lowest level of n. The tie breaker can optimally exploit this initial high gain in accuracy with a minimal sample difference of 1 in most task contexts. Thus, it corroborates that under most reasonable conditions sample-based choices call for speedy strategies.

Suffice it to mention briefly that the speed advantage is qualified, theoretically, by two boundary conditions, namely, the additive time constant required to execute the choice and the presentation speed of the sample elements. The primacy of speed decreases at extremely long decision times (reducing the direct impact of sample size on number of choices) and extremely fast presentation speed (reducing the costs of large samples). Yet, by avoiding extreme presentation modes (e.g., samples presented above the flicker-fusion frequency), and avoiding extremely unequal payoff scheme (much higher losses for incorrect choices than gains for correct choices), most naturally appearing sample-based choice tasks will produce the pronounced accuracy bias required to study metacognitive monitoring and correction processes.

Metacognitive Regulation of Speed and Accuracy in Sample-Based Choices

Though the cognitive choice process need not necessarily follow the simulated benchmark rule, there is ample evidence that decisions are sensitive to sample contents (Juslin, Winman, & Hansson, 2007; Kareev, Arnon, & Horwitz-Zeliger, 2002; Stewart, Chater, & Brown, 2006). Moreover, the dominance of speed over accuracy is so striking and counterintuitive that it seems safe to predict a distinct accuracy bias as a precondition for our metacognitive research project. The expected initial accuracy bias is also in line with recent reviews of suboptimal trade-off strategies in perceptual decisions, which also tend to be too cautious (Bogacz et al., 2010; Evans & Brown, 2017; Starns & Ratcliff, 2010, 2012). Especially when the trade-off is explicitly built into the payoff structure, so that total payoff is contingent on both accuracy and number of all decisions completed in a fixed period, speed tends to be neglected and accuracy receives too much weight even in sensorimotor tasks that were formerly believed to produce optimal performance (Evans & Brown, 2017).

Notwithstanding this convergence, speed regulation in sample-based choice is inherently different from trade-offs on perceptual tasks. In perceptual decisions, speed and accuracy are not bound to a common scale, sample size. To decide whether the majority of symbols in a scatter diagram move in one or the other direction (Evans & Brown, 2017), whether one recognizes a written word, or to answer an arithmetic calculation, participants do not have to overcome a strong initial bias. Rather than correcting for a wrong strategy in a secondary metacognitive loop, they regulate speed during the primary cognitive process. They stop processing the stimulus and make a decision when perception or calculation is extensive enough to perceive a solution. The likelihood of a correct response reaches a clear peak during a typical time window. Regulation of speed then amounts to figuring out the lower time boundary that renders the first solution visible. There is no comparable peak in the increase of accuracy over time in sample-based decisions, which never make the correct solution visible. On the contrary, the experience of ambiguity and conflict often increases with increasing *n*, and smallest samples often provide strongest evidence in favor of one choice option (Fiedler, Renn, & Kareev, 2010; Prager, Krueger, & Fiedler, 2018).

Expectations and Preview of Experiments

We do not expect our participants to solve the trade-off analytically. We rather expect that most participants will choose their initial strategy intuitively. Because the strong advantage of speed over accuracy is counterintuitive, this intuition will be clearly biased toward accuracy and against speed. Due to the ubiquitous nature of speed-accuracy trade-offs in real life decisions and their scarcity in previous research on JDM this initial demonstration of a clear-cut accuracy bias and speed neglect is of theoretical and practical interest.

However, although this initial bias is predictable on normative grounds, an open empirical question that is the focus of the present research is whether participants can detect and correct for the initial bias through metacognitive monitoring and control. During the sequential choice task, they have a chance to experience the difficulty of multiple pairs of samples of variable size; they have ample opportunities to try out different strategies affecting sample size, accuracy, and resulting payoff. They can actively explore, by varying the amount of information prior to their decision, whether more or less information will increase total payoff per time-unit or total payoff per amount of information and adapt their strategy accordingly.

Thus, having established the basic accuracy bias and the viability and usability of the entire task set-up in three parallel experiments (1a through 1c), the major part of the present research is devoted to studying the effectiveness of metacognitive control. Would monitoring and control processes help participants to ameliorate or even overcome the dysfunctional accuracy bias? For a sound test of this crucial issue, it is of course essential that participants are provided with appropriate feedback and sufficient metacognitive cues necessary for successful monitoring and control. Let us be explicit about the available cues that participants can use to detect the inefficiency of accuracy-driven strategies and the higher payoff gained through faster strategies.

Analogous to a typical calibration experiment (Koriat, 2012; Winman & Juslin, 1993), participants can translate their metacognitive feeling of choice difficulty (induced by varying of Δ_p across trials) into an inference of subjective confidence, which is however not expressed as a confidence rating but as careful sampling and decision speed. Fast choices reflect higher confidence than cautious choices after more extended sampling. Even without further feedback or decision aids, the sequential choice task strongly resembles a typical metacognition paradigm.

Throughout all experiments, a timer reminds participants of the remaining time available to complete more choices. When immediate trial-by-trial correctness feedback is introduced, participants can assess the relation between sample size (n) and accuracy (a) across trials, and they can extract the function relating sample size to time required per choice. In Experiments 1 and 2, manipulations of the upper sample limit enable participants to experience samples of different size and to recognize the clearly stronger impact of sample size on speed than on accuracy. Another opportunity to experience the advantage of faster strategies is induced in Experiment 4 using a different format of the sampling task. In Experiment 3, zero payoff for wrong choices (i.e., a gain-only payoff scheme) affords a compelling incentive for exploring speedy strategies at no risk, whereas aggravated punishment (larger negative payoff) for wrong decisions highlight the insight that faster choices based on smaller samples afford a viable strategy even under such loss-aversive conditions. Thus, throughout a large number of choice trials and across a variety of different task conditions, participants have manifold opportunities to monitor and control their speed and accuracy.

Nevertheless, we did not expect that even optimal feedback and diverse helpful decision aids would eliminate the prevailing accuracy bias. This pessimistic expectation is based on a growing body of evidence on deficits in metareasoning (Dunning, 2011; Larrick & Soll, 2008; Peer & Gamliel, 2013; Svenson & Eriksson, 2017) and on metacognitive myopia (Alves & Mata, 2019; Fiedler, 2000, 2008, 2012), especially for the role of sample size on judgments and decisions (Fiedler, Hofferbert, & Wöllert, 2018; Fiedler, Kareev, et al., 2016; Price, Kimura, Smith, & Marshall, 2014).

The notion of evaluability (Hsee, 1996) offers an explanation why metacognitive monitoring and control may fail. In the context of the trade-off choice task, speed is clearly less evaluable than accuracy. Whether a choice is correct or not is an absolute criterion that can always be evaluated immediately. Speed, however, can only be evaluated in relative terms, at the end of a sequence of trials; there is no absolute criterion to classify singular choices as either fast or slow. Participants must therefore sacrifice what they experience to be an immediately rewarding strategy, namely careful sampling rewarded by high accuracy, to detect the advantage of another, speed-oriented strategy.

In the General Discussion, we provide a summary of the empirical evidence on metacognitive quality control in sample-based choices. In addition to strong and persistent evidence on the robustness of the conservative accuracy bias in sample-based decisions, we discuss several emerging insights about the underlying monitoring and control mistakes that prevent decision makers from exploiting the greatest part of possible payoff. With a final outlook on how the massive bias in speed-accuracy regulation could be overcome in future research, we close the article.

Experiments 1a, 1b, and 1c

To establish the accuracy bias and its generality, we first conducted the same basic experiment in three different labs with different participant groups. Using symmetric unit payoff (±1 ECU for correct and incorrect choices, respectively) and choices varying over three difficulty levels (Δ_p = .1 vs. .2 vs. .4) and two upper sample-size limits ($n_{\max}$ = 18 vs. 12), we applied the basic set-up to students of psychology (Experiment 1a) and economics (Experiment 1b) and to depressed and control patients (Experiment 1c). Although economists may be less conservative than psychologists, who should be less conservative than depressed patients (Fiedler et al., 2010; Quiroga, Hernández, Rubio, Shih, Santacreu, & Howard, 2007), they should all exhibit the same basic bias.

Method

Participants and design. Forty female and male psychology students participated in Experiment 1a, recruited via the *Studienportal* platform (software hroot: Bock, Baetge, & Nicklisch, 2014) at Heidelberg University. Experiment 1b included 52 students in economics (all from Heidelberg University). Experiment 1c included 29 depressed patients and 31 nondepressed control patients who were under treatment in the psychiatry department of Heidelberg University. In all three experiments, two repeated-measures factors were manipulated within participants: choice difficulty (Δ_p = .10 vs. .20 vs. .40) and sample-size upper limit ($n_{\max}$ = 12 vs. 18).

Materials and procedures. Experiments 1a through 1c were conducted simultaneously, using the same materials and the same software that controlled the experimental procedure. Instructions announced a stock-market game. Participants were told that for each round of the game a sample of recent changes in stock quotations of two stocks, A and B, was shown. Their task would be to select the overall superior stock. On each trial of a sequential task, they would be shown, for a randomly drawn sample out of the last 100 trading days, on how many days the value of two stocks, A and B, had increased (1, presented in green) or decreased (2, presented in magenta). Based on these two binomial samples of 1 and 2 symbols, they had to select the one stock that was better across all 100 trading days. They would gain 1 ECU for each correct choice and lose 1 ECU for each incorrect choice. It was explicated that the better option can be determined objectively from the universe of 100 stock value changes. Thus, making correct decisions, they were told, called for an inference from a restricted sample to a total record across all 100 days. It was explicitly stated that the overall time for playing the game was limited. Instructions emphasized the trade-off between speed (smaller samples) and accuracy (larger samples). Neither strategy was presented as more favorable.

Participants could initiate a new round with a new sample of stock quotations by pressing the space bar. For as long as they kept pressing the space bar, sequential presentation of quotations (1 and 2) continued. Every 500 ms a new value appeared, alternating between stock A and B, presented on the left and right side of the screen, respectively. Arrows remained on screen for as long as sampling continued. Depending on the upper sample limit manipulation between trials (n_{max} = 18 vs. 12), participants were informed that they could consider the stock value changes on maximally 18 (vs. 12) days. This limit was presented in the center of the screen, surrounded by a pie-chart indicating how much information was already shown and how much was left until the upper limit was reached.

The "true" success rates p(A) and p(B) of options A and B (i.e., the probabilities of stock value increase), and hence the differential success rates $\Delta_p = |p(A) - p(B)|$ were manipulated, such that subsets of choices involved easy (Δ_p = .4), intermediate (Δ_p = .2), or difficult choices (Δ_p = .1). Correct choices of A when $p(A) > p(B)$ or of B when $p(A) < p(B)$ were rewarded with a positive unit payoff (+1 ECU); incorrect choices yielded a negative unit payoff (-1 ECU). The allocation of trials to difficulty levels (Δ_p) and upper sample-size limits (n_{max}) was random, under the constraint that all six combinations of three Δ_p -levels and two n_{max}-levels had to be included in each subset of six successive trials. The winning probabilities $p(\uparrow|A)$ and $p(\uparrow|B)$ of both options varied in the range of .9 to .1 (with a step size of .1), within which a certain Δ_p was generated, were determined at random, just as the allocation of A or B to the superior condition. The temporal and spatial presentation of the n (pairs of) symbols on the left and right side of the screen were also randomized.

As soon as participants terminated sampling by releasing the space bar or, when they had reached the maximal sample size, a decision screen appeared on which they were asked to press key *F* or *J* to indicate that they favored Stock A over Stock B or vice versa. Immediately after the response, participants could start the next sampling phase by pressing the space bar again.

Throughout the session, the proportion of total time that had elapsed appeared graphically as the filled proportion of a bar outlined at the top right of the screen. Thus, although participants were not told the total time—set to 20 min—as a numerical quantity, they were permanently reminded of the limited time and could easily infer that they could save time by keeping samples small. When the time had elapsed, they were thanked and answered any questions about the purpose of the experiment. As payment they received their point balance converted at a rate of 0.05€ per ECU.[5]

Results

For this and all following experiments, the results section starts with a decomposition of the overall payoff into the speed and accuracy component, pooling across all participants and experimental conditions. We then examine the impact of experimental conditions. Of special interest is the slope of the functions relating speed and accuracy to sample size.

A distinct methodological asset of the sample-based choice task is the availability of an actuarial measure of the sample contents, independent of the choice data. By analyzing the intermediary sample data, we can check whether participants' sampling behavior created sufficient variation in n and in Δp^* by inspecting the sample-based correlations between n and in Δp^* and the impact of the first few sampled observations on the resulting n.

Overall performance. The average number k of trials completed by the psychology students who participated in Experiment 1a was M_k = 195.15 (SD_k = 85.90). The average accuracy score a (i.e., average + 1 vs. – 1 outcome) was M_a = .87 (SD_a = .09). This corresponds to an average correctness proportion c of M_c = .94 (SD_c = .05), suggesting that most choices were correct, and most sample pairs should have been experienced as easy to discriminate. The average sample size n (for each sample in a pair) was M_n = 7.43 (SD_n = 3.13). The resulting mean efficiency or total point balance $Eff = a \cdot k$ amounts to M_{Eff} = 166.43 (SD_{Eff} = 60.20).

Speed-accuracy trade-off. Turning to the speed-accuracy trade-off, the issue of major interest, it is instructive to compare the impact of increasing sample size n on accuracy a, on the one hand, and on the number k of completed trials, on the other hand. We thus plotted individual participants' mean a and k scores and their resulting efficiency products $Eff = a \cdot k$ as a function of individuals' mean sample size n. These plots appear in Figure 2 for comparison.

Starting from the left chart, we can see that a indeed increased with increasing n. A positive correlation $r_{n,a}$ = +.40, p = .011, between n and a indicates that participants who drew larger samples were, on average, more accurate. However, a glance at the middle chart shows that k, the number of choices completed, decreased much faster with increasing n than a increased. The strong negative correlation $r_{n,k}$ = -.88, p < .001, between n and k clearly overrode the moderate positive correlation $r_{n,a}$ = +.40. Whereas a (left chart) increased no more than 15% to 20% across the range of individual n scores (on the horizontal axis), the decrease in k was several hundred percent. As a consequence, the total payoff in the right chart (i.e., the product $Eff = a \cdot k$) was almost perfectly correlated with k, $r_{k,Eff}$ = +.98, p < .001, but inversely related to a, $r_{a,Eff}$ = -.91, p < .001. The data points for Eff (right chart) mirror those for k (middle chart). Thus, consistent with our normative considerations and with the simulation in Figure 1, solving the speed-accuracy trade-off in sample-based choices calls for a distinct emphasis on speed. Increasing n to improve accuracy is highly counterproductive.

The highest Eff was attained by those participants who reached the highest k scores by drawing the smallest samples. The one participant (on the left end of all three plots) who emphasized speed more than all others, with an average n less than 2, gained the absolutely highest Eff score, raising the performance-dependent payoff to 382 ECU = 19.10€. This minimalist strategy was superior to all strategies that gave more weight to accuracy.

[5] At the time of the experiment the exchange rate EUR to USD was about 1.1.

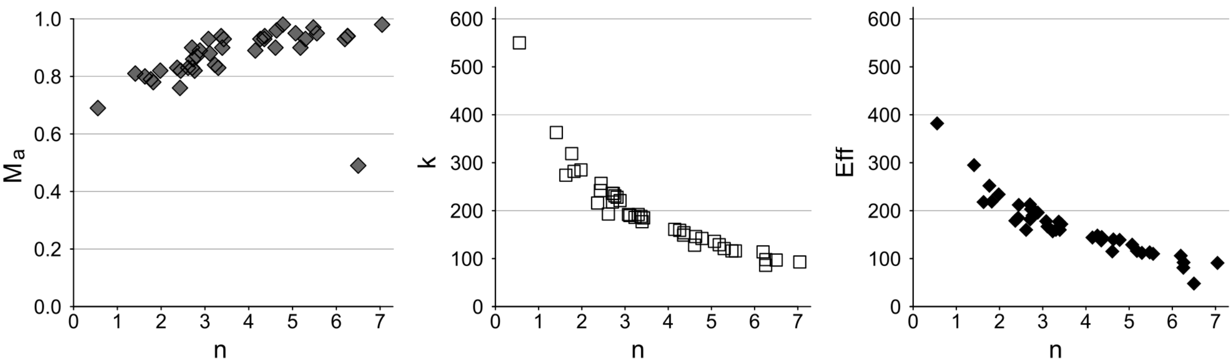

Figure 2: From left to right: Individual mean accuracy $a = c - (1 - c)$ (with c indicating the correctness rate), number of trials k completed within 20', and efficiency $Eff = k \cdot a$ plotted as a function of individual average sample size n.

The entire pattern generalized across all three participant groups (Experiment 1a, 1b, and 1c) and from the first to the second half of trials (see Table 1), contrary to an insufficient-learning account. The overall accuracy was high; average M_a scores of .87 and .88 (for first and second half of trials) correspond to correctness proportions of .93 and .94, respectively. Average sample sizes of M_n = 7.35 and 7.63 (with *SD*s ~3) started too conservatively and did not change this cautious strategy over an extended series, missing the possible payoff by many hundred percent.

Thus, our analyses corroborate the premise that sample-based choices induce a marked accuracy bias and speed neglect. Decision makers dramatically failed at exploiting the speed-accuracy trade-off, drawing samples that were too large by magnitudes. Neither fatigue or motivation loss nor learning (Khodadadi, Fakhari, & Busemeyer, 2014) moderated this pattern, which remained stable from first to second half of trials. Table 1 also highlights the robustness of this basic finding across three parallel experiments (1a through 1c) with different participant groups. Although economists (M_n = 6.61) drew somewhat smaller samples than psychologists (M_n = 7.43) and depressed patients (M_n = 7.89) were most cautious, all descriptive statistics were remarkably similar across participant groups of Experiment 1a to 1c.

Impact of experimental manipulations within participants. How did the within-participants manipulations of choice difficulty (Δ_p = .10 vs. .20 vs. .40) and upper sample limit ($n_{\max}$ = 18 vs. 12) affect the dependent measures? For an empirical answer based on Experiment 1a, with psychology students as participants, we conducted two analyses of variance (ANOVAs) for M_n and M_a. Suffice it to mention that Experiments 1b (economists) and 1c (psychiatric patients) yielded similar results (see Table 1); they will be documented in a subsequent publication.

Table 1
Decision Performance for Experiments 1a, 1b, and 1c Across Trials (Total) and Separately for the First and Second Half of Trials

Measure	Experiment 1a Psychology students			Experiment 1b Economics students			Experiment 1c Depressed patients			Experiment 1c Control patients		
	Total	1st Half	2nd Half	**Total**	1st Half	2nd Half	**Total**	1st Half	2nd Half	**Total**	1st Half	2nd Half
M_n	**7.43**	7.35	7.63	**6.61**	6.55	6.79	**7.89**	7.92	7.94	**7.16**	7.13	7.31
(SD_n)	**(3.13)**	(3.02)	(3.17)	**(3.07)**	(2.74)	(3.45)	**(3.00)**	(2.82)	(3.25)	**(3.16)**	(3.02)	(3.35)
M_a	**.87**	.87	.88	**.86**	.86	.87	**.87**	.86	.88	**.88**	.87	.88
(SD_a)	**(.09)**	(.09)	(.09)	**(.08)**	(.08)	(.09)	**(.09)**	(.10)	(.09)	**(.07)**	(.08)	(.07)
M_k	**195.2**	97.5	97.7	**225.0**	108.3	116.7	**175.7**	86.2	89.5	**200.0**	98.5	101.4
(SD_k)	**(85.9)**	(39.4)	(51.5)	**(108.8)**	(44.1)	(68.4)	**(58.7)**	(25.7)	(34.1)	**(90.4)**	(40.0)	(55.8)
M_{Eff}	**166.4**	82.9	83.5	**187.4**	91.1	96.2	**151.1**	73.7	77.3	**170.0**	83.9	86.1
(SD_{Eff})	**(60.2)**	(29.3)	(35.0)	**(71.5)**	(30.3)	(43.8)	**(45.8)**	(20.5)	(26.5)	**(60.6)**	(28.8)	(36.5)
$r_{n,a}$	**.40**	.34	.44	**.72**	.64	.71	**.24**	.17	.30	**.82**	.69	.86
$r_{n,k}$	**–.88**	–.89	–.82	**–.90**	–.89	–.87	**–.95**	–.93	–.95	**–.89**	–.89	–.84
$r_{a,Eff}$	**–.34**	–.19	–.42	**–.73**	–.54	–.75	**–.08**	.06	–.18	**–.76**	–.47	–.84
$r_{k,Eff}$	**.98**	.97	.98	**.98**	.98	.98	**.97**	.95	.98	**.99**	.98	.99

Note. M_n = mean sample size *n*; SD_n = standard deviation of sample size *n*; M_a = mean accuracy *a*; SD_a = standard deviation of accuracy *a*; M_k = mean number of completed choices *k*; SD_k = standard deviation of number of completed choices *k*; M_{Eff} = mean efficiency *Eff*; SD_{Eff} = standard deviation of efficiency *Eff*.

The ANOVA of average sample size M_n (for the 3 × 2 different subsets of choices varying in difficulty and sample-size limit) yielded strong main effects for choice difficulty, $F(2, 78) = 126.055$, $p < .001$, $\eta^2 = .764$, and for upper limit, $F(1, 39) = 24.344$, $p < .001$, $\eta^2 = .384$.[6] Sample size M_n increased markedly with increasing difficulty from 6.410 (3.127) for $\Delta_p = .40$ to 7.555 (3.355) for $\Delta_p = .20$ to 8.290 (3.508) for $\Delta_p = .10$ and with increasing $n_{\max}$ from 6.823 (2.645) for $n_{\max} = 12$ to 8.014 (3.953) for $n_{\max} = 18$, respectively. Treating Δ_p as a quantitative variable resulted in $F(1, 39) = 159.224$, $p < .001$. The interaction was weak but also significant, $F(2, 78) = 3.762$, $p = .028$, η^2=. 088, reflecting a reduced influence of the upper sample limit at the easiest level of $\Delta_p = .4$, at which M_n was generally small so that the upper limit $n_{\max}$ was rarely reached.

In the ANOVA of accuracy scores M_a, the main effect of task difficulty was also highly significant, $F(2, 78) = 119.326$, $p < .001$, $\eta^2 = .754$, reflecting increasing accuracy with increasing Δ_p, from .763 (.122) to .893 (.111) to .962 (.087). Treating Δ_p as a quantitative factor produced, $F(1, 39) = 145.357$, $p < .001$. However, no main effect obtained for upper sample limit, $F(1, 39) = 0.027$, $p < .870$, $\eta^2 = .001$, suggesting that the added extra sample size in the $n_{\max} = 18$ condition was in vain. Although M_n increased by nearly 20% from $n_{\max}$ =12 to $n_{\max}$ =18, the corresponding accuracy levels remained the same, $M_a = .873$ (.136) and .872 (.136), respectively. The interaction was negligible, $F(2,78) = 0.159$, $p = .692$, $\eta^2 = .004$.[7] These

[6] The corresponding *F*s in Experiment 1b were as follows: $F(2,102) = 124.094$, $p < .001$, η^2 =.709; $F(1,51) = 19.559$, $p < .001$, $\eta^2 = .277$; and $F(2,102) = 7.182$, $p < .001$, $\eta^2 = .123$. The *F*s in Experiment 1c were as follows: For depressed patients $F(2,56) = 106.595$, $p < .001$, $\eta^2 = .792$, $F(1,28) = 17.726$, $p < .001$, $\eta^2 = .388$ and $F(2,56) = 9.225$, $p < .001$, $\eta^2 = .248$, and for controls, $F(2,60) = 95.986$, $p < .001$, η^2 =.76; $F(1,30) = 18.501$, $p < .001$, $\eta^2 = .381$; and $F(2,60) = 3.552$, $p < .035$, $\eta^2 = .106$.

[7] Similar results were obtained in Experiment 1b, $F(2,102) = 196.555$, $p < .001$, $\eta^2 = .794$, $F(1,51) = 8.353$, $p = .006$, $\eta^2 = .141$, and $F(2,102) = 1.972$, $p = .166$, $\eta^2 = .037$ (for 1b), for depressives in Experiment 1c, $F(2,56) = 84.806$, $p < .001$, $\eta^2 = .752$, $F(1,28) = 0.981$, $p = .330$, $\eta^2 = .034$, and $F(2,56) = 2.446$, $p = .129$, $\eta^2 = .080$; and for controls, $F(2,60) = 109.870$, $p < .001$, $\eta^2 = .786$, $F(1,30) = 0.518$, $p = .477$, $\eta^2 = .017$, and $F(2,60) = 0.265$, $p = .610$, $\eta^2 = .009$.

consistent results testify to the effectiveness of both manipulations and to the sincerity of the experimental participants.

Note that it is impossible by design to conduct analogous repeated-measures ANOVAs for *k* and *Eff*, because all 3 x 2 choice types had to be included in each subset of six successive trials, thus equating the overall number *k* of choices across the 3 x 2 design. Still, the impact of both within-participant factors on the trade-off is evident (in Figure 3) from a comparative plot of average trial latency (M_l) and accuracy (M_a) as a function of mean sample size (M_n). The horizontal axis refers to overall M_n per participant, conceived as a personality attribute rather than separate sample sizes calculated for all 3 x 2 trial types.

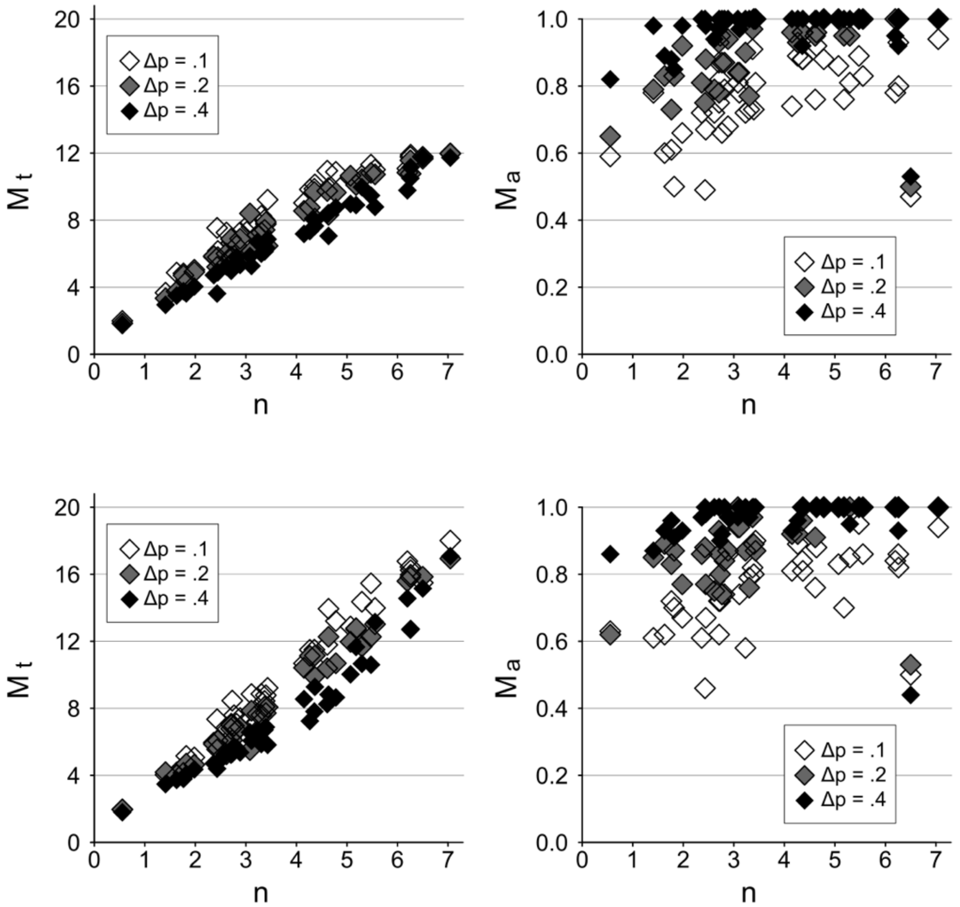

Figure 3. Average time M_l per trial (left charts) and accuracy M_a (right charts) as function of difficulty Δ_p and upper limit (n_{max}=12, upper charts, vs. n_{max}=18, lower charts).

Figure 3 highlights that accuracy was generally high, reaching a ceiling (approaching M_a = 1.00) for the easiest choices (Δ_p = .40; black diamonds). For all difficulty levels, the increase in M_a was very modest. As M_n was more than tripled, M_a improved by little more than 20%, in sharp contrast to the roughly linear increase in mean time per trial M_t with growing M_n. As M_n increased (e.g., from 3 to 13), the required time increased by the same magnitude (e.g. from ~ 4 s to 16 s) whereas the gain in accuracy M_a remained disproportionally small. The left panel of Figure 3 shows that participants at increasing levels of M_n took increasing choice latencies at increasing difficulty levels Δ_p (shadings of diamonds). This time investment was not worthwhile because accuracy (especially at Δ_p = .40) reached a ceiling anyway (right panel). Increasing the upper limit from $n_{\max}$= 12 to $n_{\max}$=18 caused an additional waste of time (see slope of M_t plot); this manipulation had virtually no effect on accuracy M_a (see right panel).

Covariation of accuracy and sample size across choices, within participants. Why did participants waste so much time in gathering too large samples and thereby greatly underperform on the speed-accuracy trade-off? An analysis of the covariation in sample size (n) and accuracy (a) across all choices within participants affords a telling answer to this crucial question. For each individual participant, we computed correlations across all j = 1, 2... k choices between the following indices: accuracy a (+1 vs. –1), sample size n , signed difference Δ_j* of proportions of positive outcomes in the sample (giving positive signs to differences pointing in the same direction as the true Δ_p), signed primacy difference ${}_5\Delta_j$* in the first maximally 5 items per sample, the unsigned difference $|\Delta_j*|$, and the unsigned primacy difference $|{}_5\Delta_j*|$.[8] As a precondition for these correlational analyses, it should be mentioned that n_j varied sufficiently across trials. The standard deviation of n_j within the average participants amounts to 3.12.

Correlations between all five indices, computed across trials within individuals, revealed a consistent pattern. Correlations $r(n_j,|\Delta_j*|)$ between sample size n_j and strength of sampled (unsigned) differences $|\Delta_j*|$ were clearly negative, M_r = –.526, $t(39)$ = 17.810, p < .001, reflecting a sensible corollary of the law of large numbers: smaller sample pairs yielded larger absolute differences. However, because the sign of most observed differences was correct, that is, because the average individual correlation between Δ_j* and $|\Delta_j*|$ was highly positive, $r_{\Delta j*,|\Delta j*|}$ = –.375, $t(39)$ = 13.66, p < .001, small samples tended to exaggerate the actuarial sample difference coded in the correct direction. The average individual correlation $r(n_j,\Delta_j*)$ amounts to .788, $t(39)$ = 28.28, p < .001. This counterintuitive negative relation between sample size and sensitivity, which is characteristic of self-truncated sampling (Fiedler, et al., 2010; Prager et al., 2018), highlights the tragedy that underlies the accuracy bias. Increasing n is a waste of time because sampled evidence and accuracy hardly increase with time or sample size.

At first glance, the negative correlation seems to reflect an adaptive strategy, namely continue sampling and increasing n_j when the sampled evidence $|\Delta_j*|$ is weak. This seemingly plausible strategy indeed offers an explanation of the inefficient trade-off regulation. A clearly negative correlation between n_j and $|{}_5\Delta_j*|$, which amounts on average to M_r = –.444, $t(39)$ = 13.000, p < .001, says that participants continued sampling, raising n_j to higher levels, when the sampled evidence $|{}_5\Delta_j*|$ in the first five observations was weak. They stopped sampling earlier when evidence was strong. Tragically, though, the seemingly adaptive strategy was in vain. Increasing initially weak samples hardly served to enhance accuracy; the average correlation $r(n_j,a_j)$ within participant j was negligible, M_r = –.024, $t(39)$ = 1.676, p = .102.

Thus, what looks like adaptive regulation turned out to be counterproductive, causing a waste of time with growing amount of information. What participants must learn is to shift from a plausible to a counterintuitive strategy, namely, choosing the leading option at small n, regardless of the evidence gathered so far. To master the trade-off, it is necessary to disregard the individual-trial level and instead consistently rely on indicators of efficiency at the overall task level to be able to efficiently monitor and improve one's strategy. No matter how weak the primacy input is, it is unlikely that $|\Delta_j*|$ will increase with growing n. The average

---

[8] Because participants did not know the true direction of an observed difference beforehand, their tendency to increase n when no clear-cut difference is apparent must rely on unsigned cues, $|\Delta_j*|$ and $|{}_5\Delta_j*|$.

correlation $r(|_5\Delta_j^*|,|\Delta_j^*|)$ between early and final sample differences was as high as $M_r = .832$, $t(39) = 32.567$, $p < .001$.

Discussion

In summary, then, Experiments 1a, 1b, and 1c testify to a robust and predictable accuracy bias in the present version of speed-accuracy trade-off, when sample size jointly constrains the regulation of speed and accuracy in sample-based choices. Indeed, the tendency to overweight accuracy and to neglect speed was immense. The conservative bias toward accuracy and the corresponding waste of possible payoff were dramatic. The positive impact on accuracy of interindividual and experimentally induced variation in sample size was overridden by the stronger negative impact of sample size on speed (number of completed choices). As a result, the overall payoff was proportional to speed but negatively related to accuracy, with the obvious implication that sample size should be reduced considerably.

However, most participants ignored this normative rule, foregoing the possibility to increase their total payoff from roughly 4€ to 20€ or more. This memorable failure to exploit the trade-off raises two questions. First, how can we explain the dysfunctional accuracy bias? Could it reflect an unfair or unsolvable task? We believe that participants had a fair chance to exhibit their metacognitive regulation competence. Analogous to countless calibration experiments, in which participants must translate a feeling of difficulty into a confidence rating, the present participants could transform their feeling of choice difficulty into an appropriate sample size. Their inability to optimize this mapping of Δ_p onto n constitutes a fair measure of calibration, which is similar to that of traditional metacognition research. Moreover, there was sufficient variation is samples size, part of it induced by the upper limit manipulation, to enable participants to extract correlations between samples size, accuracy, and time required.

However, analyses of the covariation across all choices within participants suggested an interesting explanation for the deficit. Rather than reflecting sloppy responding or a lack of motivation, the underperformance appeared to originate in a seemingly plausible regulation strategy. Making sampling truncation contingent on evidence strength in the sample seemed reasonable but was indeed detrimental. Individual *Eff* suffered greatly from the accuracy-motivated tendency to increase n_j when early sample evidence was weak, as evident in the negative average correlation $r(Eff, r[n,|_5\Delta_j^*|]) = -.33$, $p = .038$. The tendency to increase n when the primacy sample difference was small correlated negatively with total payoff *Eff*.

The second question is whether the well-motivated but fruitless accuracy bias can be corrected through appropriate monitoring and control aids. Is it possible to overcome the accuracy bias when the task is presented in a richer format or when specific decision aids are provided? We devised three further experiments to find an empirical answer.

Task comprehension. First, it would be important to rule out the possibility of insufficient understanding of task instructions. Did participants understand that the ultimate performance criterion depends on a competition of speed and accuracy, or did they simply miss this crucial point, due to suboptimal instructions or insufficient motivation? Several findings in Experiments 1a-c make this interpretation unlikely, such as the sensible influence of task difficulty and sample limit on speed and accuracy. Still, we attempted to check on task comprehension.

Explicit feedback. The second task condition that might have hindered optimal trade-off performance is certainly the absence of explicit feedback. To be sure, the lack of feedback did not render the metacognitive task unfair or impossible to solve. Just as confidence ratings in a typical calibration study must rely on metacognitive feelings of item difficulty, participants in

Experiments 1a through 1c could base their self-determined sample size on feelings of choice difficulty. After all, sample size affords a sensible measure of confidence. Moreover, participants' own choices may have provided implicit, self-constructed feedback (Elwin, Juslin, Olsson, & Enkvist, 2007) about the superiority of one option. It is nevertheless possible that explicit and immediate correctness feedback might be crucial for effective trade-off performance. Accordingly, we provided explicit feedback about the correct choice in all further experiments.

Intrinsic difficulty to evaluate speed. A third inhibitory condition could be the unequal evaluability (Hsee & Zhang, 2010) of accuracy and speed. Whether a choice is correct (accurate) or not is an unambiguous criterion that can be evaluated in absolute terms. In contrast, speed is indeterminate at the level of individual choices but can only be evaluated across a series of multiple choices on a relative scale, in comparison to the speed of other serial tasks. No absolute criterion exists to classify speed as too low, optimal, or too fast.

Because of this intrinsic difficulty to evaluate speed, participants may be overly suggestible to the provision of upper sample-size limits, which may be misunderstood as anchors or "nudges" suggesting presumably reasonable sample-size magnitudes. Indeed, two different possibilities must be considered. The upper limits provided in the preceding experiments (18 vs. 12) may have been higher than participants' spontaneously chosen samples sizes. Or, any upper limit may unduly restrict the natural truncation decisions, which may often result in even larger sample sizes.

To examine both possibilities, in Experiment 2 we eliminated the sample limit in one experimental group and reduced the upper limit from 18 versus 12 to 12 versus 8 in the other group. Performance differences between conditions should clarify the impact of unwanted demand effects associated with upper limits. No sample limit was used in Experiment 3 and 4, which introduced further helpful cues to overcome the accuracy bias.

Experiment 2

Thus, we introduced explicit feedback and eliminated the impact of potentially misleading sample-size anchors in a no-upper-limit condition, while reducing the (variable) limit to n_{max} = 12 versus 8 in another group. A post-experimental questionnaire assessed participants' understanding of the trade-off and their metacognitive assessment of own performance.

Method

Participants and design. A total of 78 participants were recruited on Studienportal. They were aged 23.31 on average (SD = 7.34); 61 were female. The experiment was the third out of four encompassed in a 1-hr session (involving contingency learning, evaluative conditioning, motivated self-ratings). Participants received a basic payment of 5€ plus an additional amount contingent on Eff. In one condition sample sizes were limited to $n_{max} = 8$ versus 12 on different trials. No upper limit was set in the other experimental condition. Choice difficulty ($\Delta_p = .10$ vs. .20 vs. .40) was manipulated within all participants of both conditions.

None of the participants showed extended choice RTs (more than four standard deviations from the mean of participants' medians). Every participant's number of correct responses was tested in a binomial test against guessing with equal weights (i.e., choosing each option with a probability of .5). Assuming $\alpha = 5\%$, all participants significantly deviated from guessing and were thus included in the analysis.

Materials and procedures. The task setting, stimulus materials, and instructions were the same as in Experiment 1, except for the following changes. The total amount of time allotted to stock decisions was reduced from 20 min to 12 min. Instructions and stimulus display were modified slightly to accommodate the new sample-limit manipulation. In the condition without upper limit, all references to maximal sample size were deleted; the screen was divided by a thin line instead of the countdown circle. Sampling could continue until the screen area reserved for arrows was full (36 arrows per side). In the other condition, n_{max} varied randomly between 12 versus 8 on different trials.

Immediate feedback was provided in addition to a constant display of the current balance in ECU in the top center of the screen. After each choice, the new payoff-score flickered for 300 ms and was temporarily shown in red (after a false decision) or green (after a correct decision) for another 700 ms. Subsequently, the next trial started automatically. In a post-experimental questionnaire, participants answered two open-ended questions concerning their understanding of the trade-off ("Which were the two means you had to weigh in order to achieve the highest possible gain?") and the strategy used to resolve it ("Did you pursue a certain strategy to handle this trade-off? Please outline your approach"). On the back of the page, they also rated how clear the trade-off between speed and accuracy was to them (1 = very unclear, 5 = very clear) and how much thought they had given the trade-off between accuracy and speed (1 = not at all, 5 = very much). They indicated how clear it was to them that they had to weigh between fast and risky versus slow and accurate strategies, and they estimated their percentage of accurate choices, first across all choices and then for choices based on samples smaller than 8 arrows on either side. Finally, they rated the clarity of the correctness feedback (very incomprehensible vs. very comprehensible) in response to the question "Was the display of the feedback and balance comprehensible enough to detect correct and incorrect decisions?" and stated whether they had participated in an experiment of this kind before (yes/no).

Results and Discussion

Overall performance. During the 12-min session the average participant achieved M_k = 85.19 (SD_k = 41.69) trials. The average accuracy was somewhat reduced, M_a = .60 (SD_a = .14), corresponding to an average correct-choice proportion of M_c = .80 (SD_c = .07). The mean efficiency ($Eff = a \cdot k$) was M_{Eff} = 47.53 (SD_{Eff} = 14.35). Meanwhile, average sample size was M_n = 7.70 (SD_n = 3.95), which was slightly less than in the first experiment but in the same range.

Speed-accuracy trade-off. Overall, strategies were similarly conservative as in the first experiment. In Figure 4, accuracy rate M_a, the number of completed trials M_k, and achieved efficiency scores *Eff* are plotted as a function of individual sample size M_n. As the left plot illustrates, M_a increased when participants relied on larger samples $r_{n,a} = .71, p < .001$. However, number of trials M_k decreased with increasing M_n, $r_{n,k} = -.75$, $p < .001$, consistent with the expected trade-off. Again, the joint impact of speed and accuracy on total payoff M_{Eff} was dominated by M_k, as evident in a high positive correlation of M_{Eff} with k, $r_{k,Eff} = +.70, p < .001$, but a weak negative relation with M_a, $r_{a,Eff} = -.11$, $p = .352$. Participants who drew smaller samples achieved clearly higher payoffs $r_{n,Eff} = -.68, p < .001$.

Stability over time. To uncover any learning effects that the explicit feedback might have fostered, all performance indicators were again computed separately for the first and second half of the entire session time (i.e., before and after 360 s).[9] Table 2 summarizes the pertinent results.

Contrary to a possible learning improvement, sample size increased further from the first to the second half of the series from M_n =7.30 to M_n = 8.34, $t(77) = 3.05$, $p = .003$; 67% of all participants followed this trend. Increasing M_n did not serve to enhance accuracy, which actually showed a (hardly noticeable) decrease from M_a = .61 to M_a = .59 (shown by 60% of all participants), $t(77) = -1.34$, $p = .184$. Thus, despite the explicit feedback and the more modest accuracy during the first half, increasing sample size in the second half was again in vain.

Despite the increase in M_n, 56% of all participants managed to slightly increase the number of completed trials k from M_k = 40.76 to M_k = 43.44, $t(77) = 2.46$, $p = .016$, reflecting a learning effect on decision latencies. The mean additive constant for decision execution declined from 1,179 to 766 ms, $t(77) = 12.67$, $p < .001$. This trend was visible in 96% of all participants. Yet, because of the simultaneous increase in sample size M_n, the decrease in time per trial did not affect efficiency *Eff*, M_{Eff} = 23.63 to M_{Eff} = 23.31, $t(77) = .31$, $p = .755$ (consensus 50%).

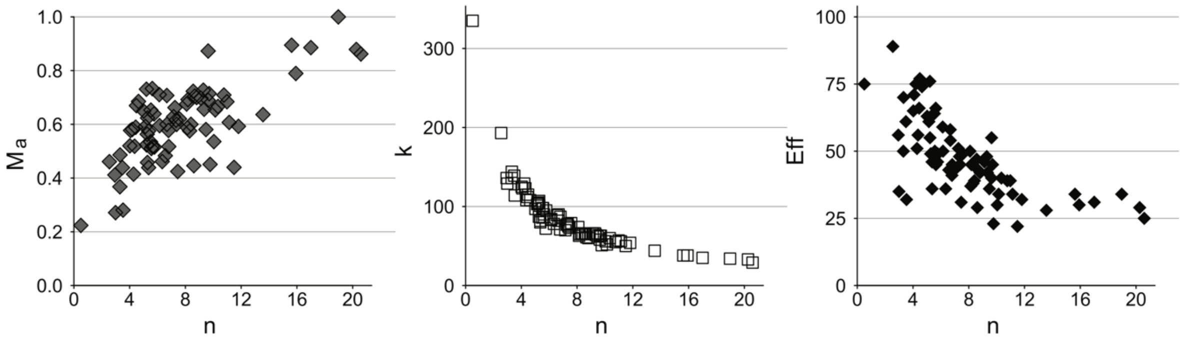

Figure 4. Mean individual accuracy $M_a = M_c - (1 - M_c)$, number k of trials completed in 12′, and efficiency *Eff* plotted as a function of mean individual sample size n.

Table 2.

Means (Standard Deviations in Parentheses) of Sample Size n, Accuracy a, Number of Completed Decisions k, Efficiency Scores Eff, and Selected Correlations Between These Variables, Calculated for the Total Series and for the Early and Late Half (First and Last 6 Minutes) Separately

Measure	Total	Early	Late
M_n (SD_n)	7.70 (3.95)	7.30 (3.77)	8.33 (4.83)
M_a (SD_a)	.60 (.14)	.61 (.15)	.59 (.18)
M_k (SD_k)	85.19 (41.69)	40.76 (18.65)	43.44 (23.82)
M_{Eff} (SD_{Eff})	47.53 (14.35)	23.63 (7.98)	23.31 (8.94)
$r_{n,a}$	.71	.50	.59
$r_{n,k}$	-.75	-.73	-.73
$r_{a,Eff}$	-.11	.20	.08
$r_{k,Eff}$	.70	.64	.67

[9] For counting the number of trials k, the very last trial was not included, because after the 720 s, the task was not interrupted but participants finished their current trial.

Impact of within-participants manipulations. To examine the impact of task difficulty (Δ_p = .1, .2, or .4) within participants, the nesting of upper limits called for hierarchical linear analyses, treating task difficulty as predictor and sample size M_n and accuracy M_a as criteria.

An ANOVA of M_n revealed a strong influence of task difficulty, $F(1, 77) = 79.91$, $p < .001$, $\eta^2 = .509$; 57% of all participants exhibited the ordinal pattern $M_{n(.1)} > M_{n(.2)} > M_{n(.4)}$ and only 1% exhibiting a reverse pattern. The same was true for accuracy M_a (right chart in Figure 5), $F(1, 77) = 468.03$, $p < .001$, $\eta^2 = .859$, with 74% of participants meeting the ordering $M_{a(.1)} < M_{a(.2)}$, $M_{a(.4)}$; nobody following the opposite ordering. In the limited condition, $n_{\max}$ (12 vs. 8) was an additional within-participants factor. Two *t*-tests served to examine its influence on M_n and M_a. Unsurprisingly, sample size was clearly higher (by almost 20%) when the limit was 12 per side ($M_n = 6.49$, $SD_n = 1.97$) than when it was 8 ($M_n = 5.47$, $SD_n = 1.17$), $t(31) = 5.96$, $p < .001$, $d = 1.05$, as shown by 88% of all participants. In this range, accuracy also increased (in 69% of all participants) from $n_{\max} = 8$ ($M_a = .54$, $SD = .13$) to $n_{\max} = 12$ ($M_a = .63$, $SD = .12$), $t(31) = 3.553$, $p < .001$, $d = 0.63$. Yet, due to the simultaneous increase in M_{n}, the total payoff differed hardly between $n_{\max} = 8$ ($M_{Eff} = 52.03$, $SD = 11.99$) and $n_{\max} = 12$ ($M_{Eff} = 51.88$, $SD = 12.08$), $t(31) = .55$, $p = .587$, $d = .10$. As in Experiment 1, a higher upper limit led participants to increase their sample size when they were given the chance to do so. This helped them to slightly increase their accuracy but did not affect their efficiency *Eff*.

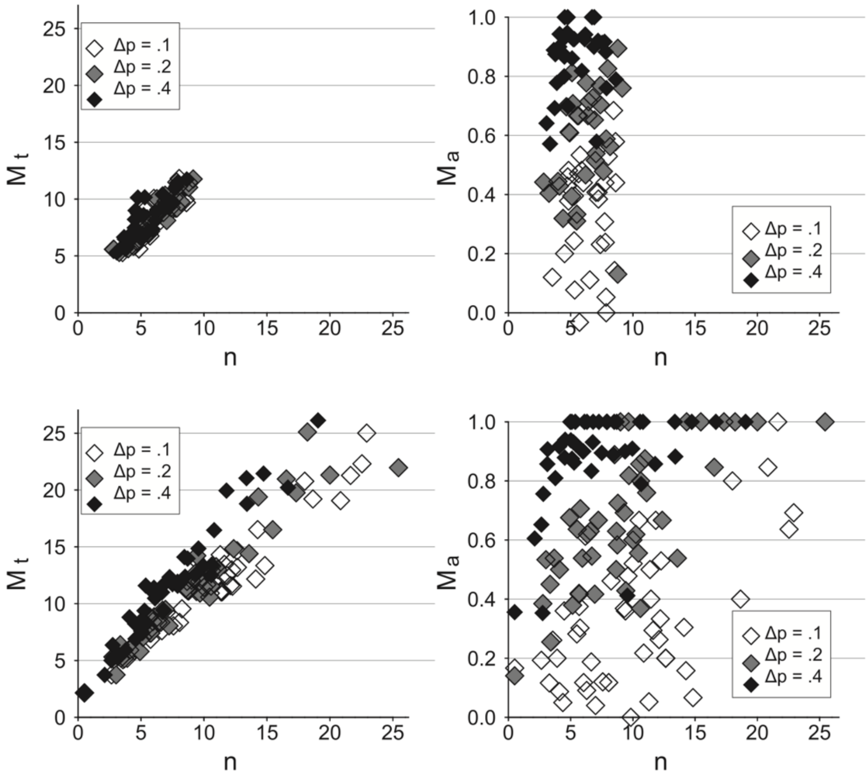

Figure 5. Average time M_t used per trial (left panels) and average accuracy M_a (right panels) as a function of task difficulty Δ_p, for the limited (n_{max} of 12 or 8; upper panels) and the unlimited task condition (lower panels).

Unlimited condition. The comparison of the limited and unlimited task conditions is of particular interest. In the absence of any upper limit, participants gathered even larger samples ($M_n = 8.90$, $SD_n = 4.63$) than in the limited condition ($M_n = 5.97$, $SD_n = 1.54$), $t(76) = 3.43$, *p*

< .001, d = .79. This increase in M_n came along with a nonsignificant tendency to complete fewer trials (M_k = 80.93, SD_k = 51.17) than in the limited condition (M_k = 91.31, SD_k = 21.24), $t(76) = -1.08$, p = .282, d = –.25, but little change in accuracy, (M_k = .58, SD_k = .11 vs. M_k = .61, SD_k = .16.), $t(76) = -.964$, p = .338, d = –.22. As a consequence, total payoff was lower in the unlimited (M_{Eff} = 44.85, SD_{Eff} = 14.97) than in the limited condition (M_{Eff} = 52.59, SD_{Eff} = 12.02), $t(76) = 2.43$, p = .017, d = .56. This analysis corroborates the conclusion that participants used every chance to (try to) increase accuracy and to decrease speed and thereby to further impair their performance.

Post experimental questionnaire. Responses to the open-ended comprehension question showed that 71% of participants expressed having an acceptable good grasp of the trade-off, 82% indicated to have had a clear or very clear grasp, 90% evaluated the feedback as clear or very clear. Strategy reports revealed that 19% of participants classified themselves as focusing on accuracy, 8% as using a balanced strategy, and only 5% as putting a focus on speed. A Mann–Whitney test showed that self-reported accuracy focusers (Mdn_n = 16.56) had drawn larger samples than self-reported speed focusers (Mdn_n = 10.5), T = 37.00, p < .001, r = –.35.

Although this seems to reflect partial insight into the accuracy bias, the vast majority (65%) interpreted the purposefully vague wording of the question as referring to a naïve sampling operation (e.g., counting green arrows). Two participants (3%) did not fill in their questionnaires.

Correctness self-ratings revealed that most participants overestimated the large-sample advantage. Whereas 81% of all participants underestimated their correctness across all trials by on average 8 percentage points (pp; M_{rated} = 72%, $SD_{r.}$ = 15pp, M_c = 80%, SD_c = 7pp), $t(75) = -6.40$, p < .001, d = .73; 88% underestimated their actual correctness on small samples (< 8 pairs) $M_{rated|<8}$ = 58% ($SD_{r.|<8}$ = 21pp), $t(75) = -9.20$, p < .001, d = 1.08, by on average 22pp (SD = 12pp). Underestimation was clearly larger for small samples than across all samples, $t(72) = 5.29$, p < .001, d = .62; 74% of all participants (73 without missing data) showed this pattern. The correlation between estimated and actual correctness was r = .69 for all choices but dropped to r = .24 for small samples.

Discussion and Interim Conclusions

To recap, in Experiment 1a through 1c, when the metacognitive regulation of speed and accuracy had to rely exclusively on the feeling of item difficulty, the monitoring process had failed to recognize the dominance of speed and the control process failed to exploit the possible payoff dramatically. Thus, although the instructions had explicitly focused on the trade-off, and although a difficulty cue similar to those of typical calibration experiments was available, there was no sign of metacognitive correction of a highly dysfunctional accuracy bias.

In Experiment 2, though, metacognitive functions could rely on several additional cues. The explicit feedback allowed participants to assess correlations across trials between n_j and c_j, or else between n_j and a_j, or correlations between the feeling of difficulty and n_j, c_j, and t_j (the time taken for a choice), or between the feeling-of-difficulty and the given sample differences $|\Delta_j|$*. Moreover, the feedback allowed them to compare the obtained correctness c_j and the sample size n_j or time t_j invested on subsets of trials with different upper limit or difficulty. Despite all these additional metacognitive cues, there was little sign of improvement.

The manipulation of limited versus unlimited sampling showed that the accuracy bias and speed neglect observed in Experiments 1a through to 1c were not due to too high upper sample limits nudging participants into oversampling. If anything, in Experiment 2 free sampling without an upper limit led to even larger sample sizes than in the limited condition.

More generally, the failure of explicit correctness feedback to trigger better speed regulation and the ironic finding that accuracy tended to decrease from Experiment 1 to Experiment 2, suggests that outcome feedback may inhibit rather than facilitate the meta-cognitive process. Outcome feedback may be too radical (making accuracy much more salient than time or information costs) and too late (to influence the sampling truncation decision). From hindsight, participants not only underestimated accuracy in general but they underestimated the accuracy on small samples relative to large samples. They fell prey not only to preconceptions favoring large samples but also misjudged the correlations between sample size and accuracy compared with what they had experienced. These findings suggest that the deficit cannot be exclusively due to insufficient evaluability of speed but must, to some extent, originate in illusory correlations of sample size and a clearly evaluable accuracy criterion.

Last, the post-experimental questionnaire of Experiment 2 served to rule out the suspicion that the low trade-off performance in Experiment 1 may have been the result of ill-understood task instructions or insufficient motivation. Although the generally high accuracy scores and their sensitivity to experimental manipulations testifies to careful responding and high motivation, the questionnaire responses speak against the possibility that participants did not understand the trade-off structure of the payoff scheme.

Experiment 3

Experiment 3 was designed to assess participants' sensitivity to the payoff scheme. We manipulated the payoffs for correct and incorrect choices to rule out another alternative account. Prevailing conservative strategies that place more weight on accuracy than on speed may primarily reflect loss aversion (Kahneman & Tversky, 2000) rather than a metacognitive deficit. If the negative subjective value of losing 1 ECU is clearly lower than the subjective value of winning 1 ECU is high, then the high weight placed on accuracy may be deliberate. The accuracy bias may constitute a remedy to loss aversion, independently of the metacognitive regulation of the relative dependence of speed and accuracy on samples size. For a straightforward experimental test of this idea, we created a completely riskless condition in which participants could only win on correct trials; no loss was possible after incorrect choices. Such a change in payoff structure should eliminate any loss-aversion effect.

Payoffs were manipulated across two blocks of trials. In the first block, all participants were exposed to the same task setting as in the unlimited sampling condition of Experiment 2 (with unit payoff of ±1 ECU). In the second block, the reward for correct choices remained invariant (+1 ECU) but the negative payoff for incorrect choices was eliminated in one condition (0 ECU) but tripled (–3 ECU) in another condition. This manipulation enabled us to examine two issues, namely, whether the no-loss condition eliminates the accuracy bias, and whether the high-loss payoff scheme renders the increased sample size worthwhile.

Method

Participants and design. A total of 80 participants (63 female) with a mean age of 23.4 (*SD* = 4.73) were recruited for an experiment that was conducted in the second position of an hour-long session consisting of four experiments (concerning contingency learning, political persuasion, moral licensing). Some participants met a university program requirement, while others received a monetary base reward of 5€ for the whole session. In addition, all participants received a performance-contingent payment (3€ on average). One participant, whose median choice latency (after terminating the sample) was more than four standard deviations longer than the mean of all participants' medians, was excluded from analyses. Three further participants were excluded because their correctness did not differ significantly from guessing during the first block (tested in a binomial test with $\alpha = .05$). In total, 76 participants were included in the analysis.

Two experimental conditions that had no upper sample-size limit varied in terms of payoff for correct and incorrect choices. During the first of two blocks, participants in both conditions received the same unit payoff (+1 ECU for correct and incorrect responses, respectively) as in all previous experiments. In the second block, both conditions continued to receive –1 ECU for correct choices, while the payoff for incorrect choices was changed to be –3 ECU in the high-loss condition as compared with 0 ECU in the reward-only condition.

Materials and procedures. The task setting was the same as in previous experiments except that two 10-min blocks were used to implement the payoff parameter change, with the second block preceded by additional instructions explaining the change. The feedback during the second block changed from 1 ECU being deducted from the balance for incorrect decisions to 0 or 3 ECU deducted, depending on the condition. Instead of a paper-and-pencil questionnaire, four questions were appended to the experimental computer dialog. The first question inquired about possible strategy changes: "Was there a difference in your strategy for resolving the trade-off between the first and second half of the experiment? If so, how did your strategy change?". The remaining questions had participants estimate their accuracy for Block 1, for Block 2, and finally their small-sample accuracy across both blocks: "How often do you think did you choose the correct stock in the first half/second half?" and ". . . across both blocks when only considering samples that were comprised of less than 8 days, that is, less than 8 arrows per options?".

Results and Discussion

Block 1: Overall performance. In the first block, participants went through $M_k = 59.30$ decisions, on average ($SD_k = 25.19$), attaining a mean accuracy of $Ma = .67$ ($SD_a = .14$), which corresponds to an average correctness rate of $M_c = .83$ ($SD = .09$) based on samples sized $M_n = 9.62$ on average ($SD_n = 5.08$). They achieved a mean overall payoff of $M_{Eff} = 37.86$ ($SD_{Eff} = 14.25$). Recall that the first block represents an almost identical replication of the unlimited condition of Experiment 2, except for the slightly shorter total time. The results were indeed very similar (see Table 3). If anything, sampling strategies tended to be more cautious (i.e., relying on larger *n*).

Block 1: Speed-accuracy trade-off. In the first block, individual participants' total payoff *Eff* again most strongly correlated with number *k* of decisions completed, $r_{Eff,k} = .79$, $p < .001$, but only weakly with accuracy, $r_{Eff,a} = .12$, $p = .320$. Interpersonal variation in sample size again drove these results (see Figure 6). Although participants who drew larger samples tended to be more accurate, $r_{n,a} = .32$, $p < .005$, increasing *n* reduced the number of trials more strongly, $r_{n,k} = -.82$, $p < .001$. As a result, the correlation of sample size *n* and overall payoff *Eff* was again strongly negative, $r_{Eff,n} = -.74$, $p < .001$.

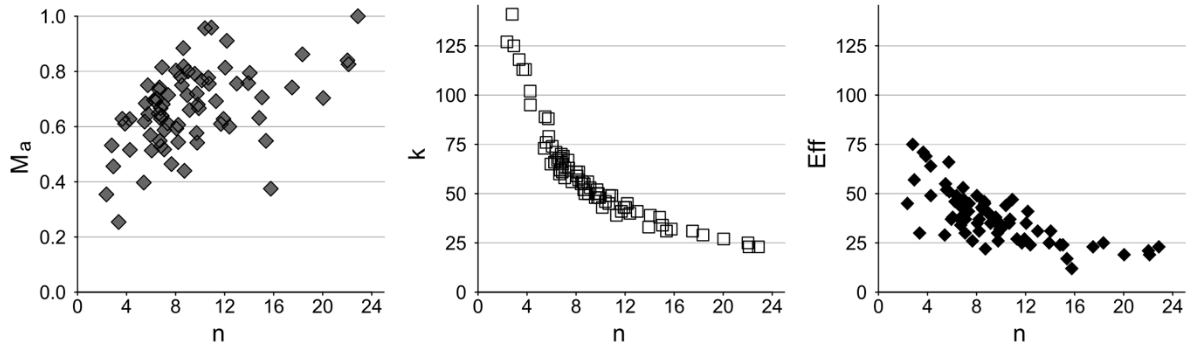

Figure 6. Individual mean accuracy $M_a = M_c - (1 - M_c)$, number of trials k completed in 12', and efficiency *Eff* in the first block of Experiment 3 as a function of mean individual sample sizes *Mn*.

Block 1: Impact of task difficulty. As in prior experiments, task difficulty Δ_p (treated as a linear predictor) had a strong influence (see Figure 7) on sample size M_n, $F(1, 75) = 108.77$, $p = .001$, $\eta^2 = .592$ (with 76% of all participants showing the ordinal pattern $M_{n(.1)} > M_{n(.2)} > M_{n(.4)}$ and nobody following the opposite pattern), and on mean accuracy rates *Ma*, $F(1, 75) = 624.13$, $p < .001$, $\eta^2 = .893$ (with 87% showing the ordinal pattern $M_{a(.1)} < M_{a(.2)} < M_{n(.4)}$ and no one following the opposite pattern). Again, this systematic pattern testifies to high motivation, task comprehension, and sensitivity to effectively manipulated sample aspects.

Block 2: Impact of payoff manipulations. For the second block, we needed to adjust the calculation of accuracy a, since gain and loss were asymmetrical. We calculated $a = c * payoff_{gain}$ - $(1\text{-}c)* payoff_{loss}$. In concrete terms, the loss-payoff needed to be set to either 0 or –3, while the gain-part of the formula remained as before. Mixed ANOVAs were conducted to examine the joint impact of the twofold manipulation of payoffs over blocks (1 vs. 2) and between the high-loss and the reward-only conditions. As the first block was identical, the payoff manipulation should be manifested in an interaction of both factors. The noncomparable payoffs in both conditions precluded an ANOVA of the total payoff measure M_{Eff}.

Table 3 provides all relevant descriptive statistics and test statistics. The Critical Blocks × Conditions interaction was mainly due to further increase in sample size and striving for high accuracy in the high-loss condition rather than due to decreasing sample size and less conservative sampling in the reward-only condition. In the high-loss condition, a roughly 50% increase (in M_n from 10.2 to 15.5; see Table 3) allowed M_c to increase slightly (from .83 to .87). This kept accuracy at a level clearly above zero (M_a changing from .67 to .46), despite the loss-prone payoff. However, the new accuracy of .46 is hardly different from the accuracy (of $M_a = .41$) that they would have obtained had they continued to achieve $M_c = .87$ with smaller sample size of $M_n = 10.2$ in spite of the higher loss. Thus, the reasonable amplification of cautious sampling in the high-loss condition was hardly worthwhile. In the reward-only condition, a roughly 20% decrease of M_n (from 9.9 to 7.0) came at the cost of a hardly noticeable decline in M_c (from .83 to .78), which however even increased M_a due to the no-loss context (change from .67 to .78).

Both conditions underperformed severely. Especially the reward-only condition remained too conservative when no loss was possible. As in all prior experiments, total payoff was positively related to k but less so to a. In the high-loss condition, the readily adopted conservative shift was mostly in vain. Generally, the fastest participants earned the highest payoffs. Participants did not exploit the chance to greatly enhance their total payoff.

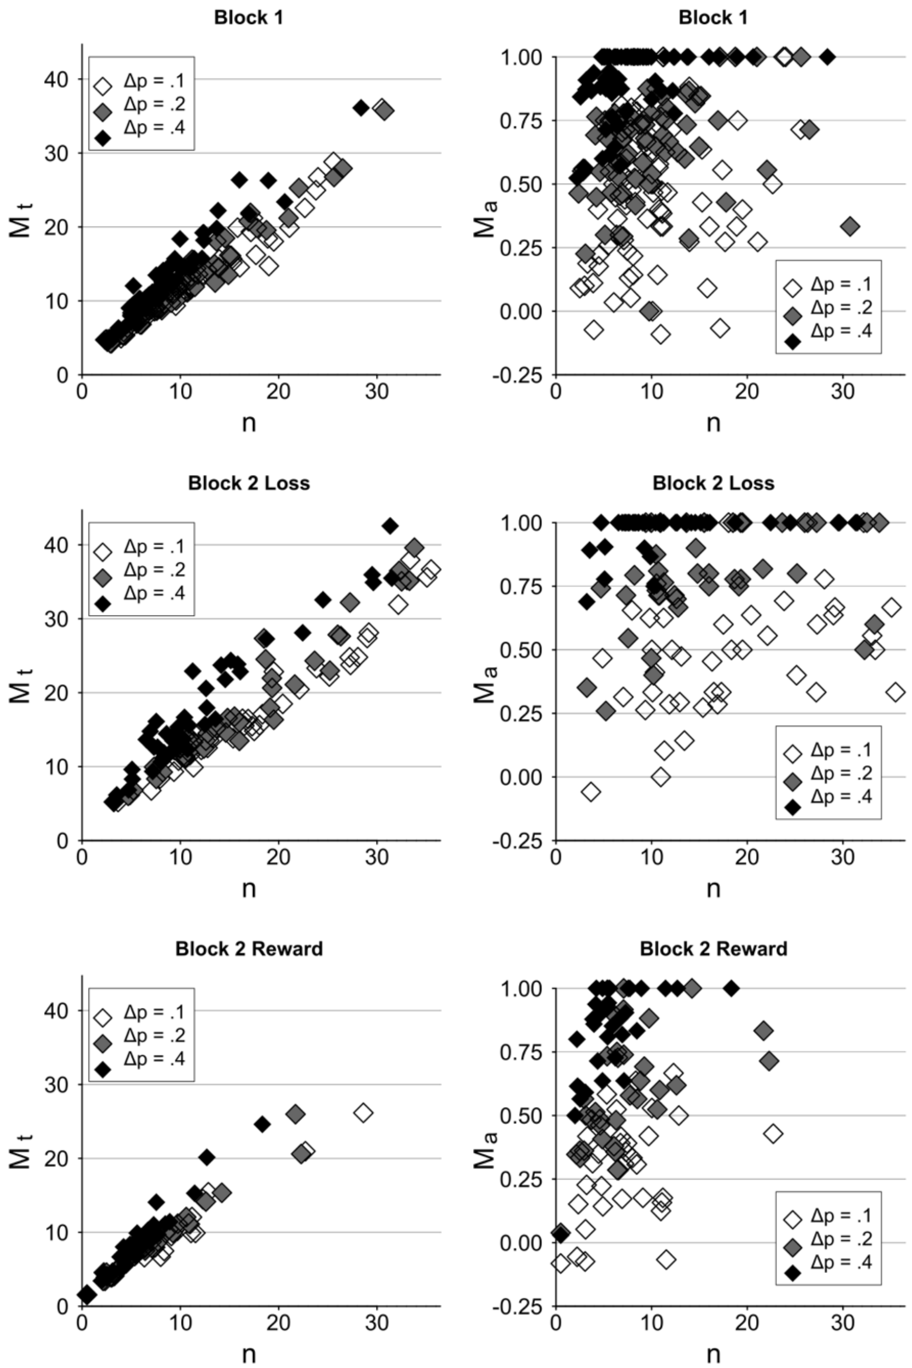

Figure 7. Average time M_t used per trial (left panels) and average accuracy scores M_a (right panels) as a function of task difficulty Δ_p, for Block 1 (upper panels) and for the high-loss condition (middle panels) and the reward-only condition (lower panels) of Block 2.

Table 3

Means (Standard Deviations) for Sample Size n, Accuracy a, Number of Completed Decisions k, Efficiency Scores Eff, and Intercorrelations, Calculated Separately for Blocks and Payoff Conditions

DV	Condition	Block 1	Block2	Interaction F	p-value	η^2
		M (SD)	M (SD)			
n	Loss	10.19 (5.66)	15.51 (8.04)	$F(1,74) = 70.12$	< .001	.086
	Reward	8.91 (4.24)	6.98 (4.51)			
c	Loss	.83 (.08)	.87 (.07)	$F(1,74) = 28.84$	< .001	.110
	Reward	.84 (.06)	.78 (.08)			
a	Loss	.65 (.17)	.46 (.30)	$F(1,74) = 29.94$	< .001	.120
	Reward	.69 (.11)	.78 (.08)			
k	Loss	57.52 (25.42)	42.69 (21.47)	$F(1,74) = 35.88$	< .001	.089
	Reward	61.50 (25.10)	92.38 (63.71)			
Eff	Loss	35.43 (13.12)	16.31 (13.81)	$F(1,74) = 75.94$	< .001	.265
	Reward	40.85 (13.12)	68.06 (32.72)			
$r_{n,a}$	Loss	.28	.42			
	Reward	.48	.48			
$r_{n,k}$	Loss	-.83	-.83			
	Reward	-.81	-.66			
$r_{k,a}$	Loss	-.44	-.54			
	Reward	-.52	-.80			
$r_{Eff,a}$	Loss	.21	.78			
	Reward	-.13	-.72			
$r_{Eff,k}$	Loss	.72	-.15			
	Reward	.89	.97			

Note. Test statistics refer to Blocks × Payoff Condition interactions.

Questionnaire responses. The open-ended strategy question was scored for whether a strategy change was indicated and, if so, for the direction of the strategy change. A large majority of participants (78.2%) indicated to have changed their strategy for the second block.

Only 7.7% reported no change; the remaining 14.1% did not give any indication of their strategy focus. In the reward-only condition, 85.2% of those who indicated a change reported it to be a change to a faster strategy, with the remaining 14.8% split evenly between change to a balanced strategy and a change toward a more accuracy-focused strategy. In the high-loss condition, reported strategies changed almost exclusively (97.1%) to more accuracy-focused strategies, only one participant indicated a change to a speedier strategy. These introspective data leave no doubt that participants in both conditions understood the trade-off and what strategy shift was called for by the payoff structure they encountered. However, the performance data highlight their metacognitive inability to regulate the impact of strategy shifts on the resulting payoff.

Participants estimated their overall correctness rate to be M_{rated} = 74% on average ($SD_{r.}$ = 14pp) for Block 1 and M_{rated} = 77% ($SD_{r.}$ = 13pp) for Block 2, which was again below their actual performance (Block 1: M_c = 83%, SD_c = 7pp; Block 2: M_c = 83%, SD = 9pp). Separate correctness estimates for small samples (< 8 arrows on each side) were again subject to vast underestimation: $M_{rated|<8}$ = 55% ($SD_{r.|<8}$ = 23pp) compared with M_c = 86% (SD_c = 5pp). Their tendency to underestimate their own accuracy was much higher for small samples (31pp) than for all samples (8pp). A two-way mixed ANOVA with estimate type (overall vs. small-sample estimates, for both blocks) as a within-groups factor and payoff conditions (loss vs. reward-only) as between-groups factor was run on estimate deviations from actual correctness rates. The degree of underestimation was clearly stronger for small samples than across all samples, $F(1, 74) = 93.87$, $p < .001$, $\eta^2 = .309$.

The lack of a significant interaction, $F(1, 74) = 1.90$, $p = .172$, $\eta^2 = .009$, indicates that the payoff manipulation did not moderate the tendency to underestimate accuracy on small samples more than on big samples. There was also no significant main effect of condition, $F(1, 74) = 2.34$, $p = .130$, $\eta^2 = .020$. As in Experiment 2, then, the excessive accuracy bias came along with a general underestimation of accuracy, and a grave underestimation for small samples.

Thus, the results of Experiment 3 corroborate the contention that participants did understand the speed-accuracy trade-off and the payoff scheme of the sequential choice tasks. They not only responded to the changing payoffs from the first to the second block, but their introspective reports almost unanimously show that they were aware of the strategic implications of the changing payoff. However, despite the alerting influence of the payoff manipulation, the resulting improvement was very modest and the strategic shift in the appropriate direction was far from adequate. Even when no loss was possible in a reward-only condition, participants remained cautious and reluctant to adopt clear speed-oriented strategies, despite the feedback highlighting their high accuracy. The amplified accuracy bias in the high-loss condition was largely in vain. Had participants not increased their sample size, their payoff would have been comparable.

Experiment 4

The findings presented so far demonstrate that people are not responsive to speed-inviting cues, they readily follow accuracy-inviting cues, but in either case they lack an appropriate metacognitive sense of speed and accuracy. Nevertheless, it may be possible to ameliorate the metacognitive deficit in more auspicious task settings that minimize an initial oversampling bias. In the present sequential choice task in particular, the dysfunctional over-sampling may be partially due to a superficial aspect of the sampling procedure. Recall that samples grew automatically if participants kept pressing the space bar. This may have created a sort of inertia effect; participants may have waited too long until they released the space bar to stop an ongoing sampling process. Samples may be smaller when every new element of the sample must be solicited actively by a new keystroke.

A final experiment was devoted to testing this idea. In two experimental conditions, sampling either continued automatically (as before) or had to be solicited elementwise. Although a reduced accuracy bias in the latter condition would not undo or excuse the metacognitive deficit, it might demonstrate that the strength of the accuracy bias is partially due to a procedural feature. Yet, convinced that such a superficial feature only played a subsidiary role, we expected the oversampling tendency to persevere, even though shifting to an active sampling procedure might somewhat reduce the average sample size.

Method

Participants and design. In total, we recruited 104 participants, of whom 80 were female; 95 were students and 17 students of Psychology; the average age was 24.9 (*SD* = 7.49). For the first 73 participants, the experiment was embedded, in second position, in a series of four studies conducted in an hour-long session (involving belief updating in observing sport results, impression formation from personality traits, and the kama muta feeling)[10]. The other 31 participants were part of another 1-hr session, working on the present experiment second in a series of three studies (involving impression formation from personality traits and kama muta). Participants received a basic payment of 5€ (or course credit) plus a performance-contingent reward (of approximately 3€ for the average participant).

One participant had to be excluded from analyses for heavily extended response latencies (median response latency exceeding the average median response latencies by more than four standard deviations). Another three participants' correctness rates did not differ significantly from guessing (tested participant-wise, binomial test with $\alpha = .05$; cf. preceding experiments). In total, 100 participants were included. Two separate groups were exposed to the same standard procedure as in the unlimited condition of Experiment 2, with unit payoff (±1 ECU) and within-participants manipulation of choice difficulty ($\Delta_p = .10$ vs. .20 vs. .40).

Materials and procedures. Whereas in one group passive sampling did not stop until the space bar was released, active sampling in the other group was contingent on a new key press of the space bar for every new element. In the passive-sampling condition, the instructions, stimulus materials, and stimulus presentation format and procedures were all the same, administered by the same software, as in Experiment 2 (unlimited condition) or Block 1 of Experiment 3. The presentation format in the active-sampling condition was modified; an extra keystroke was required to elicit another pair of observations for both investment funds with a minimal refractory period of 300 ms between successive pairs. In all other respects (generation of samples, alternating presentation of arrows per option, format, etc.) the procedure was the same across experimental conditions. At any given position between a minimum of 1 arrow and a maximum of 72 arrows (36 per side, in an invisible grid of 6 x 6) participants could truncate sampling by directly pressing one of two response keys (left and right arrow key to choose the left and right option, respectively).

Results and Discussion

Overall performance. The average sample size was slightly reduced in the active-sampling condition ($M_n = 7.49$, $SD_n = 3.61$), compared with the passive-sampling condition, ($M_n = 8.23$, $SD_n = 3.35$). Although consistent with intuition, this decline in oversampling was not substantial. The tendency of the active-sampling group to speed up a little more than the passive-sampling group served to somewhat increase the numbers of trials completed within the available time of 12 min, $M_k = 92.46$ ($SD_k = 47.63$) versus $M_k = 76.22$ ($SD_k = 23.06$), and to slightly reduce accuracy scores, $M_a = .60$ ($SD_a = .14$) versus $M_a = .63$ ($SD_a = .11$), corresponding to mean correctness proportions of $M_c = .80$ ($SD_c = .07$) versus $M_c = .81$ ($SD_c = .06$). As a consequence, the mean total payoff $Eff = a \cdot k$, $M_{Eff} = 52.64$ ($SD_{Eff} = 17.83$) versus $M_{Eff} = 47.16$ ($SD_{Eff} = 11.86$), also increased slightly from active to passive sampling. Thus, the manipulation was apparently effective. Active sampling served to ameliorate the oversampling bias.

[10] Kama muta is a technical term in emotion research for the 'being moved' emotion.

Most active samplers' self-determined presentation speed of the *n* elements per choice trial (i.e., the average time taken per observation) was shorter than the 500 ms/item in the passive sampling condition (see Figure 8). This increase in stimulus presentation speed, in addition to the decrease in average *n*, helped to increase the number *k* of trials completed and, hence, the total payoff *Eff* of active compared with passive samplers.

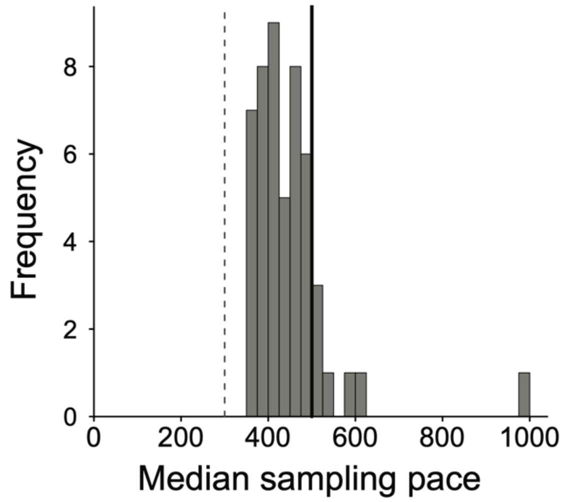

Figure 8. Individual median sampling pace (in ms) for active sampling. The dashed line marks the refractory time of 300 ms; the thick black line indicates the fixed passive sampling pace.

Speed-accuracy trade-off. However, despite the reduced oversampling in the active-sampling condition, the improvement in metacognitive speed regulation remained very modest. In both conditions, *Eff* depended more strongly on *k* than on *a*. Figure 9 shows that although accuracy among passive and active samplers increased with sample size, $r_{n,a} = .49$, $p < .001$ and $r_{n,a} = .60$, $p < .001$, respectively, number of completed choices decreased more strongly with sample size, $r_{n,k} = -.92$, $p < .001$ and $r_{n,k} = -.75$, $p < .001$. Consequently, total payoff *Eff* not only correlated negatively with mean individual *n*, for passive samplers ($r_{n,Eff} = -.81$, $p < .001$) as well as for active samplers ($r_{n,Eff} = -.60$, $p < .001$). *Eff* also showed no correlation with *a* (passive: $r_{a,Eff} = .00$, $p = .998$; active: $r_{a,Eff} = .02$, $p = .916$) and correlated positively with *k* (passive: $r_{k,Eff} = .80$, $p < .001$; active: $r_{k,Eff} = .74$, $p < .001$). Again, the relation of *n* to *Eff* (right panel) largely resembled the relation of *k* to *Eff* (middle panel). Thus, regardless of the slightly smaller sample size among active than passive samplers, the massive accuracy bias and speed neglect persisted.

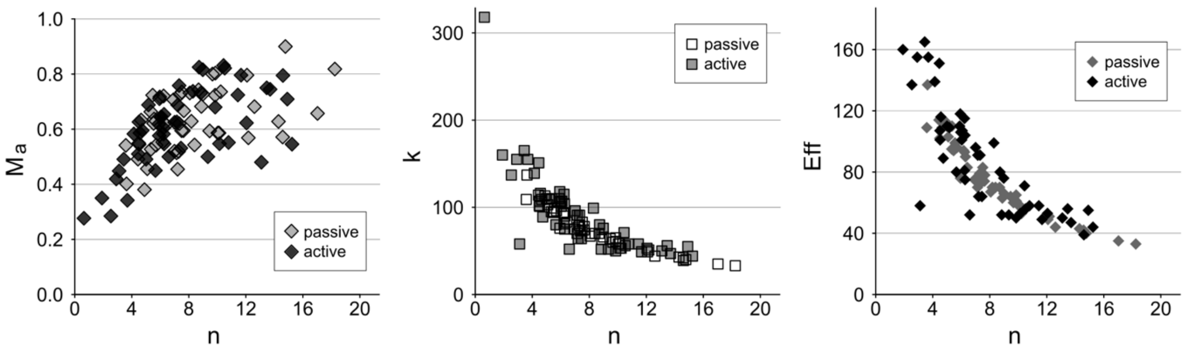

Figure 9. Individual mean accuracy $M_a = M_c - (1 - M_c)$, number of trials *k* completed in 12 minutes, and efficiency *Eff* as a function of mean individual sample size *n*.

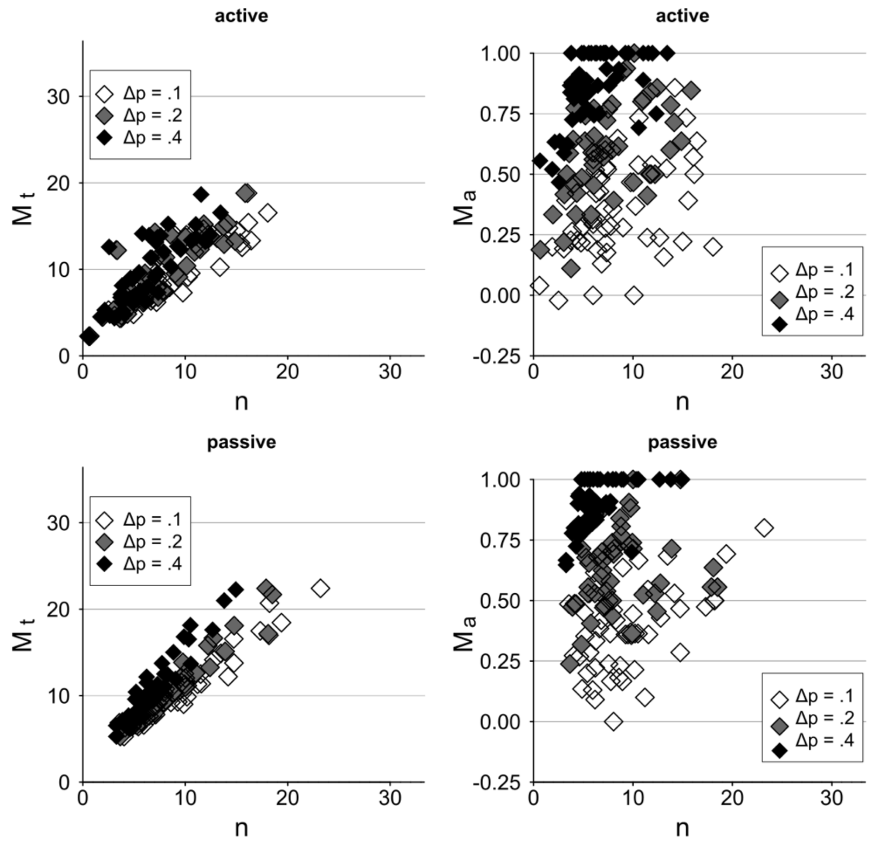

Figure 10. Average time M_t per sample (left panels) and average accuracy scores M_a (right panels) as a function of task difficulty Δ_p, for active (upper panels) and passive (lower panels) samplers.

The comparison between conditions was complicated by the fact that the active sampling condition had the opportunity to increase or decrease their sampling pace (i.e., the rate at which arrows appeared during the sampling phase). We therefore computed the average effectiveness *Eff* per sampled arrow—a measure which serves to compare effectiveness independent of individual sampling pace. The same overall pattern is visible: active samplers had a slightly but not substantially higher *Eff*/arrow of M = .05 (SD = .03) compared with passive samplers' M = .04 (SD = .01). The persistent failure to maximize *Eff* is thus not at all peculiar to passive sampling.

Impact of task difficulty within participants. Within-participants analyses of choice time (left panels of Figure 10) and accuracy (right panels) as a function of choice difficulty (Δ_p = .10 vs. .20 vs. .40) corroborate the robustness of findings across conditions.

General Discussion

Starting from the demonstration of a predictable accuracy bias in sample-based choices, the reported experiments and experimental conditions testify to the metacognitive inability to understand and correct for the serious failure to regulate speed and accuracy on this particular trade-off task. Before we discuss potential reasons for the persistence of the accuracy bias, we quantify the performance loss on a reasonable standard scale. As mentioned at the outset, we use a simplified Bayesian decision rule (Edwards, 1965), the tie breaker, as a benchmark.

A Normative Benchmark

The tie-breaker strategy offers a simple normative standard and a sensible way of positioning each participant's suboptimal performance. A computer algorithm assessed the payoff that each individual participant would have obtained (given the individual decision time) had they stopped sampling at a range of critical differences between the two options. The optimal critical sample difference that maximized the total payoff turned out to be very small, ranging between the first difference in a pair of arrows for both choice options and even a single arrow indicating a gain or loss of the first arrow in a pair. Table 4 reports the average individual *n*, *a*, and total *Eff* in comparison with the tie breaker's optimum for all conditions (of Experiments 2, 3, and 4) without upper sample size limit.

It is evident from Table 4 that participants gathered far too large samples, three to five times as large as the tie-breaker mean *n*. Their oversampling did indeed serve to increase their accuracy; they were about twice as accurate as the tie breaker. However, this accuracy gain was overridden by a comparably greater reduction of speed. They missed the opportunity to complete so many more choices that their total payoff was less than half the payoff they could have reached by applying a simple algorithm (see the last two columns). Under unit payoff conditions ((±1 ECU), the average participant achieved, on average, between 37% and 41% of their individual benchmark payoff. Even the best participants' performance (rightmost column) hardly exceeded 70% of the optimum. The marked accuracy bias persisted, despite regular feedback and hints to speed dominance.

Table 4

Mean Individual Performance Compared to Optimal Tie-Breaker Performance

Experimental condition	Mean *n* (*SD*)	Optimal *n* (*SD*)	Mean *a* (*SD*)	Optimal *a* (*SD*)	Mean *Eff*/ optimal *Eff* (*SD*)	Maximal *Eff* / optimal *Eff*
Experiment 2						
No upper limit	8.90 (4.63)	2.44(.48)	.61 (.16)	.27 (.14)	.39 (.12)	.67
Experiment 3						
Block 1	9.62 (5.08)	2.44 (0.78)	.67 (.14)	.28(.15)	.41(.13)	.72
Block 2 High-Loss	15.65 (8.09)	8.47(5.10)	.48 (.27)	.55 (.19)	.48 (.19)	.77
Block 2 Reward-Only	6.98 (4.51)	0.50 (0)	.78 (.08)	.63 (.05)	.17 (.06)	.37
Experiment 4						
Active	7.49 (3.61)	1.40 (1.06)	.60 (.14)	.39 (.13)	.37 (.12)	.61
Passive	8.23 (3.35)	1.69 (1.34)	.63 (.11)	.39 (.16)	.37 (.13)	.70

Note. Optimal values take individual's reaction times and confusion errors into account. The two right-most columns indicate the actually achieved payoff relative to the potential maximum. Only data from conditions with unconstrained sample size were included.

Forgone Chances for Correction

Granting that cognitive biases and violations of normative rules are ubiquitous and hard to avoid, rational behavior depends heavily on a subsequent metacognitive stage during which adaptive agents can monitor and control their own cognitive performance. Central to our metacognitive perspective is the conviction that a comprehensive approach to rationality must go beyond the mere occurrence of cognitive biases. It must also account for the extent to which adaptive agents can diagnose and correct for irrational behavior through metacognitive control. Thus, we next discuss the forgone chances to overcome the accuracy bias.

Even without explicit accuracy feedback in Experiment 1, participants had a fair chance to attune sample size to their metacognitive feeling of choice difficulty, analogous to confidence ratings in a calibration experiment. As the manipulation of Δ_p produced sufficient variation in difficulty, participants could easily try out faster strategies on the easiest choice tasks, for which they attained virtually perfect accuracy. Likewise, variation in samples size increased through manipulation of the upper sample size limit. When regular accuracy feedback was available (from the second experiment onward), the stimulus input allowed participants to assess correlations between sample size and accuracy, between sample size and required time, and correlations between all these variables and the feeling of choice difficulty. Moreover, they could assess and compare the average sample size and the average accuracy of choices with different upper limit. When changing payoff parameters enforced increases or decreases in sample size, or when the manipulation of active versus passive sampling motivated faster ground speed, there were further opportunities to experience the benefits of less conservative sampling strategies.

Yet, none of these manipulations and decision aids succeeded in inducing a fundamental improvement in speed regulation. Even when negative payoffs for incorrect choices were fully eliminated such that participants could only win but never loose, a half-hearted decrease in sample size was too weak to overcome the accuracy bias. Conversely, they hardly profited from a further increase in sample size when the negative payoff for incorrect choices was tripled.

Understanding the Metacognitive Deficit

Yet, all these chances to detect and exploit the speed advantage in sample-based choices were forgone. Participants ignored or neglected metacognitive cues that should have facilitated speed; they readily sampled even more information when metacognitive cues favored accuracy.

Illusory belief in the law of large numbers. The conspicuous inability to recognize the speed advantage served to distort the participants' assessment of the relationship between sample size and accuracy. They slightly underestimated their accuracy, but they generally had a good idea of its range. They most strongly underestimated choice accuracy for small samples. Thus, a lack of insight into the joint dependency of speed and accuracy on sample size lay at the heart of the problem. Although the actual correlation between (self-truncated) sample size and the size of the sampled differences was negative (across, choices, within participants), participants fell prey to misconceptions about the correlation between sample size and accuracy. Small samples remained so small (i.e., were truncated so early) because they happened to exhibit strong evidence from the very beginning (Prager et al., 2018). In contrast, large self-truncated samples are conflict-prone and devoid of clear-cut evidence. To some extent, then, the failure to correct for a serious accuracy bias seems to reflect a deeply entrenched, overgeneralized belief in the law of large numbers.

What renders optimal speed regulation counterintuitive is, apparently, that a strategy that optimizes overall success calls for sacrificing accuracy. To maximize total payoff, participants must provide fast and risky responses on those trials in which observations give a first hint but do not clearly suggest an obvious choice. This gap between optimal behavior on single trials versus whole sequences of trials, which is familiar from game theory, was central to the important research program on melioration 3 decades ago (Herrnstein & Prelec, 1991). Even though the typical melioration task structure bears little similarity with ours beyond the presence of what one could call a trial level and a sequence level, the parallel in suboptimality is striking.

Evaluability. What must have further exacerbated the bias in favor of accuracy and against speed is an intrinsic feature of the sequential design for which Hsee and Zhang (2010) have coined the term evaluability. The accuracy feedback at trial level is much easier to evaluate and more noticeable than the feedback about the proportion of the total time elapsing (i.e., the time cost of the information). Whether a choice is correct or not can be evaluated immediately and unequivocally on every trial. In contrast, speed and efficiency can only be assessed at the aggregate level, at the end of a longer trial sequence. Even then, there is no natural benchmark to distinguish fast from slow responding. Speed is a relative rather than an absolute attribute, which can only be evaluated in comparison to other levels of speed (Hsee, 1996).

This design feature could be manipulated in principle. Pursuing this line of reasoning, we are currently conducting new experiments that manipulate the relative evaluability of speed and accuracy. For instance, in one experiment (McCaughey, Prager & Fiedler, 2020), participants receive online feedback about the information costs of growing samples indicated in the same format and currency as the payoff they receive for a correct choice. Inspired by Hsee and Zhang's (2010) notion of evaluability, we expect the accuracy bias to decrease considerably if the costs of insufficient speed on each choice trial can be evaluated as transparently as the payoff gained for accurate choices. In any case, an intriguing empirical question is whether and to what degree it is possible to reduce the accuracy bias by increasing the evaluability of speed.

Intrinsic difficulty of nonlinear functional relationships. With reference to another strand of literature, the failure to regulate the function $Eff = a \cdot k$ relates to the often-noted impaired understanding of nonlinear functional relationships (e.g., Larrick & Soll, 2008; Hogarth, Mukherjee, & Soyer, 2013; Irwin, Smith, & Mayfield, 1956; Juslin, Nilsson, & Winman, 2009; Peer & Gamliel, 2013; Svenson & Eriksson, 2017). If our participants did erroneously assume a linear relation between sample size and accuracy, underestimating both the fast initial rise and the later flattening of accuracy when increasing sample size, this might indeed explain their failure to understand the lesser contingency of accuracy than speed on sample size.

Surplus value of accuracy. The social surplus meaning of accuracy implies carefulness and responsibility; it may be valued higher than speed, which has problematic connotations such as risk, carelessness, or even sloppiness. The bias in favor of accuracy may to some degree reflect such a quasi-moral norm. If so, a relative shift toward speed may perhaps occur when the context (e.g., during a gambling task) gives more positive connotations to speed.

Metacognitive myopia. Apart from these specific problems, the persistence of a strong accuracy bias in sample-based choices fits neatly into a growing body of evidence on metacognitive myopia, conceived as a general impediment to rational JDM, due to a lack of critical metareasoning (Fiedler, 2012; Fiedler, Ackerman, & Scarampi, 2019; Spiller, 2019). Metacognitive myopia refers to the reluctance to reflect on the constraints and sources of invalidity in the environment and in one's own higher order cognitive functions, even when these constraints are incontestable and easy to understand. Pertinent research has shown, for instance, that people mistake fully redundant repetitions for useful evidence (Unkelbach, Fiedler, & Freytag, 2007), they continue to follow obviously wrong communications (Fiedler, Krüger, Koch, & Kutzner, 2020), invalid advice (Fiedler et al., 2018), or even assumptions that they have themselves (correctly) negated (Fiedler, Walther, Armbruster, Fay, & Naumann, 1996; Mayo, Schul, & Burnstein, 2004). Their estimates are biased toward fully irrelevant or impossible numerical anchors (Ariely, Loewenstein, & Prelec, 2003), and they fall prey to obviously false news despite undisputable counterevidence (Schwarz, 2018), the debunking of which is often incomplete (Chan, Jones, Hall Jamieson, & Albarracin, 2017).

Metacognitive myopia was particularly shown to cause illusions related to sample size (confusion of relative proportions and absolute sample size; Fiedler et al., 2019) and aggregation levels (Fiedler, Freytag, & Meiser, 2009). Melioration (Herrnstein & Prelec, 1991) may also fit into the category of metacognitive myopious behavior. We believe that participants' far-from-optimal performance in this speed-accuracy trade-off paradigm is the result of a metacognitive regulation failure that can be placed in the category of metacognitive myopia, making such a paradigm ideal for the experimental analysis of the moderators and limitations thereof.

Outlook on Adaptive Cognition, the Extant Literature, and Future Research

It should be clear that our metacognitive perspective does not predict a universal accuracy bias. Myopia can prevent people from critical quality control and correction of any bias, whether it favors accuracy, speed, or any other sort of bias. There is no reason to expect a constant bias across all speed-accuracy trade-offs for all task domains, nor do we propagate that performance in such trade-off tasks is generally maladaptive. There are good reasons, however, to believe that the choices based on sequentially experienced samples (Hertwig & Pleskac, 2010) will exhibit a similar accuracy bias and speed neglect as in the present experiments. In this regard, the distinct trade-off task constructed for the present purpose may turn out to be of interest for researchers working in this area.

There can be no doubt that faster strategies can be induced through extraneous interventions or incentives. Phillips, Hertwig, Kareev, and Avrahami (2014), for example, induced speed through social competition. When agents competed for the right to play a gamble, they managed to make quicker choices than without competition. However, one should not mistake such an externally enforced strategy change as evidence for better metacognitive understanding. The crucial question is whether speed enforced by social competition carries over to new decision tasks and helps to overcome metacognitive myopia on future tasks when the extraneous incentives are no longer at work.

The present findings appear to be at variance with previous results demonstrating close-to-optimal speed-accuracy trade-off performance on perceptual as well as conceptual tasks (Jarvstad et al., 2012). Similarly, in their discussion of why participants often base decisions on small samples (typically seven to nine observations per option), Hertwig and Pleskac (2010) stipulated that decision makers must have a good grasp of the relation between accuracy and sample size, and they must be very conscious of the time cost of longer sampling. Still other researchers observed a basic tendency to respond too quickly on a task that rendered payoff dependent on short versus long delay (Ashby & Gonzalez, 2017).

Yet, it is not too surprising that different strategic biases are obtained in different paradigms. Another strand of literature testified to an accuracy bias in speed-accuracy trade-offs based on perceptual decisions (Starns & Ratcliff, 2010, 2012; Evans & Brown, 2017; Bogacz et al., 2010). Evans, Bennett, and Brown (2019) made the excellent point that by varying the task structure and thereby the optimal strategy, the predominant accuracy bias can be turned into a more speed-oriented strategy bias.

From such a broader perspective, the present findings are not at all incompatible with other findings in the extant literature. Closer inspection reveals that no prior research has ever tackled speed-accuracy trade-offs in a sample-based choice paradigm, in which sample size offer a joint measure of speed and accuracy and the trade-off is explicitly built into the payoff scheme.

References

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

Alves, H., & Mata, A. (2019). The redundancy in cumulative information and how it biases impressions. *Journal of Personality and Social Psychology* 17(6), 1035–1060. https://doi.org/10.1037/pspa0000169

Ariely, D., Loewenstein, G., & Prelec, D. (2003). "Coherent Arbitrariness": Stable Demand Curves without Stable Preferences. *Quarterly Journal of Economics, 118*(1), 73. https://doi.org/10.1162/00335530360535153

Ashby, N. J., & Gonzalez, C. (2017). The Influence of Time Estimation and Time-Saving Preferences on Learning to Make Temporally Dependent Decisions from Experience. *Journal of Behavioral Decision Making*, *30*(4), 807-818. https://doi.org/10.1002/bdm.2006

Bell, D.E., Raiffa, H., & Tversky, A. (Eds.). (1988). *Decision making: Descriptive, normative and prescriptive interactions*. Cambridge University Press.

Beilock, S. L., Bertenthal, B. I., Hoerger, M., & Carr, T. H. (2008). When does haste make waste? Speed-accuracy trade-off, skill level, and the tools of the trade. *Journal of Experimental Psychology: Applied*, *14*(4), 340–352. https://doi.org/10.1037/a0012859

Bernoulli, J. (1713). *Ars conjectandi, opus posthumum* [The art of conjecturing, posthumous work]. Thurneysen Brothers.

Bock, O., Baetge, I., & Nicklisch, A. (2014). hroot: Hamburg Registration and Organization Online Tool. *European Economic Review, 71*, 117-120. https://doi.org/10.1016/j.euroecorev.2014.07.003

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*(4), 700–765. https://doi.org/10.1037/0033-295X.113.4.700

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

Chan, M. P. S., Jones, C. R., Hall Jamieson, K., & Albarracin, D. (2017). Debunking: A meta-analysis of the psychological efficacy of messages countering misinformation. *Psychological Science*, *28*(11), 1531-1546. https://doi.org/10.1177/0956797617714579

Dawes, R. M., Faust, D., & Meehl, P. E. (1989). Clinical versus actuarial judgment. *Science*, *243*(4899), 1668-1674. https://doi.org/10.1126/science.2648573

Dijksterhuis, A., & Nordgren, L. F. (2006). A theory of unconscious thought. *Perspectives on Psychological Science*, *1*(2), 95-109. https://doi.org/10.1111/j.1745-6916.2006.00007.x

Dunning, D. (2011). The Dunning-Kruger effect: On being ignorant of one's own ignorance. In J. M. Olson & M. P. Zanna (Eds.), *Advances in experimental social psychology, Vol 44.* (pp. 247–296). Academic Press. https://doi.org/10.1016/B978-0-12-385522-0.00005-6

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

Einhorn, H. J., & Hogarth, R. M. (1981). Behavioral decision theory: Processes of judgment and choice. *Annual Review of Psychology*, *32*, 53–88. https://doi.org/10.1146/annurev.ps.32.020181.000413

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

Erev, I., Ert, E., Plonsky, O., Cohen, D., & Cohen, O. (2017). From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological Review*, *124*(4), 369–409. https://doi.org/10.1037/rev0000062

Evans, N. J., Bennett, A. J., & Brown, S. D. (2019). Optimal or not; depends on the task. *Psychonomic Bulletin and Review, 26*(3), 1027-1034. https://doi.org/10.3758/s13423-018-1536-4

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

Fiedler, K. (2000). Beware of samples: A cognitive-ecological sampling approach to judgment bias. *Psychological Review*, *107*, 659– 676. https://doi.org/10.1037/0033-295X.107.4.659

Fiedler, K. (2008). The ultimate sampling dilemma in experience-based decision making. *Journal of Experimental Psychology: Learning, Memory & Cognition, 34*, 186-203. https://doi.org/10.1037/0278-7393.34.1.186

Fiedler, K. (2012). Meta-cognitive myopia and the dilemmas of inductive-statistical inference. *The Psychology of Learning and Motivation, 57*, 1-55. https://doi.org/10.1016/B978-0-12-394293-7.00001-7

Fiedler, K., Ackerman, R., & Scarampi, C. (2019). Metacognition: Monitoring and controlling one's own knowledge, reasoning and decisions. In R.J. Sternberg & J. Funke (Eds.), *The psychology of human thought: An introduction*. Heidelberg University Publishing.

Fiedler, K., Freytag, P., & Meiser, T. (2009). Pseudocontingencies: An integrative account of an intriguing cognitive illusion. *Psychological Review, 116*(1), 187. https://doi.org/10.1037/a0014480

Fiedler, K., Hofferbert, J., & Wöllert, F. (2018). Metacognitive myopia in hidden-profile tasks: The failure to control for repetition biases. Frontiers in Psychology, 9. https://doi.org/10.3389/fpsyg.2018.00903

Fiedler, K., & Kareev, Y. (2006). Does decision quality (always) increase with the size of information samples? Some vicissitudes in applying the law of large numbers. *Journal of Experimental Psychology: Learning, Memory and Cognition, 32*, 883-903. https://doi.org/doi:10.1037/0278-7393.32.4.883

Fiedler, K., & Kareev, Y. (2011). Clarifying the advantage of small samples: As it relates to statistical wisdom and Cahan's (2010) normative intuitions. *Journal of Experimental*

Psychology: Learning, Memory, and Cognition, *37*(4), 1039-1043. https://doi.org/10.1037/a0023259

Fiedler, K., Kareev, Y., Avrahami, J., Beier, S., Kutzner, F., & Hütter, M. (2016). Anomalies in the detection of change: When changes in sample size are mistaken for changes in proportions. *Memory & Cognition, 44*(1), 143-161. https://doi.org/10.3758/s13421-015-0537-z

Fiedler, K., Krüger, T., Koch, A., & Kutzner, F. (2020). Dyadic judgments based on conflicting samples: The failure to ignore invalid input. *Journal of Behavioral Decision Making*. https://doi.org/10.1002/bdm.2173

Fiedler, K., Renn, S., & Kareev, Y. (2010). Mood and judgments based on sequential sampling. *Journal of Behavioral Decision Making, 23*, 483-495.

Fiedler, K., Schott, M., Kareev, Y., Avrahami, J., Ackerman, R., Goldsmith, M., . . . Pantazi, M. (2020). Metacognitive myopia in change detection: A collective approach to overcome a persistent anomaly. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. https://doi.org/10.1037/xlm0000751

Fiedler, K., Walther, E., Armbruster, T., Fay, D., & Naumann, U. (1996). Do you really know what you have seen? Intrusion errors and presuppositions effects on constructive memory. *Journal of Experimental Social Psychology, 32*(5), 484-511. https://doi.org/10.1002/bdm.669

Fried, L. S., & Peterson, C. R. (1969). Information seeking: Optional versus fixed stopping. *Journal of Experimental Psychology*, *80*(3p1), 525-529. https://doi.org/10.1037/h0027484

Garland, H. (1990). Throwing good money after bad: The effect of sunk costs on the decision to escalate commitment to an ongoing project. *Journal of Applied Psychology*, *75*(6), 728-731. https://doi.org/10.1037/0021-9010.75.6.728

Ghisletta, P., Joly-Burra, E., Aichele, S., Lindenberger, U., & Schmiedek, F. (2018). Age differences in day-to-day speed-accuracy trade-offs: Results from the cogito study. *Multivariate Behavioral Research*. https://doi.org/10.1080/00273171.2018.1463194

Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, *1*(1), 107–143. https://doi.org/10.1111/j.1756-8765.2008.01006.x

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, *62*, 451–482. https://doi.org/10.1146/annurev-psych-120709-145346

Herrnstein, R. J., & Prelec, D. (1991). Melioration: A theory of distributed choice. *Journal of Economic Perspectives*, *5*(3), 137-156. https://doi.org/10.1257/jep.5.3.137

Hertwig, R., & Pleskac, T. J. (2008). The game of life: How small samples render choice simpler. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind: Prospects for Bayesian cognitive science* (pp. 209–235). Oxford, England: Oxford University Press.

Hertwig, R., & Pleskac, T. J. (2010). Decisions from experience: Why small samples? *Cognition*, *115*(2), 225–237. https://doi.org/10.1016/j.cognition.2009.12.009

Hsee, C. (1996). The evaluability hypothesis: An explanation for preference reversals between joint and separate evaluations of alternatives. *Organizational Behavior and Human Decision Processes*, *67*(3), 247–257. https://doi.org/10.1006/obhd.1996.0077

Hsee, C. K., & Zhang, J. (2010). General evaluability theory. *Perspectives on Psychological Science*, *5*(4), 343-355. https://doi.org/ 10.1177/1745691610374586

Hogarth, R. M., Mukherjee, K., & Soyer, E. (2013). Assessing the chances of success: Naïve statistics versus kind experience. *Journal of Experimental Psychology: Learning Memory and Cognition*. https://doi.org/10.1037/a0028522

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

Irwin, F. W., Smith, W. A. S., & Mayfield, J. F. (1956). Tests of two theories of decision in an "expanded judgment" situation. *Journal of Experimental Psychology, 51*(4), 261–268. https://doi.org/10.1037/h0041911

Jarvstad, A., Hahn, U., Warren, P. A., & Rushton, S. K. (2014). Are perceptuo-motor decisions really more optimal than cognitive decisions? *Cognition*, *130*(3), 397-416. https://doi.org/10.1016/j.cognition.2013.09.009

Jarvstad, A., Rushton, S. K., Warren, P. A., & Hahn, U. (2012). Knowing when to move on: Cognitive and perceptual decisions in time. *Psychological Science*, *23*(6), 589-597. https://doi.org/10.1177/0956797611426579

Juslin, P., Nilsson, H., & Winman, A. (2009). Probability theory, not the very guide of life. *Psychological Review*. https://doi.org/10.1037/a0016979

Juslin, P., Winman, A., & Hansson, P. (2007). The naïve intuitive statistician: A naïve sampling model of intuitive confidence intervals. *Psychological Review*, *114*(3), 678–703. https://doi.org/10.1037/0033-295X.114.3.678

Kahneman, D., & Tversky, A. (2000). Choices, values, and frames. In D. Kahneman & A. Tversky (Eds.) *Choices, values, and frames* (pp.1-16). Cambridge University Press. https://doi.org/10.1017/CBO9780511803475

Kareev Y. (2000). Seven (indeed, plus or minus two) and the detection of correlations. *Psychological Review, 107*(2), 397-402. https://doi.org/10.1037/0033-295X.107.2.397.

Kareev, Y., Arnon, S., & Horwitz-Zeliger, R. (2002). On the misperception of variability. *Journal of Experimental Psychology: General*, *131*(2), 287–297. https://doi.org/10.1037/0096-3445.131.2.287

Khodadadi, A., Fakhari, P., & Busemeyer, J. R. (2014). Learning to maximize reward rate: a model based on semi-Markov decision processes. *Frontiers in Neuroscience, 8*, 101.

Koriat, A. (2012). The self-consistency model of subjective confidence. *Psychological Review*, *119*(1), 80–113. https://doi.org/10.1037/a0025648

Larrick, R. P., & Soll, J. B. (2008). The MPG illusion. *Science*, *320*(5883), 1593–1594. https://doi.org/10.1126/science.1154983

Mayo, R., Schul, Y., & Burnstein, E. (2004). "I am not guilty" vs "I am innocent": Successful negation may depend on the schema used for its encoding. *Journal of Experimental Social Psychology*, *40*(4), 433-449. https://doi.org/10.1016/j.jesp.2003.07.008

McCaughey, L., Prager, J., & Fiedler, K. (2020). *Adapting to information costs and payoff changes in sampling-based decisions*. Unpublished research, Heidelberg University.

Pachella, R. G., & Pew, R. W. (1968). Speed-Accuracy Trade-off in Reaction Time: Effect of Discrete Criterion Times. *Journal of Experimental Psychology*, *76*(1, Pt.1), 19-24. https://doi.org/10.1037/h0021275

Payne, J. W., Bettman, J. R., & Luce, M. F. (1996). When time is money: Decision behavior under opportunity-cost time pressure. *Organizational Behavior and Human Decision Processes*, *66*(2), 131–152. https://doi.org/10.1006/obhd.1996.0044

Peer, E., & Gamliel, E. (2013). Pace yourself: Improving time-saving judgments when increasing activity speed. *Judgment and Decision Making*, *8*(2), 106–115. https://doi.org/10.2139/ssrn.2178228

Peterson, C. R., & Beach, L. R. (1967). Man as an Intuitive Statistician. *Psychological Bulletin*, *68*(1), 29–46. https://doi.org/10.1037/h0024722

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

Prager, J., Krueger, J. I., & Fiedler, K. (2018). Towards a deeper understanding of impression formation—New insights gained from a cognitive-ecological perspective. *Journal of Personality and Social Psychology*, *115*(3), 379–397. https://doi.org/10.1037/pspa0000123

Price, P. C., Kimura, N. M., Smith, A. R., & Marshall, L. D. (2014). Sample size bias in judgments of perceptual averages. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*(5), 1321–1331. https://doi.org/10.1037/a0036576

Quiroga, M. Á., Hernández, J. M., Rubio, V., Shih, P. C., Santacreu, J., & Howard, V. N. (2007). Influence of impulsivity-reflexivity when testing dynamic spatial ability: Sex and g differences. *The Spanish Journal of Psychology*, *10*(2), 294-302. https://doi.org/10.1017/S1138741600006569

Rapoport, A., & Tversky, A. (1970). Choice behavior in an optional stopping task. *Organizational Behavior and Human Performance*, *5*(2), 105-120. https://doi.org/10.1016/0030-5073(70)90008-5

Ratcliff, R., Smith, P. L., & McKoon, G. (2015). Modeling regularities in response time and accuracy data with the diffusion model. *Current Directions in Psychological Science*, *24*(6), 458–470. https://doi.org/10.1177/0963721415596228

Schwarz, N. (2018). Of fluency, beauty, and truth: Inferences from metacognitive experiences. In J. Proust & M. Fortier (Eds.), *Metacognitive diversity: An interdisciplinary approach.* (pp. 25–46). Oxford University Press.

Simen, P., Vlasov, K., & Papadakis, S. (2016). Scale (in)variance in a unified diffusion model of decision making and timing. *Psychological Review*, *123*(2), 151–181. https://doi.org/10.1037/rev0000014.supp (Supplemental)

Simon, H. A. (1955). A Behavioral Model of Rational Choice. *Quarterly Journal of Economics, 69*(1), 99–118. https://doi.org/10.2307/1884852

Spiller, S. A. (2019). Opportunity cost neglect and consideration in the domain of time. *Current Opinion in Psychology*, *26*, 98–102. https://doi.org/10.1016/j.copsyc.2018.10.001

Starns, J. J., & Ratcliff, R. (2010). The effects of aging on the speed–accuracy compromise: Boundary optimality in the diffusion model. *Psychology and Aging*, *25*(2), 377–390. https://doi.org/10.1037/a0018022

Starns, J. J., & Ratcliff, R. (2012). Age-related differences in diffusion model boundary optimality with both trial-limited and time-limited tasks. *Psychonomic bulletin & review*, *19*(1), 139-145. https://doi.org/10.3758/s13423-011-0189-3

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, *53*(1), 1–26. https://doi.org/10.1016/j.cogpsych.2005.10.003

Surowiecki, J. (2004). The Wisdom of Crowds, 2004. *New York: Anchor*.Svenson, O., & Eriksson, G. (2017). Mental models of driving and speed: biases, choices and reality. *Transport Reviews*, *37*(5), 653–666. https://doi.org/10.1080/01441647.2017.1289278

Swets, J. A., Dawes, R. M., & Monahan, J. (2000). Psychological science can improve diagnostic decisions. *Psychological Science in the Public Interest*, *1*(1), 1-26. https://doi.org/10.1111/1529-1006.001

Tversky, A., & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin*, *76*(2), 105-110. https://doi.org/10.1037/h0031322

Unkelbach, C., Fiedler, K., & Freytag, P. (2007). Information repetition in evaluative judgments: Easy to monitor, hard to control. *Organizational Behavior and Human Decision Processes, 103*, 37-52. https://doi.org/10.1016/j.obhdp.2006.12.002

Winman, A., & Juslin, P. (1993). Calibration of sensory and cognitive judgments: Two different accounts. *Scandinavian Journal of Psychology*, *34*(2), 135–148. https://doi.org/10.1111/j.1467-9450.1993.tb01109.x

APPENDIX 2: Second paper

Rivals reloaded: Adapting to sample-based speed-accuracy trade-offs through competitive pressure

Linda McCaughey, Johannes Prager, & Klaus Fiedler

Heidelberg University

Collecting an adequate amount of information for a decision is an important skill. However, previous experiments on speed-accuracy trade-offs in sample-based decisions revealed a marked oversampling bias. Most participants failed to realize that speed (number of decisions completed in a restricted period) determined total payoff more than accuracy (average correctness of choices). Various interventions aimed at making participants recognize the speed advantage failed to alleviate participants' oversampling. Yet, evidence by Philipps et al. (2014), points to social competition as a promising remedy. When faced with the threat of being pre-empted by a rival in making decisions, participants seem to reduce information search substantially. In the present research (N=101), participants had to compete with a fast (computer-simulated) rival and indeed reduced self-determined sample size compared to a control condition. However, this speed increase did not carry over to subsequent decisions sufficiently to count as genuinely adaptive. Moreover, mere exposure to a teammate using small samples either in the competitive or the standard solitary version of the task led to a similar, small reductions in sample size, suggesting that social competition does not necessarily underlie the rival-induced speed increase. Further research is needed to fully uncover the metacognitive underpinnings of the oversampling's resistance to change.

Keywords: speed-accuracy trade-off; metacognitive regulation; monitoring; sample-based decisions; competition

Adaptive decision makers need to respond suitably to their environment, its informational structures, and changes thereof (Payne, Bettman, & Johnson, 1993; Todd & Gigerenzer, 2007). Specifically, while engaging in search for information about options in order to decide between them, they need to avoid making too hasty decisions based on too little or undiagnostic information, but also take care not to overdo it and collect unnecessary additional information. In short, they need to weigh the benefits against the costs of prolonged information search. Extended information search about options generally comes with an advantage of a higher probability of choosing the best option. Yet, that advantage needs to be weighed against the time, money, and effort that must be invested to gain additional information – as well as the possibility of losing out on options altogether. For example, when booking a hotel room for a conference, one should collect enough information to make a satisfactory choice, but one should definitely avoid searching for further information about the many different options for so long that all hotels close to the conference location are fully booked.

Oversampling in speed-accuracy trade-offs. Previous research has shown that participants are to some degree sensitive to changes in both quality and financial cost of additional

---

The work underlying the present article was supported by a grant provided by the Deutsche Forschungsgemeinschaft to the third author (FI 294/30-1). Email correspondence can be addressed to linda.mccaughey@psychologie.uni-heidelberg.de. The experimental data can be found at https://heibox.uni-heidelberg.de/d/78070310d91848e8a274/

information (Pitz, 1968). They also engage in longer information search to increase decision accuracy when decision stakes are higher (Hau et al., 2008) and trade off the amount of information considered against the effort required to process (Payne et al., 1993). However, when it comes to adaptively regulating the invested time to attain information about choice options in a trade-off task, our own recent results give less cause for optimism about humans' adaptivity (Fiedler et al., 2021).

In a sample-based speed-accuracy trade-off task, our participants had to make as many correct choices as possible during a limited period of time. We asked them to make choices between two investment funds based on random samples of the funds' previous successes and failures, represented by upward and downward arrows respectively. On each trial, the arrows for the two funds appeared one by one, meaning that the sample grew over time. Hence, the longer participants waited, the more information they received to base their decision on, making it more likely that those decisions would be correct. Correct decisions consisted of identifying the better investment fund and were rewarded (with a gain of 1 point in an experimental currency unit), while wrong decisions were punished (with a loss of 1 point). However, since participants had to make as many choices as possible to maximize the opportunities for a gain within the time limit, they could not wait indefinitely. To maximize their overall payoff converted from the points in the experimental currency, they had to find a suitable balance between their decision accuracy – how often their decisions were correct and rewarded – and their speed – the number of decisions they attained – by suitably regulating the amount of information they waited for in each trial.

It is important to note, however, that increasing speed (while sacrificing accuracy) and increasing accuracy (while sacrificing speed) do not impact the overall payoff to the same extent. Increasing speed by making do with smaller samples only leads to a slight loss of accuracy in most cases, but allows for many more choices to be completed, thereby increasing the opportunities for gains immensely. Improving accuracy even a little, on the other hand, already implies a major sacrifice of speed in most cases. Such forfeiting of decisions, and hence gain opportunities, cannot be compensated by the negligible increase in accuracy, leading to a remarkable loss of overall payoff when accuracy is emphasized. Thus, maximizing the total payoff in this speed-accuracy trade-off task clearly calls for an emphasis on speed rather than accuracy.

However, participants confronted with this task have demonstrated remarkable persistence in their emphasis on accuracy instead. They used sample sizes multiple times larger than they should have, thereby foregoing substantial proportions of the possible payoff. Various interventions – from trial-by-trial feedback that allowed for performance comparisons of smaller and larger samples, to different upper sample size limits that led to choices based on more diverse sample sizes, to conditions in which they could not lose from incorrect choices – failed to elicit notably faster strategies. Participants seemed unable or unwilling to exert the necessary metacognitive monitoring and control to assess different speeds and their consequences for the payoff. This left them unable to adapt away from oversampling to exploit the advantage of speed over accuracy in this sample-based speed-accuracy trade-off. They did not seem to be able to infer from their experiences that faster strategies were more profitable.

Social competition. In contrast to the conspicuous failure of all attempts to induce small-sample strategies across four experiments reported by Fiedler et al. (2021), research by Phillips et al. (2014) points to social competition as an effective way to induce faster decisions based on less information sampling. Phillips et al. (2014) found that participants who had to decide between two lotteries and could previously acquire information about the outcomes through sampling, decreased their information search drastically when a rival participant was introduced. The competing participants could pre-empt each other's choices by choosing first.[1] However, choosing first required truncating information search at an early stage. When faced with the danger of not being able to choose for themselves because the rival participant might be faster, participants make do with comparably small samples. Phillips et al.'s (2014) simulations revealed an asymmetry between the costs of oversampling and undersampling: While sampling too much bore the risk of the opponent getting to choose, sampling too little only decreased the probability of choosing the better lottery by a little. Notably, participants behaved in line with this by most frequently drawing only as little as the first mandatory observation.

Social competition as a potential remedy to oversampling

Thus, the social competition idea implemented by Phillips et al. (2014) promises a patent recipe to elicit an emphasis on speed by inducing smaller samples. We set out to transfer this idea to our paradigm to assess whether participants' performance at the speed-accuracy trade-off could be improved when a competitive element was introduced. However, closer reflection reveals that any potentially beneficial impact of social competition on adaptive regulation of speed-accuracy trade-offs must be analyzed with respect to two issues: First, what might appear to be adaptive may only be *pseudo-adaptive* unless the trade-off improvement carries over to a new situation without competition pressure present to necessitate small samples.

Second, the social competition intervention is a complex intervention involving an interplay of different features. The competitive aspect is the most prominent, but the active engagement of participants in the competition, and most fundamentally, the exposure to smaller samples may likewise be causally relevant. Any impact of the intervention might only be due to a subset of those aspects. Specifically, participants may not need to be actively involved in the competition. Passive observation of a competitive situation might have the same impact. Going even further, the competitive aspect itself might not even be crucial. It might suffice to observe choices based on small samples, which could underlie the potential effect of more complex interventions.

Pseudo-adaptive versus genuinely adaptive behavior. Let us discuss both issues in turn, starting with the need to distinguish between adaptivity and what we want to call *pseudo-adaptivity.* Pseudo-adaptive behavior is behavior that corresponds – in our case – to an efficient strategy, and so, may appear adaptive at first glance. To count as genuinely adaptive, however, the efficient behavior needs to be exhibited in response to the relevant features of the task and not just correspond to it haphazardly. Therefore, participants merely reducing sample size to

[1] In Phillips et al.'s (2014) paradigm the consequence of pre-emption by the rival was that the participant received the remaining option. Our implementation of competition differed in that pre-emption through the rival resulted in receiving no outcome.

avoid being unable to make any choices at all cannot count as evidence for genuinely adaptive behavior. To identify genuine adaptivity and demarcate it from pseudo-adaptivity, it is necessary to demonstrate that participants' use of smaller samples carries over to subsequent choices in a new block, when competition with a fast rival is no longer required. To ensure we could assess this transfer performance, our design did not only include a first block with a competing rival, but also a second block in which all participants had to engage in the standard solitary speed-accuracy trade-off task.

Only a speedy small-sample strategy that carries over from Block 1 to Block 2 would provide cogent evidence for genuinely adaptive behavior. If there is no transfer effect to Block 2, then the potential speed-up in Block 1 would solely be a response to the rival's presence and not also a response to the overall task structure, which is what makes drawing small samples the most efficient approach. Therefore, it would be pseudo-adaptive.[2]

Decomposing social competition effects on speed and accuracy. While the comparison of Block 1 and 2 is intended to elucidate the distinction between pseudo-adaptive and adaptive small-sample usage in a trade-off task, we also included a between-participants factor to clarify another distinction. We wanted to investigate whether adding a rival to the task actually constituted a social competition intervention, that is, whether the active involvement in social competition was the main causal factor in effecting a potential sample-size reduction. This may plausibly be the case because social competition can increase attention to a task (DiMenichi & Tricomi, 2015), lead to more effort being expended (Van de Pol, et al. 2012), and may increase risk taking (Mishra et al., 2014). These implications of active engagement in social competition might promote insights about the merit of a speedy strategies for the task the potential result of which will be assessed in the Actor-Rival Condition in our experiment.

However, we assume that the most relevant aspect of the intervention is that it allows participants to experience smaller samples, which they would not encounter otherwise because they do not tend to employ them spontaneously (Fiedler et al., 2021). This allows them to monitor the success of a strategy they would not have been able to monitor otherwise, namely a fast strategy based on small samples. This might lead to an improved inference about the relative effectiveness of strategies based on different sample sizes.

This implies that the rival intervention may, at its core, not be a social competition intervention. Specifically, it might not be necessary to be involved in competitive pressure oneself for an intervention to unfold the same effect. It might be enough to observe an agent competing against a rival and therefore employing fairly small samples. We included an additional condition, the Observer-Rival Condition, offering just that. Instead of competing against the virtual rival themselves, participants observed a virtual agent who was declared their teammate compete against the virtual rival. This would allow us to assess whether the active component of the main intervention had a crucial influence on participants' sample sizes.

Going one step further, it might not even be necessary to include the competitive aspect at all – whether as active involvement or passive observation. Perhaps simply observing another

[2] Frey (2020) draws a similar distinction in his lottery-based decision paradigm including competitive pressure: the 'optimistic view' is that competitive pressure actually leads adaptive behavior, while the "pessimistic view" assumes that competitive pressure might induce a certain behavior for other reasons.

agent pursue a faster strategy based on smaller samples suffices. To test this, we added a third intervention condition, the Observer-No-Rival Condition, in which participants observed a virtual agent, likewise declared their teammate, play the standard solitary trade-off task using a strategy based on small samples. A control condition in which participants actively played the standard solitary speed-accuracy trade-off task, the Actor-No-Rival Condition, completed the design for the first block. All four conditions engaged in the standard solitary task in Block 2 to allow for the transfer test of all interventions and compare them against the baseline of the control group who had not undergone any intervention.

Hypotheses. Returning to Block 1, we expect participants playing the standard solitary task (Active-No-Rival Condition) to exhibit sample sizes as large as in previous experiments, replicating the persistent oversampling. In contrast, we expect participants who actively face the rival (Actor-Rival Condition) to use substantially smaller sample sizes than those control participants (Actor-No-Rival Condition), thereby successfully replicating the competition effect found in Phillips et al. (2014) in our speed-accuracy trade-off paradigm. The two observation conditions were passive during the first block, so there is no participant-generated data to be analyzed for those two conditions.

For Block 2, in which all conditions engaged in the standard (Actor-No-Rival) task, a transfer effect for the Actor-Rival Condition would manifest itself in smaller sample sizes compared to the Actor-No-Rival Condition. A full transfer would also mean that those sample sizes are similarly small as in Block 1. An effect unique to the active-competitive sample-based speed-accuracy trade-off task sample-based speed-accuracy trade-off task intervention would be demonstrated by a decreased Block-2 sample sizes for participants who previously actively competed with a rival (Actor-Rival Condition), but no sample size decrease caused by the other interventions in the observation conditions (Observation-Rival and Observation-No-Rival). The smaller samples sizes would imply lower accuracy, more trials, and higher payoff in the Actor-Rival Condition than in the Observer-Rival and the Observer-No-Rival Condition. Alternatively, if the active-competitive intervention's effect did not rely on active engagement or social competition to lead to a speedier strategy, one or both interventions in the observation conditions may be just as effective, resulting in similarly decreased sample size and related measures compared to the control condition.

Methods

Participants and design. We recruited 104 participants from our local participant pool. They were aged 24.51, on average (*SD* = 6.95), with a range of 18 to 64 years. The vast majority were university students (92.4 %), of which 22.86 % indicated that they studied psychology, making it the most common field of study. The participants were randomly assigned to one of four conditions. The Actor-No-Rival (A-NoR) group (*N* = 23) played the standard solitary speed-accuracy trade-off task as in Fiedler et al. (2021, Exp. 3 & 4) without any additional intervention. The Actor-Rival (A-R) Condition (*N* = 27) was modelled after Phillips et al. (2014): Participants faced a computer-controlled virtual rival who they had to compete with in the speed-accuracy trade-off task during the first block. The remaining two conditions were observer conditions, in which participants passively watched a computer-controlled virtual teammate play the game in the first block. The Observer-No-Rival (O-NoR) Condition (*N* =

26) consisted only of the virtual teammate playing the standard speed-accuracy trade-off task individually, whereas the Observer-Rival (O-R) Condition ($N = 27$) involved the presence of a virtual rival whom the virtual teammate had to play against, meaning that participants watched as two computer-controlled agents competed at the task, one representing the teammate and the other the rival.

In the second block, all conditions played the standard solitary speed-accuracy trade-off task of the A-NoR Condition. The full design is visualized in Figure 1. The experiment was conducted between two parts of an unrelated experiment on eye-witness memory, which was followed by a short unrelated questionnaire about compliance. The experiment was approved by the local department ethics committee and was not preregistered.

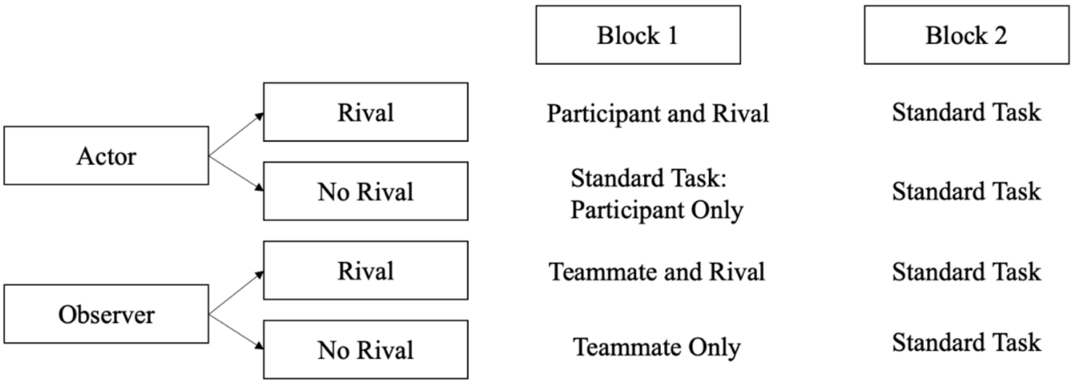

Figure 1. Design illustration with factors Actor versus Observer x Rival versus No Rival and Block 1 and Block 2

Procedure. The experiment took about 25 minutes and consisted of two 10-minute blocks and a short post-experimental questionnaire. The different conditions' interventions were implemented during the first block, while in the second block, all participants worked on the standard speed-accuracy trade-off task that was also implemented in the Actor-No-Rival Condition (used in Fiedler et al., 2021). The basic task was to choose the better-performing of two investment funds based on gradually appearing binary observations of the rate development on individual trading days (i.e., increase vs. decrease). Increases were represented by green, upward-pointing arrows and decreases by magenta, downward-pointing arrows (see Figure 2). Arrows appeared at a rate of 2 per second (500 ms per arrow), and appeared alternatingly for Investment Fund A on the left side of the screen, and Investment Fund B on the right side of the screen.

For each round of the game, the experimental software was programmed to set success rates for the two funds (i.e., the probability p_A of a green arrow appearing for Fund A and the probability p_B of a green arrow appearing for Fund B). The difference between the success rates of the funds varied across trials. It was randomly set to one of three levels, $\Delta p = |p_A - p_B| = \{.1, .2, .4\}$. Under the constraints of the selected Δp parameters, the specific p_A and p_B values were drawn randomly from $\{.1, .2, .3, .4, .5, .6, .7, .8, .9\}$. The samples displayed for the two investment funds were drawn at random from binomial distribution with these probabilities. The better of two funds was defined as having the higher of two probabilities, or success rates, independently of the actually displayed sample information. The participants' task was to decide which fund had the higher success rate. Correct decisions were rewarded by the addition of one point (measured in ECU, i.e. experimental currency unit) to the participants' balance. Incorrect decisions were punished by the subtraction of one point from the balance.

Each round of the experimental task consisted of two stages: the sampling stage (illustrated in Figure 2), during which arrows were displayed, and the decision and feedback page, displayed on a subsequent slide by the program. During both stages, the current point balance was displayed at the top center of the screen. A horizontal bar to its right indicated the proportion of time left by progressively filling as time passed (see Figure 2). Arrows were presented one by one at a pace of 500 ms, alternatingly for Investment Fund A and B. The starting side, and hence, the fund for which the first arrow appeared, was drawn at random. The arrows appeared at randomly drawn positions within the side of the screen dedicated to the respective fund. Participants had to press the space bar to initiate the sequential presentation of arrows. As soon as they released the space bar, they were directed to the next page for the decision and feedback stage.

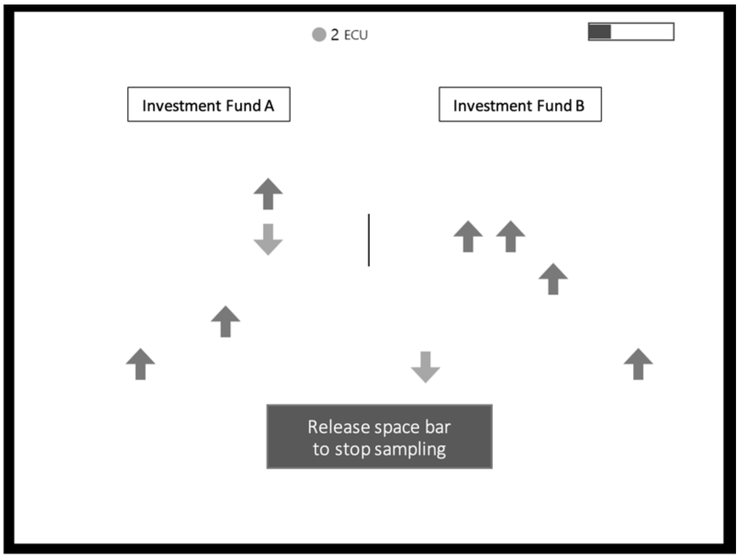

Figure 2. A screenshot of the participants' view of the sampling phase in the actor conditions (the two observer conditions did not see the instruction to release the space bar to stop sampling, as the software determined the stopping point). In the depicted situation, the space bar was pressed to initiate sampling and arrows are appearing alternatingly for Investment Fund A and B as long as the space bar remains pressed.

On the decision page, participants had to select one of two buttons to indicate their decision which was the better fund. The buttons could be selected by pressing "F" for the left option or "J" for the right option on the keyboard. The selection of the button was visualized on screen (by changing color tone and shadow of the button) and feedback on the correctness of the response was provided ("correct" vs. "false" appeared below the option buttons). Simultaneously, the balance was updated by adding +1 or –1 point (ECU) for a correct and incorrect response, respectively. The feedback and the highlighted new balance were presented for 1300 milliseconds before the next trial automatically commenced.

This was the standard version of the task that the Actor-No-Rival Condition played in both blocks and all conditions played in Block 2. The tasks faced by the other three experimental conditions in Block 1 were modifications of this standard task set-up. The Observer-No-Rival (O-NoR) Condition was kept almost identical, the only difference being that participants were mere observers and thus unable to control the sampling initiation or to make the choice. Rather, they watched their virtual teammate (represented by a picture of the robot character 'R2-D2' from the Star Wars franchise) play the game in a video-like sequence.[3] Arrow frequency and general procedure were kept parallel to the standard task, with decisions likewise visualized through a change of the selected options' on-screen button color and shadow and with feedback provided the same way (see above). Participants were informed that they would receive the teammate's payoff after the experiment. The virtual teammate was programmed to follow what we call a *difference strategy*, meaning that for each pair of arrows, it counted the successes for each option (represented by green, upward-pointing arrows). Whenever the absolute difference of the options' successes (i.e., the number of green arrows for one option minus the number of green arrows for the other option) reached a certain threshold value sampling was stopped and a decision made for the option favored by the sampled information. This threshold value varied between trials; it was drawn at random from the set of 1, 2, and 3 differences, which had the probabilities of .4, .4, and .2 to be drawn, respectively. Whenever the sample was still inconclusive (i.e., had not exceeded the threshold value) after 12 arrow pairs (12 arrows per option, 24 in total), the sampling was stopped and a random decision was made.

In the Actor-Rival (A-R) Condition, participants were active themselves and played a competitive variant of the speed-accuracy trade-off task, which featured a simulated computer-controlled rival, who was represented by a picture of the robot character BB-8 from the Star Wars franchise. The rival was controlled by an algorithm very similar to the teammate's algorithm described above, and thus followed a similar difference strategy except that the probabilities of the difference thresholds were shifted towards a slightly slower response (i.e., stopping thresholds of 1, 2, and 3 differences drawn at random with probabilities of .25, .25, and .5, respectively). The participants and their virtual rival received the same sample information in a joint sampling phase. And while sampling continued to be initiated by participants' pressing the space bar, the participant was not the only actor able to end the sampling phase. The virtual rival could also end the sampling phase whenever the algorithm had determined the rival to stop based on the shared sample before the participant released the space bar. If that was the case, the participant "lost" the decision to the rival and watched the rival make the choice (by animating the choice button as in the O-NoR Condition) and receive feedback on the decision page. They did not, however, see the rival's balance at any stage of the task. The participants' balance was, of course, not updated in response to the rival's decision, making the trials that were decided by the rival a plain loss of time for the participants.

[3] Both R2-D2, who represented the teammate, and BB-8, who represented the rival and will be introduced below, are important friendly robot characters in the Star Wars franchise. They communicate through electronic beeping noises and are therefore clearly non-humanoid but intelligent entities. There are also easy to distinguish because R2-D2 is white-blue and angular, while BB-8 is white-orange and round. Therefore, they were ideally suited as representations.

In the Observer-Rival (O-R) Condition, the participants assumed an observer role, while their virtual teammate (Star Wars robot character R2D2) played against the virtual rival (Star War robot character BB-8). The decision algorithms for the rival was almost identical to the rival in the A-R Condition, with the difference threshold drawn at random from the set of 1, 2, and 3 differences, with slightly different probabilities of .25, .2, and .5 respectively. The teammate followed precisely the same algorithm as in the A-NoR Condition (thresholds values of 1, 2, and 3 differences randomly drawn with probabilities of .4, .4, and .2). Participants observed the sampling phase, which was terminated by the first algorithm to reach its threshold. The participant was explicitly told whether the rival or the teammate ended the sampling and made the subsequent choice, which was again visualized through an animated button press and for which feedback was visible. The participants' balance was only updated when their teammate made the decision. The second block consisted of another 10-minute interval with participants of all conditions working on the standard solitary task of the A-NoR Condition, without any teammates or rivals.

The experiment concluded with a brief post-experimental questionnaire assessing participants' understanding and perception of their strategy. They were asked to rate the efficiency of different strategies, Strategy A and B, which were described as follows: "Strategy A: slow, correct decisions through large samples BUT overall less decisions and hence less gain opportunities", and "Strategy B: fast, uncertain decisions through small samples BUT all in all more decisions and hence gain opportunities". They also rated the "middle ground between Strategy A and B". They evaluated the efficiency of the three strategies separately on sliding scales with extremes labelled 'very inefficient' and 'very efficient'. They went on to indicate how big a sample a person following Strategy A and Strategy B would choose, respectively, and gave their answer by typing a number into a field., They also estimated their teammate's or their own correctness for Block 1 (depending on whether they were in an actor or observer condition) as well as their own correctness for Block 2 ("How often do you think did you/your teammate choose the better investment fund?") by typing a percentage into a field. Finally, we asked them to provide separate correctness estimate for small samples ("If you think back only to decisions that you made based on a sample of less than 8 days, i.e. less than 8 pairs of arrows, how often do you think did you choose the better investment fund in those cases?"). At the end of the session, participants were thanked and received their remuneration, which was contingent on their performance and amounted to about 5€, on average. The experimental data can be accessed at [Zenodo link will be inserted here, currently the data are available at https://heibox.uni-heidelberg.de/d/78070310d91848e8a274/]. The analysis code and the experimental materials will not be publicly available with the exception of the survey questions described above.

Results

Rationale. Recall that although we expected to replicate the persistent oversampling obtained by Fiedler et al. (2021) in the standard A-NoR Condition, we also intended to replicate Phillips et al.'s (2014) finding that a social-competition intervention in the A-R Condition affords an effective way to induce faster choice strategies based on smaller samples. Moreover, our sequential speed-accuracy trade-off task was ideally suited to test whether the competition of a rival leads to adaptive behavior, as manifested in lasting improvements of performance (even when the rival is no longer present in Block 2). Different variants of Phillips et al.'s intervention also enabled us to test the assumption that social competition is the crucial causal factor or whether variants of the intervention that merely relied on observation without involving participants in direct competition with a rival, are just as effective. Thus, our design allowed us to assess the extent to which the active-competitive intervention constitutes a lasting intervention against oversampling, on the one hand, and to which extent the intervention's effect is caused by active social competition proper, on the other hand.

Data Analysis. Participants' regulation and positioning within the speed-accuracy trade-off is captured by three indicators: The sample size *n* (defined as the number of observations for both options) for each round is the central measure, since that is the basic aspect via which participants can regulate their strategy and performance. The regulation of *n* determines the number of completed choices *k* (measuring speed) and the mean accuracy *a*, calculated as the difference between correct and incorrect choices divided by the total number of choices.[4] Speed and accuracy determine the overall efficiency, which is a multiplicative function of the number of completed choices and the mean accuracy, $Eff = k \cdot a$, and was converted into participants' actual payoff. Assuming certain common task constraints (symmetric payoff for correct and incorrect decision; stable information presentation rate), *k* decreases substantially faster than *a* increases with increasing *n*. Hence, for almost all mean sample sizes choosing faster (based on smaller *n*) is the key to a considerably higher total payoff.[5] Therefore, before we turn to the crucial question of whether profitable small-sample strategies induced in Block 1 carried over to self-determined sampling, let us first see whether the pronounced oversampling demonstrated in Fiedler et al. (2021) could be replicated in the standard A-NoR Condition.

Block 1: Replicating oversampling in the standard A-NoR Condition. With an average sample size of $n = 21.31$ ($SD = 11.40$), which led to an average accuracy of $a = .627$ ($SD = 0.158$), participants' sample sizes were similarly large as those in the original experiments and far from the optimal strategy of an average of $n = 5$ (see Footnote 4). They completed an average of $k = 53.26$ ($SD = 22.04$) choices in the given time period of 10 minutes for Block 1. The resulting mean payoff amounted to $Eff = 31.70$ ($SD = 9.53$).

The correlations between *n*, *a*, *k*, and *Eff* illustrate the task-inherent differential impact of speed and accuracy on the total payoff and how participants traded off speed and accuracy. To illustrate the relation between *n* and accuracy, consider the leftmost scatterplot in Figure 3 (filled symbols only). As individual participants' average sample size *n* increased, the average choice accuracy *a* also increased, $r_{n,a} = .59$, $t(21) = 3.31$, $p = .003$, reflecting the positive, moderate impact of sample size on accuracy. At the same time, the number of completed choices *k* decreased with *n* at a clearly higher rate (middle scatterplot), $r_{n,k} = -.815$, $t(21) =$ -

[4] This definition reflects the task structure of correct decisions leading to a gain of 1 point and incorrect decisions to a loss of 1 point. Since the payoff was the sum of those points, this definition clearly illustrates that payoff or *Efficiency* is the product of accuracy and the number of completed decisions k. An accuracy of 0 thus means that 50% of trials were answered correctly and an accuracy of 1 that all trials were answered correctly.

[5] Obviously, once very small samples sizes are reached, further reduction brings such a substantial reduction in accuracy that it is not advisable. A very profitable strategy is deciding based on the first difference between the two samples (see Fiedler et al. 2021 for further details): As soon as one observation indicates that one option is better than the other, the decision should be made. This strategy resulted in a mean sample size of about $n = 5$ in experiments in Fiedler et al. (2021).

6.46, $p < .001$. To understand how k and a jointly mediate the impact of n on *Eff*, which is the product of total k and mean a, pictured in the right scatterplot of Figure 3, it is important to note that *Eff* is strongly related to the number of completed choice k $r_{\text{Eff,k}} = .732$, $t(21) = 4.93$, $p < .001$, but is virtually unrelated to accuracy a, $r_{\text{Eff,a}} = .018$, $t(21) = 0.83$, $p = .395$. Because the number of completed trials k decreased much more strongly than accuracy a increased with increasing sample size n (Figure 3, left and center), increasing n led to a rather steep decline in *Eff* (Figure 3, right). Individual participants' total payoff correlated negatively with their average samples size, $r_{\text{n,Eff}} = -.681$, $t(21) = -4.26$, $p < .001$. Thus, although participants who drew larger samples were slightly more accurate, they were dramatically slower and thereby reduced their total payoff. This pattern is a perfect replication of the results obtained in previous research by Fiedler et al. (2021).

Figure 3. Scatterplots relating M_a, k, *Eff* to n. Open symbols refer to the A-R Condition's data (extrapolated in the case of k and *Eff;* filled symbols refer to the A-NoR Condition's data.

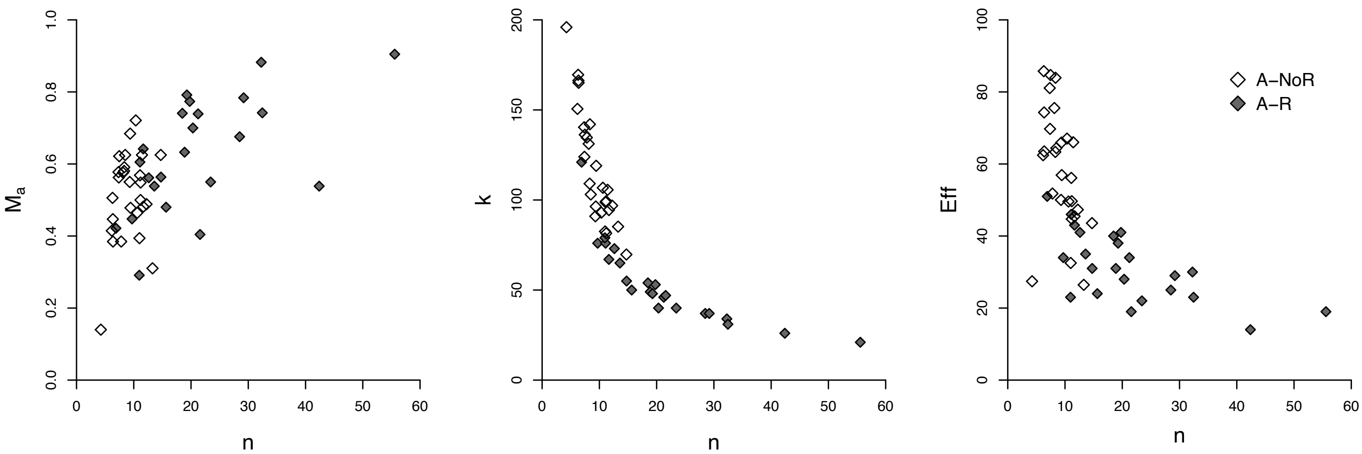

Block 1: Impact of the social competition intervention. A comparison of mean sample sizes in the experimental conditions A-NoR and A-R in Block 1 reveals that competing with a rival actually led our participants to reduce their average samples size. After all, this was their only chance to be faster than the rival in order to get the chance to make choices at all. Indeed, the reduction in sample size was quite substantial. When A-R participants were fast enough to make a choice (on 64.9% of all trials), they made do with a much lower mean sample size, $M_n = 9.24$ ($SD = 2.47$), than participants in the standard A-NoR Condition, $M_n = 21.31$ ($SD = 11.40$), $t(48) = -5.36$, $p < .001$, $d = 1.52$ (see Table 1). As a consequence of their smaller sample size, their average accuracy was also somewhat lower, $M_a = .513$ ($SD = 0.123$), compared to the standard A-NoR Condition, $M_a = .627$ ($SD = 0.158$), $t(48) = -2.85$, $p = .006$, $d = 0.81$.

Based on the A-R participants' mean n and a, we extrapolated how many choices and how high a payoff they could have hypothetically achieved under the same conditions as the A-NoR Condition, namely, without a rival who pre-empted participants' choices and therefore limited the number of choices they could make themselves.[6] These extrapolations show that while the

[6] This was done by taking participants' mean sample size (in the trials they did decide themselves) and extrapolating how many trials they could have completed in 10 minutes, which is possible because the arrows are presented at a specific frequency (2 arrows per second) and by taking into account the average decision time that needs to be added to the sampling

A-NoR Condition attained a mean of M_k = 53.26 choices (*SD* = 22.04) in the first block, the speed-up exhibited by the A-R Condition would have enabled them to complete more than twice as many choices under comparable conditions, M_k' =118.08 (*SD*= 31.84)[7]. Consequently, while the A-NoR Condition received an average payoff amounting to *Eff* = 31.35 points (*SD* = 9.57), the extrapolated payoff in the A-R Condition amounted to almost twice as much, *Eff* ' = 58.86 (*SD*=16.63). This shows that under the assumption of matched conditions the A-R participants would have outperformed the A-NoR Condition by far. This corroborates the regular finding reported by Fiedler et al. (2021): the product *Eff* = $k \cdot a$ clearly increases when sample size *n* decreases. Table 1 provides an overview of comparisons between the A-NoR and A-R Conditions for all performance measures.

Table 1

Mean n, a, k, and Eff for the A-NoR Condition as compared to mean n, a, extrapolated k' and extrapolated Eff' in the A-R Condition (standard deviations in parentheses)

Cond.	Mean *n* *(SD)*	Mean *a* *(SD)*	Mean actual *k (SD)*	Mean extrapol. *k' (SD)*	Mean actual *Eff (SD)*	Mean extrapol. *Eff' (SD*
A-NoR	21.31 (11.40)	.627 (0.158)	53.26 (22.04)	-	31.35 (9.57)	-
A-R	9.24 (2.47)	.513 (0.123)	86.74 (17.49	118.08 (31.84)	27.85 (10.02)	58.86 (16.63)

Individual scores on the same dependent measures (in the case of *k* and *Eff* the extrapolated scores) underlie the scatterplots in Figure 3. The lower sample size in the A-R Condition (open symbols) served to slightly reduce individual participants' accuracy *a* (left plot), but led to a clearly stronger increase in *k* (middle plot) and hence an overall gain in efficiency or overall payoff *Eff* (right plot). The shape of the three scatterplots highlights that *k*, rather than *a*, is the chief determinant of that *Eff*. The scatterplots also make clear that the relations between sample size *n* and *a*, *n* and *k*, and *n* and efficiency *Eff* are not fundamentally different for the two conditions. Rather, due to their smaller sample sizes, A-R participants are simply located much higher on *k*, but only slightly lower on *a*, leading to substantially higher efficiency for the A-R Condition.

In sum, the social competition intervention in the A-R Condition proved to be an effective intervention to increase speed as intended. However, the crucial question is whether employing smaller samples under competitive pressure would also lead participants to maintain their usage of small samples in the standard task, that is, whether their behavioral response would only be exhibited in reaction to the rival or carry over to a generally faster strategy on the speed-accuracy trade-off task based that persists after the rival is removed. Hence, what is even more telling than participants' behavior *while* they face the competitive pressure in Block 1, is how

time. The extrapolated number of trials can be multiplied by the actual accuracy to yield the hypothetical, extrapolated number of points.

[7] We refrained from conducting inferential tests on the extrapolated values of completed choice *k* and payoff *Eff*, as they are completely determined by *n* and *a*, which were already subjected to inferential tests.

they perform when the rival no longer threatens to intervene, in the standard task that all conditions were confronted with in Block 2.

Block 2: Analyses Overview. The Block-2 results serve to answer both of our research questions. First, to assess whether participants exhibited a lasting transfer effect from their previous actions, the A-R Condition's performance indicators – most importantly, sample size – need to be compared between Block 1 and Block 2. In addition, a regression analysis will serve to further examine the transfer effect by comparing Block 2 for the A-R Condition to the A-NoR Condition, which served as a control condition. At the same time, this regression analysis will assess whether the A-R intervention had a unique impact due to its active and competitive features/aspects by comparing the A-R Block-2 performance to the Observer-Rival Condition and the Observer-No-Rival Condition.

Block 2: Transfer from Block 1. The most important measure to answer the question of whether participants exhibited a lasting transfer is sample size *n*, since that is the variable they influence most directly (Figure 4 displays mean sample size split by condition and block). We will first compare sample size between Block 1 and 2 for the two conditions that provide participant-determined data for both blocks: A-R and A-NoR. For the A-R Condition, there is a clearly visible difference between Block-1 and Block-2 *n*, represented by the first set of solid grey bars on the left. While participants' mean sample size in Block 1 was quite similar to the rival's small sample size (grey hatched bar), as soon as the rival no longer constrained the participants behavior in Block 2, their sample size rose from a mean of 9.24 (SD = 2.47) to a mean of 15.57 (SD = 8.293), $t(26) = -4.34$, $p < .001$, $d = -0.84$.

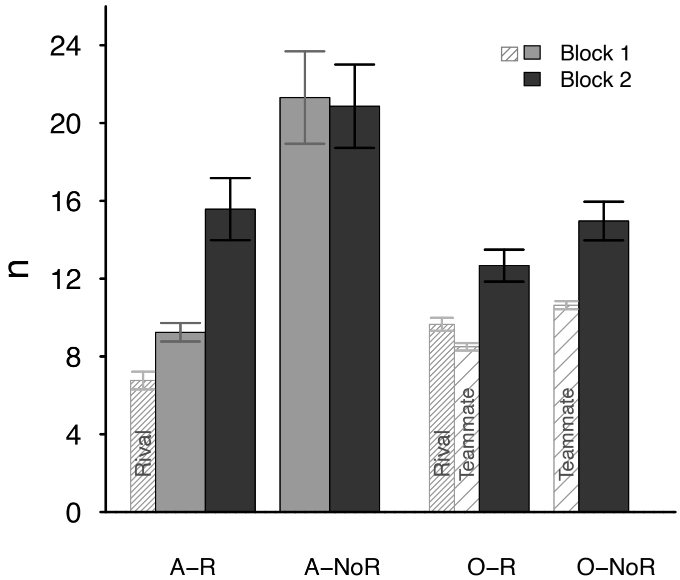

Figure 4. Average size of samples (*n*) gathered by participants in Block 1 (solid grey bars) and in Block 2 (solid black bars) compared to the sample size generated by the algorithms used to simulate rivals and teammates (hatched narrow bars).

As expected, there was no difference between the two blocks for the standard A-NoR Condition, as no intervention had been included (M_{nB1} = 21.31, SD_{nB1} =11.40; M_{nB2} = 20.87, SD_{nB2} =10.28, $t(22) = -0.55$, p = .585 d = -0.12). Turning to the comparison of conditions'

Block-2 sample sizes allows for further assessment of a potential transfer effect. This will be achieved by using the A-R Condition as reference category for the regression analysis and by including dummy-coded predictors representing the comparison between the reference category and A-NoR, O-R, and O-NoR Condition, respectively. That way, we can assess whether there are differences in the conditions' sample sizes. Specifically, the comparison of the A-R with the A-NoR's sample size *n* can show whether the active-competitive intervention in the A-R Condition had a transfer effect to the standard task in Block 2 (in comparison with standard A-NoR Condition). The test of the first predictor's standardized coefficient shows that the A-NoR Condition's *n* is different from the A-R Condition's *n.* Hence although A-R participants increased their sample size in Block 2 (M_n = 15.57; SD_n = 8.293) compared to Block 1, they still used substantially smaller samples than the standard A-NoR Condition in Block 2 (M_n = 20.87, SD_n = 10.25; β = .29, $t(48) = 2.57$, $p = .012$). So although participants were indeed moved to pursue a strategy much faster than the one they seemed to cling to so persistently in previous experiments (Fiedler et al. 2021), the return to a slower strategy instead of the maintenance of the profitable speed focus in Block 2 makes it doubtful that they realized the superiority of faster strategies, which would have allowed them to behave adaptively.

Block 2: Setting apart the social-competition intervention's influences. Nevertheless, the active-competitive intervention proved more effective at reducing participants' sample size than all previous attempts detailed in Fiedler et al. (2021). The question that such a finding should, of course, prompt is whether the uniquely active-competitive aspects of the intervention really are the crucial causal factors driving this effect, or whether there might be more general factors at play. The two additional intervention conditions in our design allowed us to investigate this question.

The active social-competition intervention can be decomposed into different features that characterize it and that might be crucial for the intervention to work. The first potentially crucial aspect is the active involvement in the competitive pressure. It might be necessary to be actively involved in the competition for it to have a lasting effect on the employed sample size. To assess whether this aspect played a crucial role for the transfer effect, we devised the O-R Condition, in which the social aspect and competitive pressure were maintained, but the active aspect was removed. Instead of actively playing themselves, the participants observed their teammate play the game. This maintained a social aspect and presumably kept engagement in the task high, because participants knew they would be awarded the teammate's payoff. The (virtual) teammate had to compete against the virtual rival, who followed almost the exact same strategy as the rival in the Actor-Rival Condition, so the informational value of the conditions was very similar for the participants, while their experience was different.

The results show that the active involvement in a competition in Block 1 had no incremental benefit for Block 2 over merely observing a competition: the A-R Condition (M_n = 15.57, SD_n = 8.23) and the O-R Condition (M_n = 12.67, SD_n =4.27) had very similar sample sizes in Block 2 that did not differ significantly, as the test of the of the standardized dummy-coded predictor representing the comparison between the O-R and the A-R Condition showed (β = -.17, $t(52) = -1.47$, $p = .145$). If anything, the negative value of the coefficient indicates that the O-R Condition's sample size tended to be even smaller than the A-R Condition's, and therefore even further removed from the standard A-NoR Condition.

We also wanted to critically assess the most fundamental aspect of the intervention, which was, of course, the competitive aspect. Instead of watching their virtual teammate play against a virtual rival, participants in the O-NoR Condition merely watched their teammate play the standard version of the speed-accuracy trade-off task employing a strategy similarly fast as the virtual rivals in the A-R and O-R Conditions. The Observer-No-Rival Condition's mean sample size (M_n = 14.96, *SD* =5.06) also did not differ from the Actor-Rival Condition's mean sample size, as the third predictor's coefficient showed (β = -.03, $t(51) = -.31$, $p = .761$). Hence, the active-competitive aspects of the intervention did not have any incremental benefit for Block-2 performance over mere observation of a fast strategy.

This means that all interventions had a comparable effect on reducing sample size in Block 2. All intervention conditions' sample sizes (A-R, O-R, O-NoR's joint M_n = 14.40, SD_n = 6.19) were lower than the standard A-NoR Condition's sample size (M_n = 20.87, SD_n = 10.25; $t(101) = 3.76$, $p < .001$, $d = -0.89$). Accordingly, it seems most plausible that the effective experience in all conditions was the fairly strong focus on speed through the use of small samples. Rather than competitive pressure leading participants to deliberately pursue a fast strategy (because they perceived it to be most efficient), it was the more basic fact that they executed or observed decisions speedier than those they would have pursued themselves, which spilled over to the standard task, regardless of metacognitive processing or insight.

While sample size *n* is the most central measure to assess because it is the only variable participants directly influence, it is interesting to examine how the differences in *n* (or lack thereof) are reflected in the remaining variables. Analogous to the regression analysis depicted above, we conducted dummy-coded regression analyses comparing all conditions individually against the reference condition A-R for accuracy *a*, number of trials *k*, and efficiency *Eff*.

Accuracy. For accuracy *a*, there were no detectable differences between any of the four experimental conditions (A-NoR vs. A-R: β = .17, $t(48) = 1.41$, $p = .161$; O-R vs. A-R: β = -.13, $t(52) = -1.11$, $p = .271$; O-NoR vs. A-R: β = .10, $t(51) = 0.85$, $p = .398$; see Table 2 for descriptive statistics). So although sample size was substantially lower in the A-R than the A-NoR Condition, this difference was not reflected in a measurable decrease in accuracy.

Speed. This was different for the number of completed trials *k*, which – inversely proportional to *n* – was greater for the A-R than the A-NoR Condition (based on the analogous regression analysis: β = -.26, $t(48) = -2.23$, $p = .028$; see Table 2 for descriptive statistics). Again, this effect was not unique to the A-R Condition's active-competitive intervention, since the number of completed trials *k* was similar in the O-R Condition (β = .12, $t(52) = 1.00$, $p = .271$), and the O-NoR Condition (β = -.04, $t(51) = -0.36$, $p = .719$). The lower *n* in the intervention conditions (A-R, O-R, O-NoR) compared to the A-NoR Condition allowed for the completion of more trials in the given time limit.

Efficiency. Finally, the speed-up in all intervention conditions did not result in higher payoffs (see Table 2). The A-R Condition did not earn noticeably more than the A-NoR Condition, as the analogous regression analysis showed (β = -.22, $t(48) = -1.87$, $p = .065$). As can be expected from highly similar *k* and *a* values across the three intervention conditions, their earnings did not differ either (A-R vs. O-R: $\beta = .02$, $t(52) = 0.17$, $p = .868$; A-R vs. O-NoR: $\beta = .03$, $t(51) = 0.27$, $p = .821$).

Table 2
Means and SDs (in parentheses) of the dependent variables for Block 2 by condition

Cond.	Mean *n* (*SD*)	Mean *a* (*SD*)	Mean *k* (*SD*)	Mean *Eff* (*SD*)
A-R	15.57 (8.29)	.613 (.131)	71.41 (25.70)	42.15 (11.51)
A-NoR	20.87 (10.27)	.663 (.119)	56.09 (24.08)	36.09 (14.14)
O-R	12.67 (4.27)	.575 (.136)	78.00 (25.13)	42.67 (9.13)
O-NoR	14.65 (5.06)	.642 (.110)	69.00 (21.88)	43.00 (10.87)

Decomposing interpersonal performance variation. The relations between those variables may again be informative; they are visualized in Figure 5 by a second set of scatterplots analogous to Figure 3 above. What the scatterplots are best suited for showing is that the participants in different conditions generally did not underlie different constraints. The relations between *n* and *a, n* and *k,* and *n* and *Eff* are very similar for all conditions, constrained by the nature of the task. The standard A-NoR Condition just tended to rely on larger samples, meaning that their symbols are shifted to the right toward the larger values on the x-axis, while still constrained by the general relation. However, because many participants in the A-NoR Condition chose larger samples, they forfeited rather large proportions of payoff. Reflecting their tendency toward larger samples in Block 2, the difference between the standard A-NoR Condition and the intervention conditions (A-R, O-R, O-NoR) displayed in the three scatterplots (Figure 5) is more blurred in comparison with the first set of scatterplots that illustrated the clearer difference between the Block-1 performance of the Actor-Rival and the Actor-No-Rival Condition above (see Figure 3). The Block-2 correlations by condition are detailed in Table S1 of the Supplementary Materials.

Interim Summary. On the one hand, there is a clear difference between the intervention conditions and the standard A-NoR Condition, with the former employing smaller samples that lead to slightly lower accuracy but many more trials and consequently higher efficiency. On the other hand, however, this difference is far less pronounced than one might expect, especially in comparison with the pronounced Block-1 difference between the A-R and A-NoR Condition(visible in Figures 3 and 4). The difference between the A-R and A-NoR Conditions' sample sizes in Block 2 (Figure 5) is much more blurred than in Block 1 (Figure 3) because the A-R Condition increased their sample size as soon as the external pressure of the rival was no longer present. Similarly, the two observation conditions exhibited sample sizes quite a bit larger than those they observed the teammate and rival employ. All in all, what seems like an adaptive behavior change at first glance may mostly reflect a superficial reaction and hence pseudo-adaptive behavior instead of indicating adaptive change based on successful metacognitive regulation.

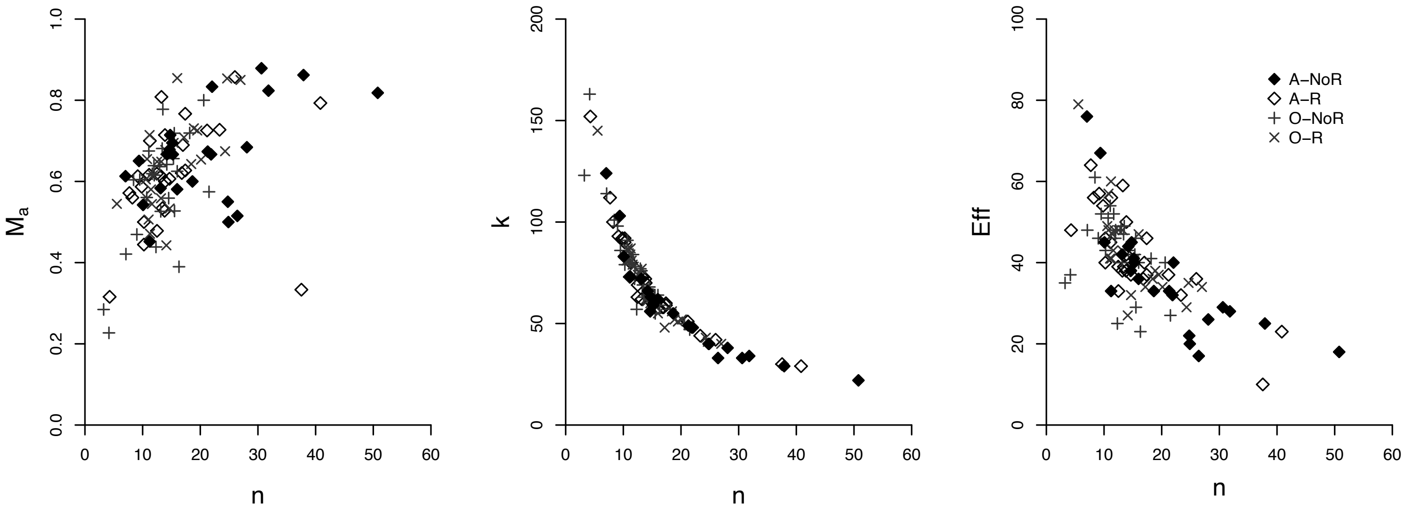

Figure 5. Scatterplots relating M_a, k, and *Eff* to *n*. Black diamonds represent Condition A-NoR (standard control condition), white diamonds the A-R Condition, pluses the O-NoR Condition, and crosses the O-R Condition.

Post-experimental questionnaire

Assessing participants' beliefs about their own performance and about different strategies can shed light on why they did not pursue faster, more efficient strategies. Since monitoring how many decisions one answers correctly can be an important step in improving one's performance, deficiencies in this area might prevent participants from improving. Therefore, participants' correctness estimates will be examined first. On the control side of metacognitive regulation, inaccurate beliefs about which strategy is most efficient can lead participants to execute an inefficient strategy outright, potentially even foregoing further monitoring. Participants' efficiency ratings of fast, moderate, and slow strategies, as well as the typical sample sizes reported for those strategies will allow us to determine whether participants' beliefs contributed to their pursuit of inefficient strategy.

Correctness estimates. Adequately estimating one's correctness rate, that is, the percentage of trials answered correctly, plays an important role for good performance at the speed-accuracy trade-off task. Overall, participants' correctness estimates for Block 2 , M = 71.29% (SD = 14.07), fell short of their actual correctness rate, M = 81.08% (SD = 6.36). The mean of the individual differences between estimated and actual correctness was 9.79 percentage points (SD = 13.51). The vast majority of participants underestimated their accuracy (77.67%). At the same time, the correlation between participants' actual and estimated correctness rate was positive and moderate, $r = .37$, $t(101) = 3.94$, $p < .001$.

Participants' underestimating their overall correctness might already have led them to draw larger samples. What is even more central to performing well than the overall correctness estimates, however, is the estimate of one's correctness for decisions based on small samples, which the questionnaire also assessed. Decision based on small samples were defined as those based on less than 16 observations for both options taken together. Participants estimated their correctness rate for decisions based on small samples to be 58.96% (SD =15.17) for Block 2, while their actual correctness rate was 77.02 % (SD = 23.92).[8]. The mean of the individual differences between their actual and estimated small-sample correctness rate was 18.06 percentage points (SD = 29.88) and a majority of 83.15 % underestimated their correctness for decisions based on small samples. Hence, participants underestimated their correctness rate based on small samples substantially. In addition, there was no correlation between estimated and actual small-sample correctness rates, r = -.12 , t(87) = -1.17, p = .245, indicating that participants were not calibrated to their actual correctness for decisions based on small samples. This makes it likely that there might have been difficulties as early as the stage of metacognitive monitoring that prevented participants from discerning the most efficient strategy.

Efficiency ratings. Central to this speed-accuracy trade-off task is the clear-cut advantage of speed over accuracy, meaning that small samples lead to substantially higher payoffs than large samples. An interesting question is whether participants recognized this advantage. Participants' efficiency ratings of a slow, moderate, and fast strategy can answer this question (see Figure 6 and Table 3). They revealed that overall participants rated the slow strategy, M = 53.79 (out of 100), SD = 23.68, as more efficient than the fast strategy, M = 44.23, SD = 21.18, t(91) = 2.49, p = .015, d = 0.26. To assess how participants' ratings differed for the fast compared to the slow strategy, we subtracted the efficiency rating for the slow strategy from the rating for the fast strategy at the individual level to form an efficiency difference rating. We subjected this efficiency difference rating to a regression analysis analogous to those used to assess the behavioral data above (i.e., dummy-coded using A-R as reference condition) to examine whether it differed between conditions. The difference between the two efficiency ratings was larger for the A-NoR Condition than for the A-R Condition (β = -.32, t(48) = -2.70, p = .009). While the A-R Condition rated the two as approximately equally efficient, the standard A-NoR Condition rated the slow strategy as more efficient than the fast strategy (see Table 3 for means and SDs). Neither the O-R Condition nor the O-NoR Condition differed from the reference A-R Condition (O-R vs. A-R.: β = .002, t(52) = -.02, p =.988; O-NoR vs. A-R: β = .08, t(51) = -0.65, p = .515), meaning that like the A-R Condition, they rated the slow and fast strategy as similarly efficient.

To assess whether participants' beliefs also corresponded to their behavior, we computed the correlation between the actually employed mean sample size and the efficiency difference rating (i.e., individual efficiency rating of fast minus individual efficiency rating of slow strategy). This correlation was moderately negative, r = -.41, t(90) = 4.32, p < .001, indicating that people who rated fast strategies as more efficient indeed tended to have smaller sample sizes. We also assessed the correlation between the efficiency difference rating and earned payoff, which was positive and also moderate, r = .40, t(90) = 4.315, p < .001. So there is a considerable correspondence between participants' beliefs and their behavior. Their beliefs, however, were clearly not based on the insight that the fast strategy was the most efficient.

Rather than the fast or the slow strategy, the vast majority of participants (67%) rated the intermediate strategy as most efficient, across all conditions. A within-participants t-test between the efficiency ratings for the moderate strategy (M = 75.62, SD = 19.31) and the

[8] Fourteen participants who did not employ a single sample of less than 16 observations were excluded from this analysis.

individually pooled ratings of slow and fast strategies (M = 49.01, SD = 13.17) revealed a highly significant difference, $t(176) = 9.78$, $p < .001$, $d = 1.02$. Participants clearly failed to recognize that a fast strategy is the most efficient.

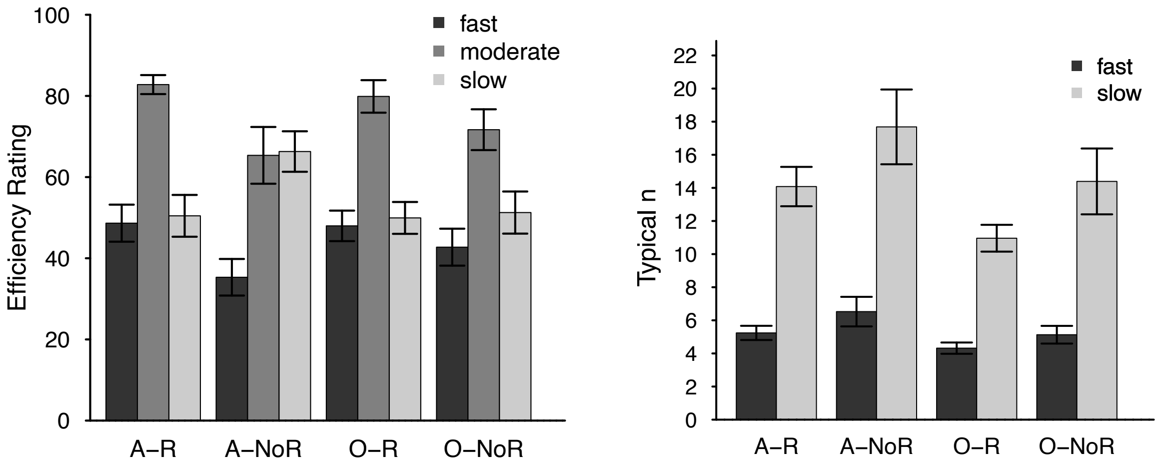

Figure 6. The left panel shows efficiency ratings of slow, moderate, and fast strategies separately for conditions. The right panel shows the sample size participants indicated as typical of a fast and slow strategy. Error bars represent standard errors of the mean.

Table 3
Means and SDs of belief about efficiency of fast, slow, and moderate strategy and reported typical n for fast and slow strategies (both sides)

Condition	Efficiency rating fast strategy	Efficiency rating slow strategy	Efficiency rating mod. strategy	Mean individual fast – slow eff. ratings	Typical *n* for fast strategy	Typical *n* for slow strategy
A-R	48.65	50.45	82.79	-1.81	10.48	28.16
	(22.83)	(25.75)	(11.70)	(34.45)	(4.29)	(11.89)
A-NoR	35.32	66.30	65.36	-30.98	13.05	35.37
	(19.64)	(21.74)	(30.48)	(39.92)	(7.79)	(19.69)
O-R	47.89	49.94	79.87	-1.96	8.64	21.92
	(18.80)	(19.62)	(20.03)	(32.12)	(3.40)	(8.09)
O-NoR	42.73	51.26	71.68	-8.53	10.26	28.78
	(21.83)	(24.88)	(24.12)	(36.62)	(5.13)	(19.08)

Note. Efficiency ratings range from 0 to 100.

Typical sample size reports. To rule out the possibility that different perceptions of what the fast and the slow strategies mean in terms of sample size invalidated the efficiency ratings, we also compared participants' typical sample size reports for fast and slow strategies across conditions.[9] We conducted a regression analysis equivalent to those above on the reported typical sample sizes, separately for fast and slow strategies. The conditions reported typical sample sizes did not differ significantly from each other, neither for the slow strategies nor for the fast strategies (all coefficients βs < .21, *p*s > .108, see Table 3 for descriptive statistics) . Hence, the interventions did not lead to significantly different perceptions of sample sizes, neither for the fast nor for the slow strategy, meaning that the efficiency ratings can be assumed to refer to roughly the same sample sizes.

General Discussion

Summary. We aimed to examine whether a social competition intervention in form of a rival could lead to genuinely adaptive behavior in a speed-accuracy trade-off task. While previous research showed that people responded suitably to a rival while it was present, we extended the scope. To discriminate between pseudo-adaptive behavior and genuinely adaptive behavior we wanted to assess whether the suitable behavioral response was only exhibited while the rival was present or whether it also transferred to the standard speed-accuracy trade-off task when the rival was no longer at work. In addition, to examine whether the rival intervention was, at its core, a social competition intervention, we also tested two variants of the intervention, which either removed the active engagement of participants or both the active engagement and the competition.

In accordance with previous findings, the results show that the rival intervention led to a substantial decrease in sample size while it was in effect, in Block 1. As soon as the rival was removed, however, participants reverted to a higher sample size, although they still performed better than the control group who never underwent an intervention. Instead of displaying genuinely adaptive behavior by continuing to pursue the efficient strategy they had engaged in themselves (or had observed the virtual agents engage in) during Block 1, participants employed substantially larger samples. Hence, the reduced sample size that the Active-Rival Condition exhibited in Block 1 can only be labelled pseudo-adaptive.

Furthermore, the active-rival intervention was not more effective in leading to this small sample-size reduction than the variants of the intervention. Hence, what appeared to be a true social competition intervention based on Block 1 was revealed to be no more effective than simpler interventions, demonstrating that social competition is not necessarily the causally relevant factor in the rival intervention. Since the small decrease in sample size was caused by all interventions, it was likely caused by an aspect they all shared. A contender for this aspect is the exposure to small samples that all interventions imposed on the participants. So while these and similar interventions could be used to bring about a small speed-up in strategy, they do not lead to the realization that fast strategies are superior, which would allow for adaptive behavior and lead to substantial improvements in task performance and payoff gains.

Hence, the intervention that Phillips et al. (2014) inspired did not prove as effective as we had hoped. We hasten to add, however, that it offered some valuable insights nevertheless. It

[9] Eleven participants' efficiency ratings and typical sample size reports were excluded from the above analyses of the questionnaire answers because they indicated a higher typical n for the fast strategy, which was explicitly described as based on smaller samples, than for the slow strategy. They very likely confused the labels for the slow (Strategy A) and fast strategy (Strategy B).

demonstrated that although participants themselves engaged in a very efficient strategy or observed a virtual agent labelled as teammate engage in it in the first block, they were unable to recognize its advantage. This is corroborated by the questionnaire data, according to which participants tended to believe that slow strategies are more efficient than fast strategies, but most of all believed moderate strategies to be substantially more efficient than both slow and fast strategies. Participants seemed to have deliberately employed larger samples, believing this to be the most efficient approach to the task.

This further supports the interpretation that the small speed increase caused by the interventions was not based on insight about the efficiency of strategies. Rather, a more superficial influence of being exposed to smaller samples may have been at play.

The question that suggests itself quite naturally at this point is: Why did participants not recognize the advantage of the fast strategy? Improving one's performance in the trade-off is fundamentally a metacognitive regulation task, so the answer is bound to lie in the area of metacognitive monitoring and control, the two fundamental relations in metacognitive regulation (Nelson & Narens, 1990).

Metacognitive regulation. Monitoring one's performance in the speed-accuracy trade-off task is not a simple feat. It involves keeping track of one's performance (indicated by one's point balance increase) in relation to a particular sample size one is employing. Although not straightforward, this process should have been facilitated by the clear strategy change from Block 1 to Block 2. Participants could have compared the last period of Block 1 to an equivalent period at the beginning of Block 2, in which almost everyone pursued a slower strategy right from the start. We firmly believe that this did not exceed our participants' cognitive capacities. So why could they not use this opportunity to compare the strategies' efficiency? Participants did seem capable of monitoring their correctness rate quite well, at least as indicated by a fairly accurate final estimate of the mean correctness rate. However, monitoring correctness alone does not suffice. Either efficiency (i.e., the point balance), or the number of decisions achieved need to be monitored in addition. It might have been the case that participants did not engage in monitoring of speed or points, leaving performance monitoring incomplete. This would also be in line with the worse correctness estimates for decisions based on small samples. If participants only monitored correctness without relating it to time or a particular sample size they were employing, they would not be able to accurately estimate their correctness based on small samples.

Upon further reflection, it may not be so incomprehensible that participants did not engage in complete monitoring. Monitoring efficiency or number of decisions is intricate because it always needs to be evaluated in relation to time. This is unusual. Although efficiency is crucial in our modern times – especially in production industries, but increasingly in all kinds of work places, including academia – explicitly monitoring one's performance in relation to time is uncommon. Most often, performance is monitored at the level of the individual task, which amounts to monitoring of accuracy. The time dimension is only noted in the extremes – when someone is surprisingly fast (and still excellent) or when someone is disturbingly slow. Continuously monitoring quality in relation to time is rare in people's lives, unless they work in a highly industrialized environment (e.g., in an assembly line). Without further assistance or technical support, it might not have been obvious to participants how important monitoring the temporal dimension is. Instead, participants focused on accuracy, which is typically the measure of performance and, in addition, takes a lot less effort to keep track of, not least because it is more *evaluable* than speed (Hsee & Zhang, 2010).

In addition, it may be the case that people value high accuracy in itself, independently of whether it maximizes their payoff. On the one hand, participants may be deliberately aiming for high accuracy because they explicitly value it. As pointed out in Fiedler et al. (2021), social norms may portray accuracy pursuit as quasi-morally superior, while a focus on speed may be seen sloppy and undesirable. On the other hand, participants may be trying to avoid negative feedback at the level of the individual trial. An efficient strategy implies making very many fast decisions of which a substantial proportion will necessarily be wrong. This might be perceived as quite uncomfortable. Especially in an experimental context, where participant payoffs are not usually excessive, feeling good about oneself during the experiment may contribute substantially to participants' utility. Assisting participants in maximizing their payoff may thus also imply changing their perception of negative feedback at the trial level and even their views of the value of accuracy more broadly.

Strengths. It should be clear by now that the main strength of this experiment is that is has opened up many fascinating new avenues for further investigation of the fundamental processes of metacognitive regulation, which apply well beyond this task. The two other strengths relate to conclusions that the experiment prevented us from drawing. First, the design's ability to discriminate between pseudo-adaptivity and adaptivity allowed us to ascertain that participants were not adaptive despite pursuing an efficient strategy in the first block. Without a test for the transfer effect in Block 2, we might have come to the conclusion that participants exhibited a behavioral response that was suitable for the task conditions, which was not the case.

Second, the intervention variants included in the design allowed us to discover that the rival intervention inspired by Phillips et al. (2014) was not a social competition intervention proper. The same small reduction in sample size in Block 2 was achieved by all interventions, even if they lacked the active engagement of participants or did away with the active engagement and competition altogether. Without these simpler alternative interventions, we might have concluded that it is social competition that causes a sample-size decrease larger than all interventions investigated in previous experiments (Fiedler et al., 2021), which was not the case. This is especially interesting from a more general research methodology perspective. These findings illustrate how important it is to pay close attention to one's manipulation and one's assumption about the causally relevant aspects of the manipulation. Manipulations are hardly ever pure in the sense that they manipulate just a single variable. Hence, while one naturally assumes that any difference in the dependent variable of interest is caused by the aspects of the independent variable that one's focus lies on, one can be quite mistaken. It might be just a single and most basic aspect of the manipulation that is at work. Interestingly, even a good manipulation check cannot safeguard anyone from drawing such overly specific inferences. What we have called diagnostic manipulation check (Fiedler, McCaughey, & Prager, 2021) is a good starting point, since it calls to rule out the most plausible alternatives.

However, a manipulation check alone cannot achieve full diagnosticity in this sense, but must be supplemented by a diagnostic research design, which allows to test alternative and especially more general explanations.

Future Directions. Although we are only at the beginning of this research endeavor, we believe that our speed-accuracy trade-off paradigm holds invaluable potential for investigating metacognitive regulation. Based on the ideas prompted by this experiments, future studies need to investigate which task features inhibit the successful monitoring necessary for arriving at the insight that faster strategies are more profitable or obstruct the control functions that apply that insight. This will reveal important metacognitive functions and their vulnerabilities, but also point to effective ways of prompting or assisting those functions. Identifying such metacognitive functions and their weaknesses and potential for improvements can be useful far beyond this particular task. We believe that the investigation of specific metacognitive functions will lead to general principles for adaptation to any task in a variety of circumstances, which naturally often include other people in a social setting. Be it in a competitive or friendly manner, metacognitive functions can be crucial for learning from others. Identifying them and their moderators can benefit a wide range of situations and the behaviors that needs to be displayed in them.

Conclusion. Effective metacognitive regulation is an important way of arriving at genuine adaptivity, which needs to be distinguished from seemingly suitable behavior that only happens to correspond to it – pseudo-adaptive behavior. Our paradigm shows one way of distinguishing between the two. We hope it encourages other researchers to think about and implement ways of drawing the distinction wherever it applies. Likewise, we hope to have conveyed the importance of using diagnostic designs to rule out alternative explanations, especially more general ones that can easily be neglected. Most importantly, we have shown that metacognitive regulation does not succeed enough to lead to adaptive behavior in our task. Since we are convinced that it can succeed in principle, we believe that it is has great potential to uncover specific metacognitive functions and their facilitators and deterrents. These should generalize to many other tasks and situations, furthering our understanding of how people use metacognitive regulation to respond adaptively to the world they live in.

References

DiMenichi, B. C., & Tricomi, E. (2015). The power of competition: Effects of social motivation on attention, sustained physical effort, and learning. *Frontiers in Psychology*, *6*(September), 1–13. https://doi.org/10.3389/fpsyg.2015.01282

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

Frey, R. (2020). Decisions from experience: Competitive search and choice in kind and wicked environments. *Judgment and Decision Making*, *15*(2), 282–303. http://journal.sjdm.org/19/190114/jdm190114.pdf

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

Hsee, C. K., & Zhang, J. (2010). General evaluability theory. *Perspectives on Psychological Science*, *5*(4), 343–355. https://doi.org/10.1177/1745691610374586

Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. *Psychology of Learning and Motivation*, *26*, 125–173.

Mishra, S., Barclay, P., & Lalumière, M. L. (2014). Competitive disadvantage facilitates risk taking. *Evolution and Human Behavior*, *35*(2), 126–132. https://doi.org/10.1016/j.evolhumbehav.2013.11.006

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

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

Todd, P. M., & Gigerenzer, G. (2007). Environments that make us smart: Ecological rationality. *Current Directions in Psychological Science*, *16*(3), 167–171. https://doi.org/10.1111/j.1467-8721.2007.00497.x

Wulff, D. U., Hills, T. T., & Hertwig, R. (2015). How short- and long-run aspirations impact search and choice in decisions from experience. *Cognition*, *144*, 29–37. https://doi.org/10.1016/j.cognition.2015.07.006

Van De Pol, P. K. C., Kavussanu, M. K., & Ring, C. (2012). The effects of training and competition on achievement goals, motivational responses, and performance in a golf-putting task. *Journal of Sport and Exercise Psychology*, *34*(6), 787–807. https://doi.org/10.1123/jsep.34.6.787

APPENDIX 3: Third Paper

The Information Cost–Benefit Trade-Off as a Sampling Problem in Information Search

Linda McCaughey[21]*, Johannes Prager and Klaus Fiedler*

The search for information before arriving at a decision is a most natural activity. Moreover, the information acquired about options has a major influence on the decision outcome. Cognitive-ecological approaches emphasise the information sample as major determinant of subsequent cognitive processing and decision outcomes and take into account that most entities that are the focus of judgements or decisions cannot be assessed directly – be it the risks involved in a certain therapy or the potential happiness derived from a consumer product. A sample of information – a subset of direct observations of proxies – is the only way to estimate said entities. While sometimes it may be realistic to assume that information is available in the environment and just needs to be processed to form the basis of our decisions, far more often the information sample is the result of an individual's active search process, for example when we choose whom or what to interact with and when to stop. Since cognitive processes necessarily influence (if not fully determine) active search, which in turn determines the information sample that forms the basis of further cognitive processing and decision making, investigating the determinants of active search is imperative.

This perspective is emphasised by other recent theoretical developments outlined in other chapters of this volume: Denrell and Le Mens's concept of hedonic sampling, for example, describes how information search (sampling) is guided by active sampling strategies that serve to attain pleasure and to avoid displeasure and pain, as in the Hot Stove Effect (see Chapter 4 by Denrell & Le Mens in this volume; Denrell & March, 2001). Chapter 9 by Harris and Custers (this volume), illustrates how reward-rich versus reward-poor environments induce different active information search strategies concerning two options as sources. In the first case, in which both options are rewarding, search strategies emphasise exploitation and thus serve to uphold biased initial beliefs about the options. In the second case, however, where both options are meagre, search strategies focus on exploration and thus successfully correct the biased initial belief.

Apart from the influences on source selection in information search, a particularly interesting and important aspect to consider is that of information amount. Information amount, defined as sample size in the case of sampling approaches, tends to be tightly related to decision quality: Usually, the more information one has accumulated about the options, the more comprehensive and accurate one's judgements of them will be – the bigger the (random) sample, the more accurate its estimate. However, information

[21] Corresponding author, e-mail address: linda.mccaughey@psychologie.uni-heidelberg.de

acquisition requires resources such as time or effort, and may even incur financial costs. Hence, one cannot simply act on the maxim 'the more the better'. Instead, the cost of the information (whatever form it may take) needs to be weighed against the benefits of the information in a cost–benefit trade-off.

1 Benefits and Costs of Information: An Inherent Trade-Off

At every step of the information search, one has to decide whether it is worthwhile to carry on searching or whether one should make the decision. Whether it is worthwhile depends on the cost–benefit trade-off, that is, whether the information's benefit outweighs its cost. Information's costs are simply higher the more one has to expend to acquire it. Its benefits, on the other hand, increase with higher expected value, which can either be increased by a higher probability of a correct choice of the better outcome or by increasing the outcomes' values themselves.

How well can we take all these aspects into account to avoid making rash decisions based on too little information, on the one hand, but also to avoid acquiring too much or overpriced information, on the other hand? Put in simpler terms: do we know how much information is worthwhile? In the remainder of this article, we first provide a review of previous attempts to assess information costs in the form of financial and time costs. As an extension of the literature on time costs, we will present our own experiments implementing a sample-based speed–accuracy trade-off that directly links speed and accuracy via sample size. The apparently robust bias towards accuracy found in those experiments will be discussed critically in light of evaluability and with reference to a recent review by Evans et al. (2019), who showed that the optimal strategy varies vastly depending on the specific task parameters and, hence, does not allow for general claims about the optimality of human performance. Our own experiments using financial information costs will be presented to corroborate this claim, showing that the investigation of specific aspects of adaptivity might be a more worthwhile research aim. In preparation for this mental journey, we will start out by looking at previous experiments investigating information costs in the following.

2 Investigating Information Costs

2.1 Financial costs

The question of whether people know how much information is worth- while captivated the attention of the Bayesian approach to judgement and decision making and cognition (Slovic & Lichtenstein, 1971) in the 1960s and 1970s, with mixed findings regarding the issue of information amount (for a review, see Connolly & Serre, 1984 The paradigm most commonly used to investigate the question is elegantly simple and relied on financial costs. The participant was shown two urns with different proportions of two colours of marbles, for example, 70 per cent red and 30 per cent white marbles versus 30 per cent red and 70 per cent white marbles. The urns represented two different Bernoulli distributions (distributions with binary outcomes and a certain probability p of one of the outcomes occurring). The experimenter would flip a coin to determine which urn, that is, which distribution, would be selected without telling the participant, whose task it was to infer which of the two urns had been selected based on a sample of marbles drawn from it. One by one, the participant could then sample observations from the urn at a certain monetary cost, each observation being a marble randomly drawn from the selected urn. They could draw as many or as few items as they wished before making a final decision, the correctness of which was rewarded (and usually, its incorrectness punished). The factors manipulated in this paradigm were mostly those that influence the trade-off between the cost and benefit of information dis- cussed above: the information cost, reward, and punishment (gain and loss) for correct and incorrect choices, and the discriminability of the choice options, that is, how similar the two urns were in terms of their distributions.

Operationalising both information cost and payoff in monetary terms, such experiments demonstrated a certain sensitivity to costs and payoffs on the participants' side. When the information cost was low or the payoff for an accurate choice was high, they tended to search for more information. When the information cost was high or the payoff low, they tended to search less, investing less money in the information search. This sensitivity was, however, limited, meaning that participants' adjustments of the sample size pointed in the right direction but fell short of an optimal strategy specified in Bayesian terms[22] (e.g., Pitz, 1968; Pitz, Reinhold, & Geller, 1969; Edwards, 1965).

Next to information costs and payoffs, one main determinant of the optimal amount to sample is the evidence conveyed by the sampled information,[23] which is determined by the discriminability of the two hypotheses or options and was also examined in the same paradigm. To attain different levels of discriminability for the two options, Fried and Peterson (1969) varied the outcome probabilities of the two options (between proportions of .6 and .4 vs. .7 and .3), with sampled information consisting of certain light bulbs lighting up. The condition that could engage in optional stopping, deciding when to stop during the sampling procedure, was shown to acquire too little information compared to the optimal Bayesian strategy. Another condition, which had to indicate in advance how many observations they wanted to acquire, performed closely to the respective optimal strategy.

Snapper and Peterson (1971) used a very similar paradigm to manipulate the discriminability of the options and hence, the evidence conveyed by the information. Again, a decision had to be made between two options as possible sources of a sequence of numbers. In this experiment, however, the two possible sources were not Bernoulli distributions with binary outcomes, but rather normal distributions with continuous

[22] Since each new observation in the sample calls for a re-evaluation of the evidence in favour of one hypothesis in relation to the evidence in favour of the contrary hypothesis, it is commonly assumed that Bayesian updating of the conditional probabilities of the two hypotheses and their posterior odds is a suitable normative standard of comparison given suitable priors – an optimal strategy (Edwards, 1965). The best minimum odds ratio at which a decision should be made in favour of the leading option, can be determined through simulations that take into account the information cost and payoff parameters.
[23] We take evidence to mean the change in the posterior odds (in favour of one over the other of the two hypotheses) that the data leads to (understood in Bayesian terms).

values as observations. This introduced another way in which the two options could be more or less discriminable, that is, via the difference between their distribution means. The further apart their means were, the less the two distributions overlapped and the more discriminable they were. And the more discriminable the two options, the less information was needed to discriminate them (to the same degree). However, participants did not adjust their sampling behaviour enough in response to the levels of discriminability, which led to oversampling in this paradigm.

In a variation of the paradigm, Hershman and Levine (1970) posed a hypothetical military reconnaissance problem, in which participants had to determine which of two missile mixes was being used. After a first free reconnaissance flight had already yielded a sample of 10 observations of missiles, participants were given the opportunity to buy a second reconnaissance flight that would yield another 10 observations. The cost of the second sample was varied and through varying the proportion of the two missiles, the evidence or diagnosticity of the first sample also varied. Participants, again, mostly overpurchased information, meaning that they tended to buy additional samples even when those were overpriced compared to how little they increased the probability of choosing the correct answer.

Overall, participants appeared to be somewhat sensitive to the different constraints of the tasks, adjusting the amount of sampled information to the payoff, the information cost and the discriminability of the options to some extent. Compared to the optimal Bayesian strategy, however, some experiments showed that participants sampled too much information, while others showed that participants sampled too little. How these mixed findings can be reconciled will be discussed later.

Explicit financial costs are not the only type of cost we are confronted with when searching for information on options. Often, information is freely accessible, and the information cost most important to consider is that of the time one spends on information search into options.

2.2 Time Costs

The cost of time is more difficult to implement in an experimental setting than monetary costs. One way of making time or waiting costly is to implement a time limit after which the potential payoff is foregone (Madan, Spetch, & Ludvig, 2015). A variation of that is to make the payoffs decrease with the passage of time, which is usually implemented to induce time pressure (Hausfeld & Resnjanskij, 2018; Payne, Bettman, Luce, 1996).

Time passes whether we want it to or not, it cannot be spent or accumulated in the same way as money can. Hence, what is actually meant by “time cost” is most often the opportunity cost of time – the cost of foregoing alternative uses for a certain amount or period of time, that is, being unable to use the time to do something else. This conception is reflected more closely in a research design in which a total session time was allocated to work through, say, 50 decisions and each minute taken longer would be punished (each minute shorter rewarded; Rieskamp & Hoffrage, 2008). Although fascinating, this research only offered limited insight into how well participants regulate the amount of information collected before deciding, since it was not related to a framework of optimal regulation. Its focus rather lay on the influence of time pressure on risk preferences (Madan et al., 2015) or on decision strategies (Payne et al., 1996; Rieskamp & Hoffrage, 2008; but see Hausfeld & Resnjanskij, 2018).

Another clever and fairly common way of implementing opportunity cost of time is through a speed-accuracy trade-off. Speed is an inverse function of required time relative to the total time allotted. When accuracy and speed need to be traded off against each other, emphasising accuracy at the expense of speed means that time is invested to achieve a certain degree of accuracy, accepting the opportunity cost of lower speed, which consists of the rounds or trials one could have worked through instead, which will be irrevocably lost.

Speed-accuracy trade-offs are very common in perceptual research, where they are often modelled by drift-diffusion models, mostly based on random dot motion tasks. Processing perceptions of moving dots to judge the direction that they tend to move in is not directly comparable to sampling as information search. Nevertheless, the typical finding that participants tend to overemphasise accuracy (waiting too long to make a decision; Evans et al. 2019) will prove relevant. In one of the few instances that are not perceptual, Jarvstad, et al. (2012), constructed speed-accuracy trade-offs not only with perceptual but also with higher-order cognitive (simple arithmetic and mental rotation) tasks, of which as many were to be completed correctly as possible in the specified time period of 2 minutes. Both types of tasks required cognitive processing to be solved, meaning that spending more time on an individual task increased the chances of being correct and obtaining the reward, but it also incurred the cost of that time not being available for later decisions, thereby decreasing the overall number of tasks completed in the time period. Hence, there was a trade-off between the speed at which the tasks were completed and the (expected) accuracy at which they were completed. Even though information search was not the focus of those studies or, in fact, any studies we know of investigating speed-accuracy trade-offs, they can be adapted in a way that makes them well suited to investigating the opportunity cost of time in information search.

3 Linking Speed and Accuracy more Directly through Sample Size

Adapting a speed-accuracy trade-off for the investigation of time costs in information search first requires that the constituting task be a decision based on *sampled* information. Since sample size has a given relation to one dimension of the trade-off, accuracy, the relation to the other dimension, speed, has to be established next. This is achieved by tying the amount of acquired information (i.e., sample size) directly to the elapsing time by having the sample increase at a steady rate of time. This specification achieves a more direct (and also somewhat more objective) relation between accuracy and speed than is present in other speed-accuracy trade-off, where the trade-off depends on the participants' processing, attention, reaction time, and many other factors.

Following these considerations, we (Fiedler, McCaughey, Prager, Eichberger & Schnell, 2020) constructed and implemented a novel speed-accuracy trade-off in sample-based choices, which allowed us to operationalize speed and accuracy in terms of the same joint scale, namely sample size. To the best of our knowledge, this is the only experimental paradigm so far that combines information sampling with a trade-off task that links speed and accuracy to the same natural measure of information amount, namely the size of a binary sample of choice options' outcomes.

At a more concrete level, the paradigm took the form of an investment game that had participants decide between pairs of stocks. The stocks were described by samples of the changes in share price, which were explicitly described as being drawn randomly from the last 100 trading days. Each sample or information was displayed in the form of an arrow that indicated a price increase when it was coloured green and pointed upward, or represented a price decrease when it was coloured magenta and pointed downwards (see Figure 1 below). The aim was to identify and choose the stock that had a larger proportion of positive share price changes within the last 100 trading days. Identifying it correctly resulted in a reward of points, while making the wrong decision implied a loss of points.

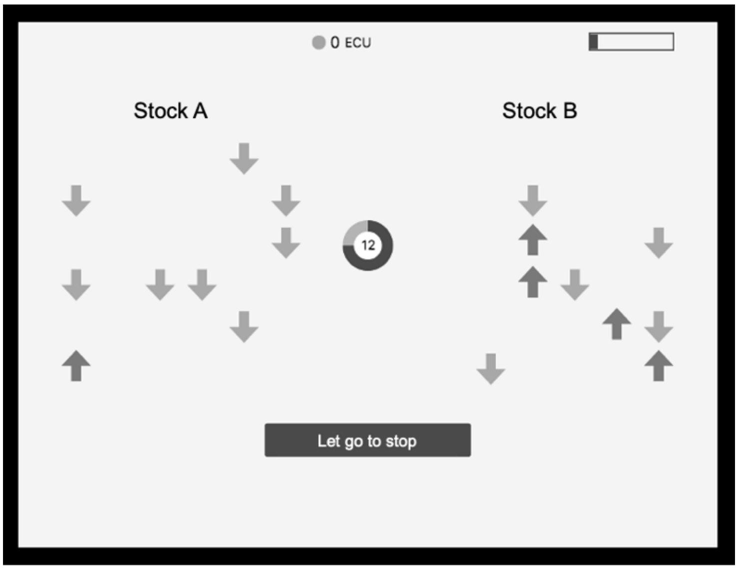

Figure 1: Graphical illustration of a fund-investment choice task developed by Fiedler et al. (2020) to investigate speed-accuracy trade-offs in sample-based choice tasks

The price change information for the stocks was implemented by drawing samples from Bernoulli distributions with certain probabilities of a price increase, $p(\uparrow)$, that could range from .1 to .9 (in steps of .1). The difference between the two probabilities (of price increases) of stock pairs A and B, $\Delta p = p_A(\uparrow) - p_B(\uparrow)$, was manipulated and could take the values $\Delta p = .1$, .2, or .4. This led to a dynamic difficulty parameter, with some trials being easier because the stocks were easier to discriminate (e.g., $\Delta p = .4$) and other rounds being harder because the stocks' *p*s were closer together (e.g., $\Delta p = .1$).

As indicated above, what made it a speed-accuracy trade-off was that participants had a limited total time (e.g., 20 minutes in Experiment 1) and that the sample information appeared at a certain rate of 500 ms per arrow. Hence, accuracy would increase with time spent waiting for more information, but the amount of rounds accomplished, that is, speed, would decrease with time spent waiting for information in a particular round. Accuracy had to be traded off against speed, since both had advantages and disadvantages with respect to the points one aimed to win by making correct decisions. Crucially, the number of sampled prior outcomes (sample size n) afforded a direct quantitative link between accuracy and speed, because n is functionally related to both trade-off components and this two-fold functional relationship is well defined and understood.

Given a presentation speed of 500 ms per arrow, the time required for a sample of n arrows amounts to $n/2$ seconds, and the total time period is sufficient for 1200 seconds (=20 minutes) divided by ($n/2$) samples of average size n (plus an additive constant for choice execution). By comparison, statistical sampling theory tells us that the standard error *SE* of the sample mean (i.e., the expected inaccuracy) decreases with increasing n according to the formula SE = Standard Deviation / $\sqrt{n}$. Thus, while the number of choices completed in the given time decreases linearly with sample size n, the corresponding increase in accuracy is clearly sublinear (i.e., the rate of inaccuracy shrinkage is $\sqrt{n}$ rather than n). Sampling n more observations costs n more time units but only yields a standard error decrease of $\sqrt{n}$. Thus, for many parameter specifications (above-chance accuracy; reasonable information costs; fairly even payoff for correct and incorrect choices), maximizing the expected total payoff (i.e., the average payoff gained from all completed choices) calls for a fast strategy, sacrificing accuracy. In other words, over a wide range of parametric conditions, speedy small-n strategies will be superior to accuracy-based large-n strategies. Although larger samples produce more accurate choices than smaller samples, after the first few items, additional items tend to decrease the number of completed choices more than they increase the average payoff per choice.

3.1 Substantial Oversampling

Yet, a whole series of experiments employing the speed-accuracy trade-off paradigm described above consistently demonstrated that participants collect far too large amounts of information. Owing to the different relations of speed and accuracy to sample size, the strategies that lead to the most points are very fast, meaning that they are based on very small samples. The reason is that each arrow requires the same amount of time to arrive, but it does not always bring the same increase in (expected) accuracy, as is illustrated in Figure 2. Earlier arrows lead to a bigger increase in accuracy – they are “worth more” - than later ones, but still incur the same cost in terms of time. Hence, instead of “wasting” the time on additional arrows for one round, in most cases one should rather invest that “arrow time” in a new round, where the arrows are “worth more” again.

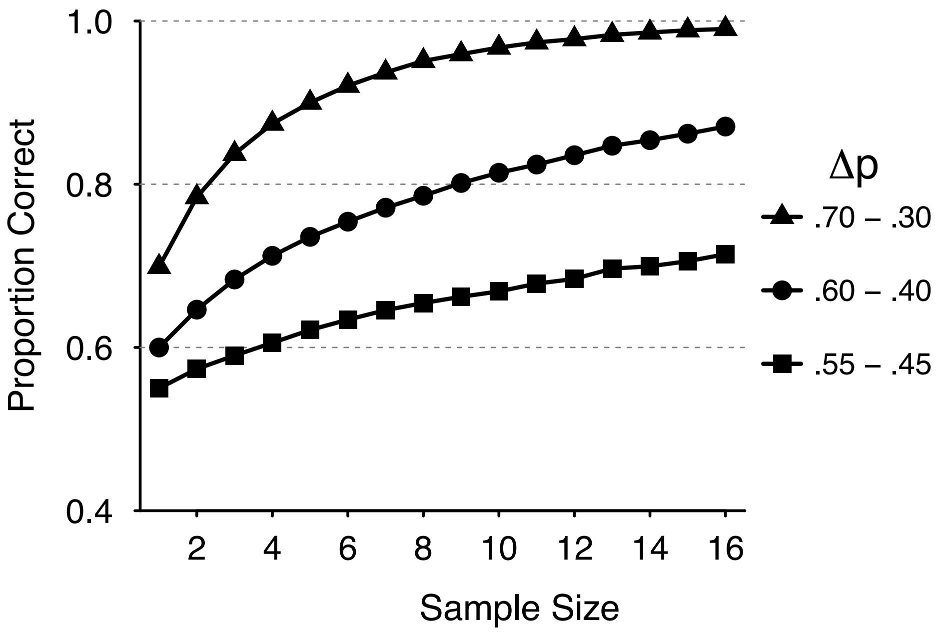

Figure 2. Simulation results of the proportion of correct decisions as a function of mean sample size n based on an algorithm that always samples n observations. Sample size increases when mean n is lower lead to a higher accuracy increase than sample size increases when mean n is higher.

Participants, however, pursued strategies that were much “slower”, focussed on accuracy, engaging in a degree of oversampling that substantially reduced the amount of financial reward at the end of the experiment. Moreover, they did not change their strategies over the course of the experiments as a function of learning by experience, nor were they responsive to various interventions that clearly highlighted the advantage of speed.

In a first experiment, there was no explicit outcome feedback to demonstrate to participants that choices based on smaller samples were almost as likely to be correct as choices based on larger samples; participants had to adjust the self-determined sample size to the experienced ease or clarity of smaller or larger samples (as in a calibration experiment calling for confidence judgments under uncertainty). However, in all following experiments, they received explicit trial-by-trial feedback, making it clear that the accuracy hardly increased with markedly increasing sample size. Yet, although they experienced that investing time into gathering very large samples produced at best a negligible gain in accuracy, such regular and transparent feedback was not successful in helping participants to reduce their persistent oversampling.

Neither could the accuracy bias of almost all participants' oversampling strategies be reduced when sample size was limited experimentally, forcing them to experience the (hardly reduced) accuracy resulting from smaller samples (i.e., when maximal sample sizes were reduced from n_{max} = 18 and 12 to n_{max} = 12 and 8). Moreover, to avoid anchoring effects and unwanted influences of externally imposed limitations, we also included an "unlimited" condition, where no explicit limit restricted self-determined sample sizes. This intervention served to show that participants were, if anything, constrained by the limits rather than nudged to sample more by them. When all limits were removed, participants tended to sample even more.

Another intervention that forced participants to make do with smaller samples was procedural in nature. When every new arrow had to be solicited by an active keystroke (rather than continuing to sample for as long as the space bar was pressed), the resulting sample size was actually reduced from about M_n = 8.23 (SD_n = 3.35) in the normal passive-sampling condition to M_n = 7.49 (SD_n = 3.61), but to such a small extent that it stands to reason that the reduction was due to the procedural nudge rather than insight into the advantage of reduced sampling.

In another experiment, indeed, the accuracy bias was shown to persist even when participants had no losses to fear because the payoff schema was such that correct choices were rewarded while incorrect changes were of no consequence. All they stood to lose was time; their only punishment was having foregone the opportunity of gaining a point. And still, they did not seem to be able to give up their fondness for accuracy, continuing to sample too much.

3.2 Pseudo-adaptive Trade-Off Regulation

While the results detailed above quite clearly demonstrate that participants' behaviour is far from optimal, results from other experiments, which we will describe below, seem to present a more hopeful picture. This picture should rather be likened to an illusion, however. The results that may suggest, at first glance, that participants are solving the trade-off well, upon further examination reveal that it was merely external pressures or blatant demand characteristics that caused the behaviour. This cannot be counted as adaptive behaviour, which is usually understood as (sustainable) behaviour that matches the structure of the environment (e.g., Todd & Gigerenzer, 2007), as the match presumably occurred mainly because of the external pressure and not a response to the task structure itself. Taking the pressure away and leaving the task structure the same should result in the same behaviour if the previous behaviour was adaptive, but should revert to less well-matched behaviour if the change was for reasons of pressure alone. The latter is an instance of what we want to call pseudo-adaptivity: behaviour that resembles adaptive behaviour but does not occur as a (suitable) response to the context, but rather, for other reasons or by happenstance.

It is also important to differentiate between optimality and adaptivity or pseudo-adaptivity in this context. Here, a suitable definition of optimality refers to behaviour that maximise reward, whereas deviations from optimality are tantamount to behaviour that fails to maximise the reward. This can be and often is interpreted normatively as “bad” or “irrational”, which we will criticise below. Adaptivity, on the other hand, requires a broader framework and cannot just be diagnosed through a simple (or even one-off) deviation from optimality. Since adaptivity implies flexibility and suitable responding to task structure and changes, it is not enough to observe behaviour that happens to be in line with the optimal strategy to infer adaptive behaviour. One needs to make sure that the behaviour actually occurs in response to the task’s characteristics and can be flexible (or is at least suitable for a range of different task characteristics or parameter values). Optimality comparisons can be very helpful, but are only a tool to assess adaptivity. Assessing adaptivity was precisely an aim of another experiment that used a variation of the speed-accuracy trade-off task. In this variation, inspired by Phillips et al. (2014), participants had to play an opponent at an adapted version of the speed-accuracy trade-off task in one condition. In another condition, they watched a teammate play the opponent before engaging in the standard speed-accuracy trade-off task themselves (McCaughey et al., 2021). To allow for a competitive version of the trade-off task, both players saw the same sample emerge concurrently, but only the player who decided to truncate the sampling process could make a decision and potentially win or lose, with the other player having to wait. The active condition saw the participants engage with the virtual opponent themselves, while in the passive condition participants watched a virtual teammate play against the opponent. In both conditions, the programmed rival’s strategy was close to the optimal strategy, which the virtual teammate was programmed to match, but which also forced participants to acquire smaller samples, thus exhibiting a strategy that was closer to optimal. Nevertheless, in the subsequent (single-player) standard trade-off task, both conditions subsequently reverted to overly accuracy-focussed strategies (although they differed significantly from the control group). This is a strong indication that the rival condition’s apparent proximity to optimality was for reasons unrelated to insight into the problem or other causes that could generate flexible, adaptive behaviour, but was rather caused by external pressure from the rival and can therefore be considered pseudo-adaptive. The pressure was such that if sampling was not truncated early enough, no chance at making a decision was attained at all. Pursuing a faster strategy can plausibly be assumed to be a superficial reaction to missing out, rather than recognition that this is a good strategy for the task, which would underlie adaptive behaviour.

4 Avoiding Premature Conclusions

Thus, strong and persistent evidence from a series of several experiments seems to converge on the conclusion that speed-accuracy trade-off performance in sample-based choices are plagued by a systematic impairment or bias that cannot be easily overcome with the help of interventions. Against the background of the relation of speed and accuracy with sample size in our paradigm, this may appear very plausible. Given that speed decreases linearly with increasing sample size, whereas accuracy increases in a clearly sub-linear manner, the reported findings could be a reflection of the well-known difficulty to understand non-linear functional relationships (Larrick & Soll, 2008;

Svenson & Eriksson, 2017). Apparently, in the context of a sample-based choice task, in which speed and accuracy are linked to each other via sample size, the optimal strategy calls for speed rather than accuracy, and relative to this optimum, almost all participants exhibit oversampling.

However, there are two important reasons why these conclusions, as plausible and empirically justified as they might appear at first, would be premature and unwarranted. The first is that drawing a general inference about the optimality or sub-optimality of human performance from a comparison with an optimal strategy (for speed-accuracy trade-offs) on closer inspection turns out to be untenable. As Jarvstad et al. (2012) have anticipated with respect to the alleged rationality of perceptual in contrast to conceptual trade-off tasks, and as later articulated explicitly by Evans, Bennett, and Brown (2019), generalised inferences based on the comparison with an optimal strategy are highly problematic because they are far less general than they first appear. As Evans et al. demonstrated, the optimal strategy – defined as the strategy that maximises the payoff for the given task – depends very much on the specific parameter settings of a task. This means that even for a speed-accuracy trade-off task, one can specify parameters such that strategies favouring accuracy would be optimal (e.g., payoff with high stakes or large asymmetries, time delay as punishment for incorrect choices or very difficult choices $|\Delta p| < .10$). Hence, assuming that people generally oversample would definitely be premature. An experiment would have to include a range of parameter settings implying different optimal strategies to support more general inferences about people's performance.

The second reason why a persistent bias against speed in the reported experiments may not reflect a fundamental cognitive or meta-cognitive deficit lies in the particular task structure and its implications for what Hsee and Zhang (2010) call evaluability, that is, how easily a certain aspect or dimension can be evaluated. One reason why sensitivity to speed may have been inhibited, creating an overweighting of accuracy, is that accuracy is much more evaluable than speed. Whether a choice is correct or incorrect, leading to a positive or negative payoff, can be evaluated naturally and immediately for each trial, on a clearly defined categorical (dichotomous) scale. However, speed differences between fast and slow choices are much less evaluable. Whether information sampling speed was "too high" or "too low" on a relative efficiency scale will at best be evident at the end of the experiment, when the total payoff is visible, or maybe from a comparison of two strategies played in two (consecutive or simultaneous) experiments. Because of this fundamental problem, the failure to solve the speed accuracy trade-off and the apparent insensitivity to speed may, to an unknown degree, reflect a task-inherent imbalance in evaluability. Whether the choice on a single trial is correct or not, is easily and naturally evaluable, but there is hardly a benchmark or enumerable scale for whether the present choice is fast or slow. To avoid this imbalance, it would be necessary to use an experimental task that operationalises the cost-benefit trade-off in dimensions that are both similarly evaluable.

5 Information Cost Experiments: Context Changes and Normative Inferences

To substantiate and empirically test these two considerations concerning task-dependent optimality and evaluability, we decided to conduct further experiments with a different experimental task that closely resembles the urn task in the Bayesian tradition described above. To keep it similar to the speed-accuracy trade-off task, participants had to decide whether the underlying population, which could take a probability of .1 to .9 (excluding .5), was above or below .5, instead of judging which urn was the source of a sample. Similar to the difficulty in the speed-accuracy trade-off task, this adds a dynamic element to a series of trials with different *p*-parameters: a population with an extreme $p = .9$, will lead to samples that more quickly support the likelihood of one hypothesis over the other. A population with a moderate $p=.6$, on the other hand, will tend to convey more conflicting/cancelling observations and will take longer to reach the same level of support for one over the other hypothesis.

This task differs from the speed-accuracy trade-off task in a few notable ways. Most importantly, it specifies information cost in terms of financial costs instead of time, so the cost-benefit trade-off is no longer operationalised in different dimensions, speed and accuracy. This leads to a number of useful consequences with regard to the first consideration, concerning the comparisons with optimal strategies. It allows for very easy changes in the cost and payoff parameters that imply different optimal strategies. Hence, it allowed us to expand the range of information costs relative to the payoffs to such a degree that optimal strategies called for either small, moderate or quite large samples. Instead of implementing just one optimal strategy in an experiment, this means that participants' behaviour can be compared to a whole range of different optimal strategies. Implementing this range in a repeated-measures design with varying cost-payoff ratios across different trial blocks is also useful because it may sensitize participants to this variation (Fischhoff, Slovic & Lichtenstein, 1979). Specifying costs and payoffs and the resulting optimal strategy is also part of the Bayesian tradition. However, this was mostly done to compare human performance to Bayesian integration. Our focus lies on changes in the parameters with the aim of studying human adaptation.

With regard to the second issue outlined above, evaluability, this new task likewise had advantages. We could quantify the information costs (for one additional observation sampled from a distribution) and the payoff for a correct choice in terms of the same measurement unit or "currency", such that participants could see, on every trial, what portion of the entire payoff for a correct decision they had to pay for each new observation. For instance, when each observation cost 10 points and a correct choice was rewarded with 100 points, they presumably understood that a sample larger than 10 ($n>10$) was more expensive than the profit to be gained from a correct choice. In addition to equal evaluability of information cost (as financial costs instead of time) and accuracy (payoff) in terms of the same "currency" with equal visibility on every trial, the salience of the information costs was further enhanced by letting participants "purchase" every additional piece of information. This means that participants were reminded of the price of an increasing sample each time they actively solicited a new observation.

6 Research Design to Illustrate Task-Dependent Optimality and Evaluability

Figure 3 provides an example of how the sample-based choice task parameters were specified in recent experiments by McCaughey, Prager, Fiedler (2021). The four columns indicate the four blocks participants worked on, which consisted of 25 or 30 decisions each, depending on the experiment. Below each block label, the ratio of the information cost per sampled item to payoff for a correct or incorrect choice is displayed. For instance, a ratio of 5:100 for B1 indicates that each sampled observation cost 5 units while the gain for a correct response and the loss for an incorrect response amounted to 100 units, with the standardised ratio of 1:20 ratio indicated in the bottom row. Likewise, the (standardised) ratios were 1:40, 1:80 and 1:10 for B2, B3, B4, respectively, thus making information most inexpensive for B3 and most expensive for B4.

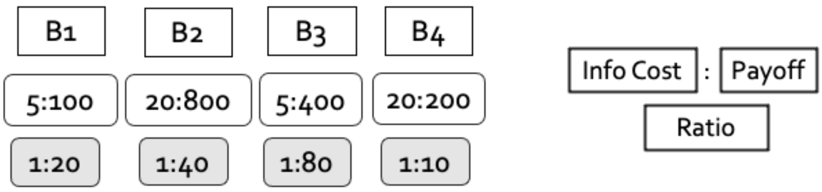

Figure 3. Schematic illustration of one of the sequences of cost parameters used in the sample-based decision task for Block 1 to 4, with displayed information costs and payoff in the middle row and the standardised ratio in the bottom row.

The information that participants could sample at the indicated cost again consisted of arrows that were either green and upward pointing or magenta and downward pointing, similar to the previously described speed-accuracy trade-off task. However, a decision had to be made on the basis of only one instead of two samples. The task indicated to participants that the arrows represented gains and losses that a fictitious slot machine had yielded in the past. Based on a sample of past outcomes, participants had to decide whether the majority of all past outcomes consisted of gains or of losses, that is, whether the dominant outcome of the slot machine were gains or losses. The slot machine's probability to yield gains (green upward-pointing arrows) could take values from .1 to .9 (in steps of .1, with the exception of .5). Making the correct choice was rewarded with the addition of points to one's balance, while an incorrect choice incurred an equivalent loss of points. To obtain information, each individual arrow had to be bought at the information cost illustrated above. As in previous experiments, larger random samples provided a better chance of making the right decision and receiving the reward instead of the loss. Yet, a larger sample also incurred higher information costs. Hence, there was a clear trade-off between the benefits and costs of information (higher expected payoff for correct choices, but more expensive sample costs), measurable in the same monetary unit at the level of each individual trial.

This trade-off depended on the ratio of information cost to potential payoff (gain/loss): roughly speaking, the cheaper the information was relative to the potential gain, the more information tended to be worthwhile to pay for, with the increase in the expected accuracy outweighing the direct cost of the information. When the payoff was low and the information expensive, the cost of information more quickly started to outweigh the additional benefit. Different ratios (of information cost to payoff) called for different strategies (in terms of sample size) in different blocks. Thus, the inclusion of a wider range of ratios in the longitudinal design allowed us to see whether the strong oversampling bias obtained in the aforementioned experiments generalizes across a wider range of task conditions and, in particular, whether the oversampling bias disappears when speed (information cost) and accuracy are comparable in evaluability.

For an empirical answer to both questions, we analysed the size of participants' samples as main dependent measure for the sequence of tasks in Figure 3. As a benchmark to evaluate participants' performance, we specified a Bayesian model for each cost ratio. The model consisted of an iteratively determined threshold that maximised the payoff. The threshold was the posterior odds ratio that the evidence in a sample had to exceed in favour of one of the two hypotheses(i.e., $p < .5$ or $p >.5$) before sampling was to be truncated and a decision made. Assessing individual participants' performance and sampling strategy relative to these benchmarks of optimal performance allows for comparisons across participants, experiments, and blocks.

A glance at Figure 4 shows that the average participant was indeed somewhat sensitive to the manipulated changes in cost ratios. From the smallest information-cost ratio on the left (1:80) to the highest ratio on the right (1:10), the average sample size decreases slightly but monotonically from 4.49 (for 1:80), when sampling was least expensive, to 3.24 (for 1:10), when information was most expensive (see light grey squares indicating average *n*). However, although the decline in average *n* with increasing cost ratios may be considered genuine evidence for some adaptive sensitivity, the extent of this sensitivity is conspicuously small. The mean of the individual correlation coefficients for the relation between sample size and ratio is $M_r = -.24$ for Experiment 1(based on 73 individuals who completed one to two blocks of 25 decisions per ratio) and $M_r = -0.17$ for Experiment 2 (based on 84 individuals who completed one block of 30 decisions per ratio).[24] Both means are reliably different from 0, but the variance accounted for by the ratio is surprisingly small, considering that the ratio should be the main systematic influences on the sample size for the optimal strategy.

While these results testify to the basic possibility that participants at least understand the general demands implied by the cost ratios (i.e., that lower information costs allow for larger samples than higher costs), it is obvious at a quantitative level that this basic

[24] The standard deviation of the individual correlation coefficients was $SD_r = 0.21$ for Experiment 1 and $SD_r = 0.23$ for Experiment 2. T-tests showed that both means differed from 0; for Experiment 1 $t(72) = -10.08$, $p < .001$; for Experiment 2 $t(83) = -7.03$, $p < .001$.

sensitivity is not sufficient for effective regulation of cost-benefit trade-offs. The sensitivity to the ratio of information cost to payoff, that is, the decline of light-grey squares from left to right in Figure 4, is by magnitudes smaller than the decline of the black points illustrating an optimal agent's sensitivity to cost ratios.

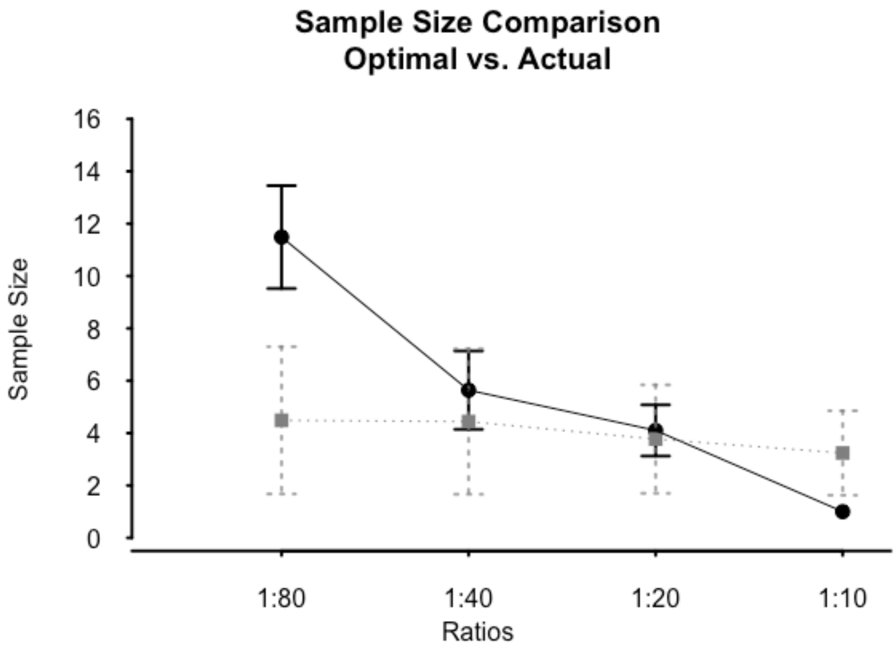

Figure 4. Graph indicating participants' mean sample sizes (grey squares with dashed bars) and corresponding mean optimal sample size (black points) for each cost ratio. Bars indicate the standard deviations.

Crucially, however, a comparison of the black and light-grey bars in Figure 4 entails clear-cut answers to the twofold theoretical question that guided this investigation. The crossover pattern for the black and light-grey points highlights the insight gained from Evans et al. (2019) that comparisons with an optimal strategy are determined more by the parameter-specific optimal strategy, which can vary greatly, than by the participants' behaviour, which only shifts mildly.

In other words, as robust as findings such as our oversampling bias in speed-accuracy trade-offs may seem, any seemingly general result based on a single (or narrow) set of parameters, may only be unidirectional simply because the parameters yield an extreme optimal strategy as comparison standard. Different parameters, implying a different comparison standard, may lead to very different conclusions.

Thus, rather than corroborating a general oversampling bias, Figure 4 shows a clear pattern of deviations from the optimal strategy in both directions that could be considered regressive: Whatever strategy is called for is apparently followed insufficiently, seemingly regressing to the mean or to some task-specific default that participants may be using. Participants sample too much information when information is expensive (ratio

1:10) and small samples are called for. They sample approximately the optimal amount when information costs are moderate (ratios 1:20 and 1:40), and finally, they sample too little when the information is the cheapest and call for large samples. If we had included only a subset of those ratios in our experiments, we might have been convinced that we have evidence for any one of those conclusions. That conclusion, however, would have been arbitrary, determined by our choice of parameters rather than participants' behaviour. In any case, just as Jarvstad et al. (2012) and Evans et al. (2019) pointed out, optimality depends on the task and its parameters. Hence, comparisons with optimality should always be interpreted carefully, and are only of limited meaning in and of themselves.

With regard to evaluability, in particular, the regressive pattern of Figure 3 lends support to the notion that obtaining a one-sided oversampling bias in previous research was at least partially due to unequal evaluability. Quantifying the negative payoff to be paid as information costs in the same currency as the positive payoff for a correct choice, and making this ratio jointly visible on every trial sensitized participants for the relative advantage of both speed and accuracy and presumably supported the two-sided (regressive) pattern.

7 Concluding remarks

Thus, as a result of various experiments and simulation studies we have conducted on information cost-benefit trade-offs in sample-based decision making, we have learned that the joint regulation of costs and benefits, or speed and accuracy, is an intricate, highly demanding metacognitive task. Mapping speed and accuracy onto the same quantitative index, sample size, highlights the different relations that speed and accuracy have with sample size, with speed decreasing more rapidly with increasing sample size than accuracy increases. Although responses to the post-experimental questionnaires indicate that people seem to understand that they have to weigh speed against accuracy, they fail to appreciate their specific relations with sample size and especially the fact that those relations differ. Hence, they cannot translate their basic understanding of the trade-off into a quantitative sampling strategy that maximizes the product of the number of complete choices (speed) and the accuracy of the expected correctness of the average choice (accuracy) in a given period of time.

However, more thorough tests of the range of possible task parameters reveal that the unequal relation of sample size to speed (or information costs) and accuracy does not necessarily imply an optimal strategy that emphasizes speed (or very low expenditure on information costs) over accuracy. Rather, the optimal strategy depends substantially on the specific task parameters and so does the classification of participants' strategies as oversampling, undersampling or 'about right'. Hence, the most important superordinate insight to be gained from the reported research is that comparisons with optimality should be interpreted with great care and are only meaningful when they are embedded in a broader framework. Our suggestion for such a framework is the study of adaptivity. Assessing humans' adaptivity to their environments and changes thereof can be aided greatly by specifying an optimal strategy for comparison. Importantly, deviations from optimality would not automatically be bad or irrational. Their interpretation would depend on what they imply for adaptivity. In our financial costs experiments, the comparison with the changes required by an optimal strategy helps to quantify the extent of the participants' change in response to the changing task parameters and demonstrates how small those changes are. The deviations are not meaningful themselves, but they become meaningful in the context of what one expects from adaptive behaviour.

Why people's adaptivity seems rather limited in this task warrants further investigation, since changing behaviour in response to changes in one's environment and task is a cornerstone of adaptive behaviour. Just as research into covertly dynamic environments (e.g., Navarro et al. 2016) is fascinating and important, because many aspects of the world are far from stable, research into responses to overt and obvious changes in the environment complements and completes the study of human adaptivity. An additional aspect of interest in the latter category is that humans might not have to be able to adapt to the changes without assistance. They need the metacognitive insight to recognise that they are in a situation that requires adaptation and that they might not perform well on their own, which would enable them to use technological tools or ask for (expert) advice. Whether or not they are able to recognise their need for help is a worthwhile question for further research – the investigation of which has already begun (Fiedler, Prager, McCaughey, submitted).

REFERENCES

Balci, F., Simen, P., Niyogi, R., Saxe, A., Hughes, J. A., Holmes, P., & Cohen, J. D. (2011). Acquisition of decision making criteria: Reward rate ultimately beats accuracy. *Attention, Perception, and Psychophysics*, *73*(2), 640–657. https://doi.org/10.3758/s13414-010-0049-7

Connolly, T., & Serre, P. (1984). Information search in judgment tasks: The effects of unequal cue validity and cost. *Organizational Behavior and Human Performance*, *34*(3), 387–401. https://doi.org/10.1016/0030-5073(84)90045-X

Denrell, J. & Le Mens, G. (2022) Chapter in this volume.Denrell, J., & March, J. G. (2001). Adaptation as Information Restriction: The Hot Stove Effect. *Organization Science*, *12*(5), 523–538. https://doi.org/10.1287/orsc.12.5.523.10092

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

Harris, C.A. (2022) Chapter in this volume.

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. https://doi.org/10.1016/0022-2496(65)90007-6

Evans, N. J., Bennett, A. J., & Brown, S. D. (2019). Optimal or not; depends on the task. *Psychonomic Bulletin and Review*, *26*(3), 1027–1034. https://doi.org/10.3758/s13423-018-1536-4

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

Fiedler, K., McCaughey, L., Prager, J., Eichberger, J., & Schnell, K. (2020). Speed-accuracy trade-offs in sample-based decisions. *Journal of Experimental Psychology: General*. https://doi.org/10.1037/xge0000986

Fiedler, K., Prager, J., McCaughey, L. Metacognitive Myopia – A Major Obstacle on the Way to Rationality. *Current Directions in Psychological Science.* Under review.

Fischhoff, B., Slovic, P., & Lichtenstein, S. (1979). Subjective sensitivity analysis. *Organizational Behavior and Human Performance, 23*(3), 339–359. https://doi.org/10.1016/0030-5073(79)90002-3

Fried, L. S., & Peterson, C. R. (1969). Information seeking: Optional versus fixed stopping. *Journal of Experimental Psychology*, *80*(3 PART 1), 525–529. https://doi.org/10.1037/h0027484

Gigerenzer, G. (2004). Striking a Blow for Sanity in Theories of Rationality. In M. Augier & J. G. March (Eds.), *Models of a man: Essays in memory of Herbert A. Simon* (pp. 389–409).

Gigerenzer, G. & Goldstein, D. G. (1996). "Reasoning the fast and frugal way: Models of bounded rationality". Psychological Review, 103, 650-669

Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: the recognition heuristic. *Psychological Review*, *109*(1), 75.

Harris, C.A. (this volume).

Hausfeld, J. & Resnjanskij, S. (2018) Risky Decisions and the Opportunity Costs of Time, ifo Working Paper No. 269 ifo Institute, Munich

Hershman, R. L., & Levine, J. R. (1970). Deviations from optimum information-purchase strategies in human decision-making. *Organizational Behavior and Human Performance, 5*(4), 313–329. https://doi.org/10.1016/0030-5073(70)90023-1

Hsee, C. K., & Zhang, J. (2010). General evaluability theory. *Perspectives on Psychological Science*, *5*(4), 343–355. https://doi.org/10.1177/1745691610374586

Jarvstad, A., Rushton, S. K., Warren, P. A., & Hahn, U. (2012). Knowing When to Move On: Cognitive and Perceptual Decisions in Time. *Psychological Science*, *23*(6), 589–597. https://doi.org/10.1177/0956797611426579

Larrick, R. P., & Soll, J. B. (2008). The MPG illusion. *Science, 320,* 1593–1594. http://dx.doi.org/10.1126/science.1154983

McCaughey, L., Prager, J., Fiedler, K. (2021). Rivals reloaded: Rivals reloaded: Adapting to competitive pressure in sample-based speed-accuracy trade-offs. *Manuscript in preparation*.

McCaughey, L., Prager, J., Fiedler, K. (2021). Adapting information search to changing task parameters. *Manuscript in preparation*.

Madan, C. R., Spetch, M. L., & Ludvig, E. A. (2015). Rapid makes risky: Time pressure increases risk seeking in decisions from experience. *Journal of Cognitive Psychology*, *5911*(June), 1–8. https://doi.org/10.1080/20445911.2015.1055274

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

Payne, J. W., Bettman, J. R., & Luce, M. F. (1996). When Time Is Money: Decision Behavior under Opportunity-Cost Time Pressure. *Organizational Behavior and Human Decision Processes*, *66*(2), 131–152. https://doi.org/10.1006/obhd.1996.0044

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

Rieskamp, J., & Hoffrage, U. (2008). Inferences under time pressure: How opportunity costs affect strategy selection. *Acta Psychologica*, *127*(2), 258–276. https://doi.org/10.1016/j.actpsy.2007.05.004

Sedlmeier, P., Hertwig, R., & Gigerenzer, G. (1998). Are judgments of the positional frequencies of letters systematically biased due to availability? *Journal of Experimental Psychology: Learning Memory and Cognition*, *24*(3), 754–770. https://doi.org/10.1037/0278-7393.24.3.754

Slovic, P., & Lichtenstein, S. (1971). Comparison of Bayesian and regression approaches to the study of information processing in judgment. *Organizational Behavior and Human Performance*, *6*(6), 649–744. https://doi.org/10.1016/0030-5073(71)90033-X

Snapper, K. J., & Peterson, C. R. (1971). Information seeking and data diagnosticity. *Journal of Experimental Psychology, 87*(3), 429–433. https://doi.org/10.1037/h0030557

Svenson, O., & Eriksson, G. (2017). Mental models of driving and speed: Biases, choices and reality. *Transport Reviews, 37,* 653–666. http://dx .doi.org/10.1080/01441647.2017.1289278

Todd, P. M., & Gigerenzer, G. (2007). Environments that make us smart: Ecological rationality. *Current Directions in Psychological Science*, *16*(3), 167–171. https://doi.org/10.1111/j.1467-8721.2007.00497.x

APPENDIX 4: Fourth paper

Adapting to information search costs in sample-based decisions

Linda McCaughey*, Johannes Prager+, & Klaus Fiedler*

*Heidelberg University +Ludwig-Maximilians-University Munich

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 search efficiently. Using a sampling-based decision paradigm with financial costs for each observation, we conducted four experiments to investigate whether and how participants adapt their information amount to the relevant aspects of the environment. Experiment 1 (N = 78) and 2 (N = 84) focussed on adaptation to different ratios of information cost and decision outcomes, revealing that participants are capable of adapting, but seem to do so mainly based on a-priori planning. Experiment 3 (N = 147) assessed the role of a second, fine-tuning process by varying whether participants received feedback or not. Surprisingly, this did not make a difference. To test whether this was simply due to unsuitable feedback or a lack of time, Experiment 4 (N = 176) implemented different feedback types and more trials, but participants still did not improve their efficiency over time. This provides crucial insights into how participants adapt and what difficulties might befall the necessary metacognitive monitoring and serves as prompt for valuable future research.

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

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 the financial resources one has to invest to gain it, the time one has to spend attaining it, or the mental effort required for its search or subsequent processing. Hence, as beneficial as additional information search is, it cannot be pursued without limit. 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 now? Do the costs of information outweigh its benefits or vice versa? If additional information only results in a small increase in probability of choosing the best option, and the best option is only a little better than the alternatives, than even low information costs may be too high a price to pay. And while it is rarely important to achieve the perfect trade-off in real life, being able to use one's resources efficiently is considered central to adaptive decision making.

Whether people can change their behavior to match changing characteristics of the decision situation such that efficient resource use is achieved, and what processes are involved, is the focus of this paper. We designed four experiments that investigated how people respond to different information costs in simple decisions based on samples of observations. The cost of each observation and the payoff for correct and incorrect decisions were explicitly stated and varied across different blocks. Across the experiments, the type and amount of feedback also varied to investigate the processes involved in participants' adaptation. The results paint a mixed picture of adaptivity, which will be discussed at length. Before we can do so, we must introduce the concept of adaptivity.

The core of adaptivity is clearest in contrast with what is usually termed rationality. While behavior has to be aligned with certain norms such as probability theory or principles of logic to count as rational, the central characteristic of adaptive behavior is its fit with the environment it is displayed in, according to Todd and Gigerenzer (2007). What we believe this 'fit' is grounded in, is the efficiency of resource use, similar to what Lieder and Griffiths (2017) call resource rationality. We understand a good 'fit' to mean sensitivity to the key characteristics of a decision situation in order to achieve a good outcome without expending too many resources on it. This is closely related to Simon's concept of bounded rationality (1956). An appropriate standard for humans needs to take into consideration the cognitive and epistemic limits and time constraints. According to this standard, humans should achieve as good an outcome as possible with as few resources as possible. Heuristic decision strategies are the prime example (e.g., Gigerenzer, Todd, & the ABC Research Group, 1999). They reduce the amount of information and hence the money, time and effort required to arrive at a decision and still achieve a very high accuracy. Trying to improve the accuracy any further would require such a great amount of resources that it is only rarely worthwhile and may even be impossible. Lieder and Griffiths (2017) found evidence that people learn which heuristic decision strategy to apply in a given situation based on their experience of the costs and benefits of the strategy in that type of situation. Even one of the most persistent biases, the anchoring bias, can be interpreted as a successful balance between the resources required to further improve an estimate and the benefit that this improvement would bring, according to Lieder and colleagues (2017).

Overall, heuristics research indicates that humans are adaptive decision makers. They achieve this by selecting the right heuristic for the environment (for example from the adaptive toolbox, Gigerenzer, Todd, & The ABC Research Group, 1999), which leads to the processing of a certain amount of information before deciding. Some accounts assume that the amount of information or evidence desired in a situation guides the selection of the strategy (Lieder & Griffiths, Lee & Cummins, 2014; Söllner & Bröder, 2015). Even so, the amount of evidence never needs to be decided on during information search. This is the case for a different type of decision, which we will call sample-based decisions and which has been argued to be very common in every-day life (Hertwig et al., 2004).

The cognitive-ecological approach (Fiedler & Wänke, 2009) assumes that, similar to sampling from a population, one needs to collect individual observations as proxies for many variables of interest, because we cannot observe them directly. Assessing the health risk entailed by a certain diet or lifestyle, or how honest and reliable a new colleague is are examples of such situations. In such sample-based decision, one has to decide about the information amount more directly. Hence, an adaptive decision maker not only needs to be able to choose a suitable heuristic depending on the environment and its information structure, they also need to search for a suitable amount of information when faced with a type of decision that calls for collecting observation in a sampling-like manner. As mentioned above, a suitable amount of information is an amount that efficiently uses the resources depending on the decision outcomes, their distribution, and the cost of information.

Experiments investigating to what extent people's reasoning and decision making resembled the Bayesian ideal (see Slovic & Lichtenstein, 1971, for a review) give a first indication of people's adaptivity in sample-based decisions. Various studies demonstrated that participants were sensitive to the different aspects of decisions tasks that influence the trade-off between benefits and costs of information, not least the information cost and the payoff for correct and incorrect decision (Pitz, 1968; Pitz et al., 1969; Snapper & Peterson, 1971; Fried & Peterson, 1969; Irwin & Smith, 1957). However, they did not provide a satisfactory explanation for how this comes about, since the findings also demonstrated that people are not well described by Bayesian updating. Hence, to what extent and especially how people are able to display behavior that is sensitive to the relevant characteristics and especially how are questions worthy of our attention.

Fascinating work on how people adapt to parameters of a trade-off has recently been conducted in a related paradigm – the perceptual speed-accuracy trade-off, in which participants have to make as many perceptual decisions as they can as accurately as they can within a certain time limit, because correct decisions are rewarded. Increasing accuracy by spending more time on the individual decisions leads to a decrease in speed in terms of the number of decisions that will be made in total, and vice versa, making it a speed-accuracy trade-off. Instead of financial information costs, information incurs a time cost and needs to be traded off against the benefits of information to achieve efficient use of resources.

Experiments using this paradigm have most frequently revealed an exaggerated focus on accuracy (e.g., Starns & Ratcliff, 2010 and 2012, but see Evans, Bennet, & Brown, 2019), meaning that participants spent too much time on individual decisions, foregoing payoff in the process. Some researchers turned their focus to investigating whether participants could overcome this suboptimal focus to use resources more efficiently (Simen et al., 2009; Balci et al., 2009). Most interestingly, Evans and Brown (2017) investigated the effects of different types of feedback. They demonstrated that participants could overcome their accuracy focus with enough time and adequate feedback, with more extensive feedback leading to more efficient performance, especially when it included explicit advice on how to improve one's efficiency. However, what processes are involved and what their requirements are has not been studied, although Khodadadi et al.'s theoretical simulation (2014) offers an interesting perspective.

Since moving between different contexts and encountering new decision situations is a crucial aspect of real-world decision making, we believe that it is of particular relevance to understand what happens when people are confronted with changes in the decision environment. An adaptive decision maker needs to be able to enter a new context and at least broadly respond to it almost immediately and then perhaps use additional experience to further improve decision outcomes.

We assume that at least two processes are at work in adapting information search to characteristics of the decision situation. First, we believe that people can make a-priori plans for an adjustment based on what they know or are told about a new situation before beginning to search in the environment. This enables them to immediately respond to a different environment. If they are familiar with the environment, this may suffice to lead to very efficient resource use, similar to when participants apply a suitable heuristic to a decision situation. However, when an environment is not familiar, presumably a secondary process additionally comes into play, which we want to call fine-tuning: while engaging with the environment, searching for information and making decisions, people should be able to improve their efficiency based on the experience they are collecting with the environment. For this process, the metacognitive perspective is very valuable (Fiedler, Ackermann, Scarampi, 2019). Metacognition refers to cognition about cognition, and is often divided into the two main functions of monitoring and control (Nelson & Narens, 1990). In a new environment, one needs to monitor one's information search strategy and its outcomes and evaluate it. The evaluation of one or ideally multiple strategies informs the control function which strategy to execute, ideally resulting in an improvement in efficiency with experience. The task structure and feedback might not always make this type of efficiency improvement possible. In earlier work with a sample-based speed-accuracy trade-off, our participants exhibited a marked accuracy bias that they failed to overcome despite multiple interventions, presumably because of an imbalance between how salient and evaluable speed and accuracy were due to the trade-off structure. However, financial-cost tasks do not suffer from this imbalance and therefore provide a good task for investigating the proposed processes in relation to information search. Also, since we wanted to investigate both processes, an experimental task representative of the real world would not be ideal, since participants may have already developed adaptive strategies for it. Instead, a somewhat 'artificial' task (with an arbitrary information structure) was more suited for observing both potential processes and their interplay.

Paradigm and task

Our paradigm required participants to make a decision about a parameter of a population from which they could acquire as many observations as they thought was sensible. In Experiments 1 and 2, participants were told that they had to decide whether a slot machine tended to yield gains or losses. The individual observations were defined as past gains and losses drawn at random from the totality of past recorded gains and losses. In Experiment 3 and 4, participants were told that they had to decide whether products had received predominantly positive or negative reviews on the whole. The observations were individual reviews drawn at random from the total population of reviews.

The observations came at a certain cost, which will be referred to as information cost in the following, while the decision resulted in the gain or loss of a certain sum, which will be called the payoff. The screen participants saw indicated the trial number and block number out of the total trials and blocks on the top right, and the information cost and payoff on the top left. In most conditions, it also displayed participants' current balance on the left. Participants could acquire information by pressing the space bar, which made an arrow appear that was either green and upward-pointing to indicate a positive observation or magenta and downward-pointing to indicate a negative observation. As soon as they wanted to make a decision, they could press the upward arrow key or the downward arrow key to indicate their decision that the population was predominantly positive or negative, respectively. For Experiment 3 and 4, instead of indicating decisions through key presses, participants saw two buttons on the screen below the arrows that represented the two decision options. The maximal amount of arrows, that is, the maximal sample size varied across experiments for better generalizability.

The arrows displayed were drawn from a Bernoulli distribution with different ranges of values for the probability parameter p for a positive arrow. This was called the population mean in the instructions to make it easier to understand for the participants. For Experiments 1 and 2, p could vary from .1 to .9 in steps of .1, excluding .5, and was drawn from those values at random for each new trial. For Experiments 3 and 4, p could take the values .3 and .7.

Optimal Strategy

To be able to assess to what extent participants' responses were adaptive, we also specified an optimal strategy that represented a highly efficient use of resources. It was determined by calculating the Bayesian posterior odds ratio for each possible sample size for a dataset of 10000 samples for each probability value the sampling distribution could have. Through iterative approximation based on the resulting balance (i.e., taking into account both information costs and payoffs), we determined the posterior odds ratio threshold that yielded the maximal payoff. We did this separately for each information-cost-payoff ratio. Using these thresholds, we generated an optimal-strategy-equivalent for each participant's actual samples, resulting in an optimal sample size, optimal accuracy, and optimal balance tailored to each participant.

Experiment 1

Hypotheses

To establish whether participants were sensitive to changes in the information cost and payoff parameters and to what extent, in Experiment 1, we varied either the information cost or the payoff across four blocks. We expected participants to be sensitive to the ratio of information costs and payoffs and change their behavior in response to the change based on a mixture of the two proposed processes. By behavior, we mainly mean sample size since that is the variable that participants can influence directly. Usually, changes in sample size also lead to changes in accuracy, but accuracy is also influenced by other factors, for example practice and attention.[25] Hence, sample size is the most important and reliable indicator of strategy change. How sample size and accuracy changes come together to result in the balance of points is also an interesting indicator of adaptivity, but again, strategy changes will be reflected less reliably.

[25] Accuracy will change with sample size when sample size changes reflect a change in evidence threshold. When sample size differences merely reflect the stochastic variation until the same evidence threshold is met, the relation should be exactly 0. Accuracy is also influenced by whether participants use an evidence-based strategy versus other strategies consistently – an increase in sample size due to a change from an evidence-based to a fixed-n strategy will have an unclear relation with accuracy. Finally, if participants gain familiarity with task, their accuracy might improve without a corresponding change in sample size, while a lapse of attention or lack of focus might lead to decreased accuracy without a clear relation with sample size. All in all, accuracy is not as reliable an indicator of strategy change as sample size.

While we expected participants to be sensitive to changes, we did not expect them to perform closely to the optimal strategy. Nevertheless, we did expect basic adaptivity from a majority of participants, as indicated by changes in sample size that track the changes of the optimal strategy. We expect qualitative alignment with the optimal strategy, while the quantitative accordance is secondary for adaptivity in an unfamiliar task.

Methods

Participants. We recruited 78 participants from our local participant pool, of which 76% of which were female, 23% male and one participant preferred not to say. They were aged 19 to 77 with a mean of 26.3 (SD = 11.6), with a vast majority of students (6 non-students, 14% psychology students). The experiment was part of a one-hour session in our laboratory and was conducted as 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 were remunerated with a variable amount dependent on their performance at the experiment at hand. On average, this variable sum amounted to 3€, but could be as high as 8€, resulting in a maximum payoff of 13€ for the whole session.

Design and Procedure. The experiment had a within-participant design with three different information-cost-payoff ratios as within-participant factor. To avoid order effects and effects of the concrete amount of points that instantiated the ratios, there were two conditions that differed in the order of the ratios and in the specific amount of points that were used to implement the ratios.

Participants proceeded through four blocks, each of which had different cost and payoff parameters that remained constant throughout the 25 decisions in each block. Figure 1 shows the order of the three different ratios and what specific information costs and payoffs were used to instantiate them. It also illustrates that the ratio changes were implemented by changing one ratio component and keeping the other constant. The decisions that participants had to make each trial were exactly as described above, with a maximum sample size of 40. The experiment was programmed in Python and Pygame.

Figure 1

Illustration of ratio sequences for Condition 1 and 2 for Blocks 1 to 4.

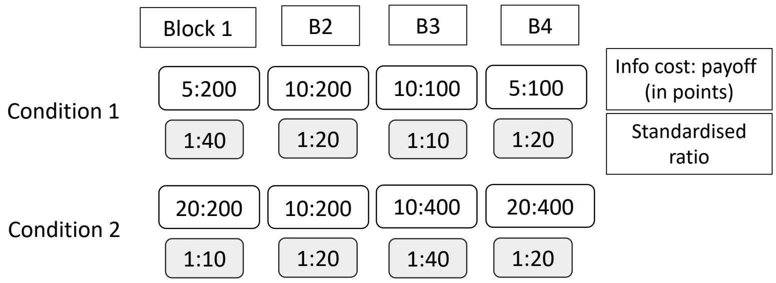

Note. Boxes with white background show the ratio of information cost and payoff instantiated by the amount of points used in the experiment, boxes with grey background show the standardized ratio.

Results

Exclusions. 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 observations for a maximum of ¾ 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.

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. Efficiency was calculated by equating the different information costs and payoffs with arbitrary ones, 10 and 100, and recalculating the balance on that basis to make all trials comparable for the purpose of demonstrating the trade-off. The accuracy plot in the left panel of Figure 1 clearly illustrates the steep initial increase in accuracy with the first few observations and the increasingly flat increase with additional observations. 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. Since the relations are clearly not linear, correlation coefficients would not be meaningful.

Figure 2

Participants' mean accuracy and efficiency as a function of sample size

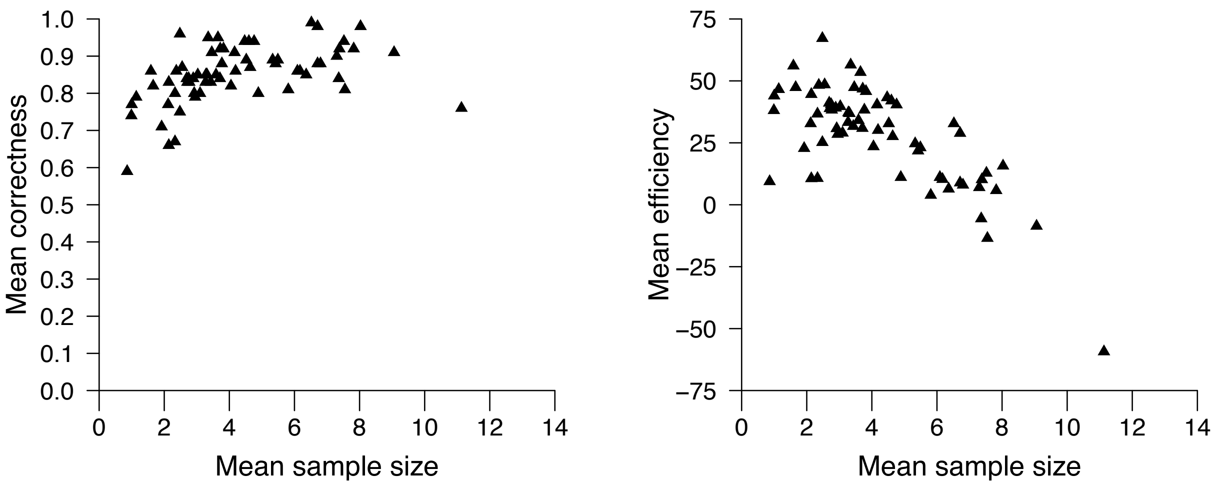

Note. Accuracy, efficiency, and sample size were averaged across all blocks, resulting in one value each for each participant.

Influence of ratio. The most interesting question is whether and to what extent participants changed their behavior in response to the announced block changes of either a different information cost or payoff. Table 1 shows the means and standard deviations of sample size, accuracy, and balance separately for blocks and conditions. For the analysis, we chose to examine the influence of ratio. We estimated a linear mixed model with sample size as dependent variable, ratio, trial, and their interaction as fixed effects and a random intercepts that varied by participants and random slopes for ratio and trial that varied by participant (using the lme4 package Version 1.1-31; Bates et al., 2015). Ratio 1:10 was coded as 1, ratio 1:20 as 2 and ratio 1:40 as 3 because we assumed that each halving of the ratio would be perceived equivalently, especially since each step was encountered in a block change. The interaction term did not contribute to model fit, as a likelihood ratio test comparing the full model to one without the interaction showed ($\chi^2(1) = 1.0137$ p = 0.314). A model that included ratio and trial as fixed effects did not fit the data better than a model with just ratio ($\chi^2(1) = 0.15$, $p = .704$), while the other way around, a model that included ratio and trial yielded better model fit than just trial ($\chi^2(1) = 33.84$, $p < .00$)1. Hence, participants indeed responded to the ratio changes by implementing larger sample sizes when information cost was cheaper or payoff was higher, and vice versa, with sample size increasing by 0.84 with every coded ratio increase ($\beta = 0.84$, $SE = 0.13$, $t = 6.56$). The sample size did not systematically change across trial numbers for any of the ratios. Figure 3 shows the mean sample size per ratio.

Table 1

Descriptive statistics by condition and block

Condition	**Block**	**Ratio**	**Points**	M_n (*SD*)	M_a (*SD*)	M_{bal} (*SD*)
1	*1*	*1:40*	5:200	4.93 (2.91)	.87 (.10)	122.79 (31.80)
	2	*1:20*	10:200	4.14 (2.38)	.86 (.10)	102.33 (29.76)
	3	*1:10*	10:100	3.23 (1.67)	.83 (.08)	33.41 (14.49)
	4	*1:20*	5:200	3.73(2.10)	.85 (.08)	51.42 (11.78)
2	*1*	*1:10*	20:200	3.65 (1.33)	.83 (.09)	57.48 (42.58)
	2	*1:20*	10:200	4.12 (2.05)	.85 (.07)	97.84 (34.70)
	3	*1:40*	10:400	5.31 (3.33)	.87 (.07)	239.95 (57.29)
	4	*1:20*	20:400	4.70 (2.38)	.86 (.08)	195.08 (66.77)

Note. Mean sample size (M_n), mean accuracy (M_a) and mean balance (M_{bal}) per condition per block with respective standard deviation (*SD*) in brackets.

Since changes in sample size should lead to accuracy changes as part of the trade-off, we also examined accuracy with a binomial generalized linear mixed model with ratio as fixed effect. We did not include trial as fixed effect because we expected the accuracy data to be too noisy to show a change across trials. Including both random intercepts and random slopes for participants led to a singular fit. Assuming that the model was overfitting the data, we decided to drop the random slope for ratio. The likelihood ratio test comparing the full model with a reduced model that did not include the fixed effect of ratio showed that including the ratio significantly improved model fit, ($\chi^2(1) = 9.34$, $p = .002$). Each increase in ratio, making information less expensive, resulted in the odds of decisions being correct increasing by 10% ($\beta = .10$, $SE = 0.03$, $z = 3.05$).

Figure 3

Optimal and Actual Sample Size and Accuracy by Ratio

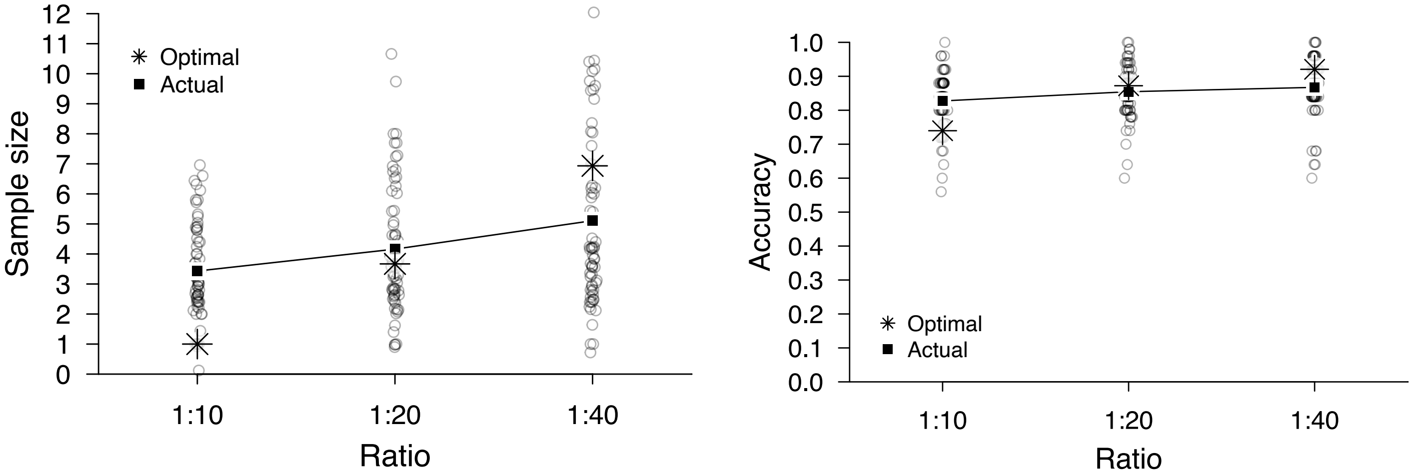

Note. Participant sample size means are depicted by black squares and mean optimal sample size by asterisks. Individual participants' values are represented by grey circles.

Analyzing the influence of ratio on balance requires a different approach since the two are necessarily confounded. Different ratios, especially if they are instantiated by different payoffs, lead to a very different balance just because the payoff is different. Even if one standardizes the balance so that the payoff is the same for all ratios and the ratio is maintained by adjusting the information cost, that difference in information cost should still substantially influence the balance across ratio.[26] The best way to assess the balance and its change across ratios is in comparison with the optimal strategy, which will be described below.

Table 2

Descriptive Statistics of Actual and Optimal Dependent Variables and their Difference

Ratio	*n opt* (*SD*)	*n* (*SD*)	*n diff* (*SD*)	*a opt* (*SD*)	*a* (*SD*)	*a diff* (*SD*)	*bal opt* (*SD*)	*bal* (*SD*)	*bal diff* (*SD*)
1:10	1 (0)	3.43 (1.54)	-2.43 (1.54)	.750 (.098)	.827 (.085)	-.088 (.113)	55.4 (34.1)	44.9 (33.3)	10.5 (35.2)
1:20	3.67 (0.66)	4.16 (2.17)	-0.49 (2.19)	.872 (.063)	.855 (.078)	.018 (.082)	122.5 (51.6)	110.1 (50.9)	12.4 (24.3)
1:40	6.93 (1.19)	5.11 (3.10)	1.83 (3.21)	.921 (.047)	.867 (.088)	.053 (.092)	195.5 (70.9)	178.7 (74.43)	16.8 (43.3)

Note. Mean optimal and actual sample size and their difference (*n*, *n opt,* and *n diff*), mean optimal and actual accuracy and their difference (*a opt*, *a*, and *a diff*), and optimal and actual mean balance and their difference (*bal opt*, *bal*, and *bal diff*), with respective standard deviation (*SD*) in brackets, separately for ratios.

Optimality comparison. Table 2 displays the means and standard deviations of actual and optimal sample size, accuracy, balance and standardized balance and the respective differences

[26] Fixing all payoffs and information cost to an arbitrary value as done for efficiency above also distorts the results, since it ignores all strategy differences resulting from the different ratios.

between optimal and actual measures. To assess the difference between optimal sample size and participants' sample size, we constructed a dummy variable 'optimality' to indicate whether a sample size was optimal or participants' actual sample size (coded 0 for actual and 1 for optimal) and fitted a linear mixed model with sample size as dependent variable and optimality, ratio, and their interaction as fixed effects. We also included a random intercept and a random slope for optimality that varied by participant. A model that also included a random slope for ratio yielded a singular fit, so we decided to drop it.

The full model provided a better fit than a reduced model that only included ratio as fixed effect, as shown by a likelihood ratio test ($\chi^2(2) = 224.7$, $p < .001$). Compared to the actual sample size, the optimal sample size was smaller by 3.63 assuming a ratio of 0 ($\beta = -3.63$, $SE = 0.35$, $t = -10.36$). For actual sample size (i.e. when optimality = 0), every increase in ratio brought an increase of 0.84 in sample size ($\beta = 0.84$, $SE = 0.08$, $t = 11.22$). However, the interaction between optimality and ratio showed that for optimal sample size, every increment in ratio brought an additional 1.58 increment ($\beta = 1.58$, $SE = 0.11$, $t = 14.96$), meaning that for the ratio coded 1 (1:10), the optimal sample size was estimated to be about 1.5 observations lower than the actual sample size, but about equal for ratio 2 (1:20), and higher by about 1.5 observations for ratio 3 (1:40).

To see whether the differences in sample size also translated into differences in accuracy, we fitted a binomial generalized linear mixed model with accuracy as dependent variable, the same dummy-coded optimality and ratio as fixed effects, and a random intercept that varied by participant as well as random slopes for ratio and optimality that likewise varied by participants. The results show that at a hypothetical ratio of 0 optimal decisions were less likely to be correct than actual decisions ($\beta = -1.10$, $SE = 0.11$, $z = 12.57$, $p < .001$). With every increase in ratio, actual decisions were more likely to be correct ($\beta = 0.90$, $SE = 0.21$, $z = 3.74$, $p < .001$). With increasing ratio, the likelihood of being correct changed differently for optimal and actual decisions, with optimal decision increasing more ($\beta = 2.26$, $SE = 0.29$, $z = 7.76$, $p < .001$).

To appropriately test for differences between optimal and actual balance, we needed to make the different scales comparable. This was achieved by standardizing the payoff to 100 and adjusting the information cost to 10, 5 and 2.5 to maintain the ratios 1:10, 1:20 and 1:40, respectively. That way the different implementation of the ratios in different points, which were sometimes confounded with ratio, would not influence the results. We fitted a linear mixed model to this standardized balance data with the dummy optimality, ratio, and their interaction as fixed effects. We also included a random intercept that varied by participant and a random slope for optimality that varied by participant. Including an additional random slope for ratio led to a failure to converge or a singular fit, presumably due to overfitting the data, so we chose not to include it. A likelihood ratio test between this full model and a reduced model without the interaction showed that including the interaction did not improve fit ($\chi^2(1) = 0.04$, $p = .841$). Comparing this first reduced model to a further reduced model that did not include the dummy optimality with a likelihood ratio test revealed that including optimality did increase the fit of the model to the data ($\chi^2(1) = 17.42$, $p < .001$). Testing the influence of ratio in the equivalent way, a likelihood ratio test between the first reduced model that included optimality and ratio as fixed effects to a further reduced model that included only optimality showed that ratio also had an influence on balance ($\chi^2(1) = 275.74$, $p < .001$). Hence, the actual balance differed from the optimal balance, but this differences was not moderated by ratio, as is visualized in Figure 4.

Figure 4

Actual and Optimal Balance by Ratio

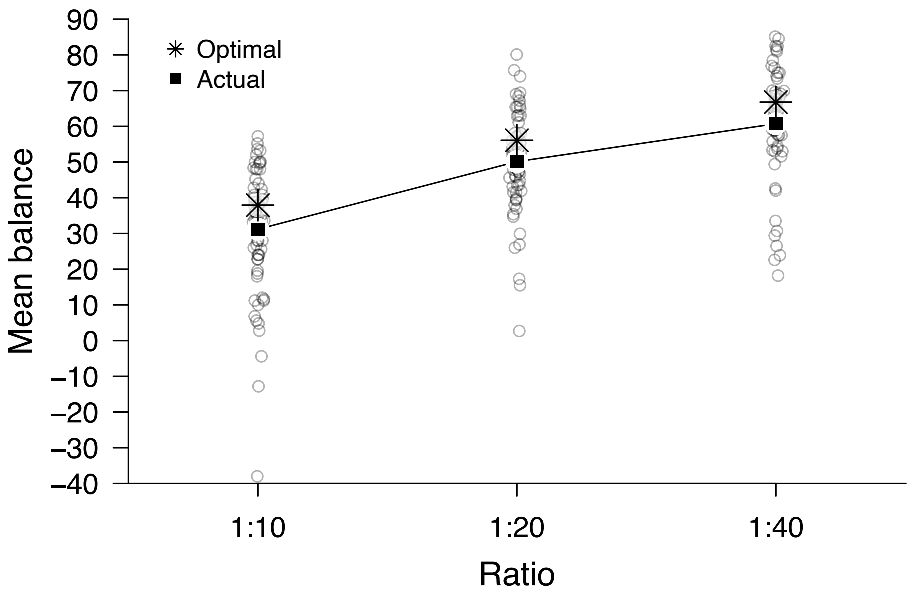

Note. Participants' mean balance is depicted by black squares and mean optimal sample size by asterisks. Individual participants' values are represented by grey circles.

Adaptivity of changes. A relevant question for adaptivity, especially the a-priori planning process that we presume to play a role, focusses on the direction of change in sample size in response to the new ratios – independently of the absolute deviation between participants' and optimal sample sizes. It is a fundamental skill to be able to increase versus decrease information search in response to the requirements of the decision situation, even when sample sizes are not of the right size in absolute terms yet.

Table 3 displays the proportion of participants who changed their sample size in line with the direction of the optimal-sample-size change, which demonstrates high sensitivity to the requirements of the task. This sensitivity is presumably made possible in large part due to the a-prior planning of strategy change. The changes in the opposite direction were categorized according to whether they approached the optimal sample size or diverged from it and show that often when participants did not follow the direction of the change of the optimal sample size, they instead came closer to the optimal strategy in absolute terms. This might reflect the secondary fine-tuning process based on the experience with the ratio, but might also be accidental.

Table 3

Proportion of sample size changes in the same direction as the optimal strategy per block

Cond.	Blocks and ratios	Adaptive *n* change	Other changes that…
1 (*N*=35)	B1 to B2 1:40 to 1:20	0.77 (27)	…decreased distance from optimal strategy: .14 (5) …increased distance from optimal strategy: .03 (1) No change: .06 (2) (n closer to *opt n* in B2 for all)
	B2 to B3 1:20 to 1:10	0.8 (28)	…decreased distance from optimal strategy: .05 (2) …increased distance from optimal strategy: .11 (4) No change: .03 (1) (n closer to *opt n* in B2 for all)
	B3 to B4 1:10 to 1:20	0.8 (28)	…decreased distance from optimal strategy: .05 (2) …increased distance from optimal strategy: .05 (2) No change: 0.09 (3) (n closer to *opt n* for B3 for all)
2 (*N*=32)	B1 to B2 1:10 to 1:20	0.63 (20)	…decreased distance from optimal strategy: .06 (2) …increased distance from optimal strategy: .28 (9) No change: 0.03 (1) (*n* closer to *opt n* for B1)
	B2 to B3 1:20 to 1:40	0.84 (27)	…decreased distance from optimal strategy: 0 (0) …increased distance from optimal strategy: .09 (3) No change: 0.06 (2) (*n* closer to *opt n* for B2)
	B3 to B4 1:40 to 1:20	0.63 (20)	…decreased distance from optimal strategy: .19 (6) …increased distance from optimal strategy: .09 (3) No change: 0.09 (3) (*n* closer to *opt n* for B2)

Note. The table displays the proportion of participants who changed their sample size in line with the direction of the optimal strategy change per block and condition (cond.), denoted adaptive *n* change. The remaining participants are categorized into changes that decrease or increase the distance from the optimal strategy, meaning that the sample size in the current block is closer to or further from the optimal sample size of the current block than the sample size in the previous block from the current block's optimal sample size. Proportions are based on all participants in the condition. Numbers in brackets indicate the number of participants.

Discussion

Our participants definitely displayed adaptivity to the different ratios, both in sample size and accuracy. Since this was intended as unfamiliar task, the significant deviation of their sample size from the optimal sample is not diagnostic for adaptivity. The high proportion of changes in the direction of the optimal sample size further support adaptivity. Since only one component of the ratio changed at a time and the blocks were only interrupted by the announcement of the new block and the new ratio, they were easy to compare. Hence, we assumed that the a-priori planning process played a large role in bringing about the strategy change that participants displayed. To test whether such a process was indeed at play, we designed an experiment in which the blocks were harder to compare, making the a-priori adjustment of strategy more difficult. This would also allow us to gain insights into the extent to which the experience-based fine-tuning process contributed to the change.

Experiment 2

Hypotheses

To examine to what extent adaptation relied on a-priori planning and to what extent the adjustment was based on the fine-tuning through experience, we implemented ratios that changed in both components between blocks and added a distractor task to make it more difficult to compare the blocks' ratios.

We assumed that participants would still be able to adapt to the different ratios, as indicated by ratio influencing sample size significantly. However, as should be evident from lower proportions of adaptive changes, we expected participants to have greater difficulty adapting, indicating that the adaptation is based on the a-priori process to a large extent. We will refrain from analyzing accuracy and balance again, as they are so heavily dependent on sample size that they do not add much insight beyond the analysis of sample size.

Method

Participants. Using our local participant pool, we recruited 84 participants to take part in an hour-long session consisting of multiple studies. Their ages ranged from 18 to 51, with a mean of 22.7 (*SD* = 4.8). About 62% were female, 37% male, and one person who preferred not to say. The vast majority were students of different subject areas (with two non-students) with the largest group being students of psychology (31%). The experiment at hand was conducted as second of three experiments. The other two experiments investigated contingency reasoning and impression formation. The participants received 5€ plus a variable amount that depended on their performance at this experiment. Psychology students could choose to receive course credit instead of the 5€, but still received the performance-dependent monetary bonus.

Design and Procedure. Participants had to work through four blocks that differed in ratio. The set of ratios was expanded to include a fourth (1:80), which resulted in four levels of the within-participant factor, with each block implementing one ratio. The order of the cost ratios differed for Condition 1 and 2. This is important to ensure that potential differences are not just due to order effects such as insufficient adjustment from the previous ratio. They are displayed in Figure 5. The changes from block to block where achieved by changing both the information cost and the payoff, meaning that neither parameter was kept constant in successive blocks. Blocks consisted of 30 decisions and the maximum sample size per decision was 40. The experimental program was based on Python and Pygame. Between blocks, participants had to solve a 2x2 sudoku, which took them about one to two minutes and was intended to be a distraction from the main task. Sudokus were chosen as easy numeric puzzles that we assumed would have a higher likelihood of interfering with the memory for the numeric ratios.

Figure 5

Illustration of ratio sequences for Condition 1 and 2, for Blocks 1 to 4.

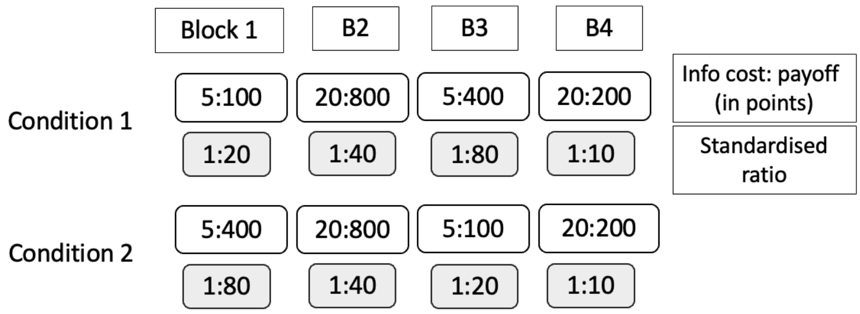

Note. Boxes with white background show the ratio of information cost and payoff instantiated by the amount of points used in the experiment. Boxes with grey background show the standardized ratio.

Results

Exclusions. Five participants had to be excluded because their accuracy rate did not differ from guessing, as indicated by binomial tests. No further participants had to be excluded due to low sample sizes, although the five already excluded participants also met this exclusions criterion. None had to be excluded based on slow response times. This left a total sample size of 79.

Change across ratios. Whether participants adapted to the ratios across the blocks is once more the most interesting question. Table 4 gives an overview of the mean sample size, accuracy, and balance for the four blocks, separately for the two conditions.

Table 4

Descriptive Statistics for Sample Size, Accuracy, and Balance by Block and Condition

Condition	Block	Ratio	Points	Mean *n* (*SD*)	Mean *a* (*SD*)	M *bal* (*SD*)
1	1	1:20	5:200	3.98 (1.46)	.838 (.098)	47.6 (17.2)
	2	1:40	20:800	4.44 (1.55)	.859 (.065)	485.2 (100.7)
	3	1:80	5:400	4.76 (1.98)	0.876 (0.062)	277.3 (47.8)
	4	1:10	20:200	3.52 (1.31)	0.841 (0.070)	66.1 (30.5)
2	1	1:80	5:400	4.78 (3.19)	0.833 (0.093)	242.1 (69.9)
	2	1:40	20:800	4.48 (2.60)	0.849 (0.081)	468.5 (112.0)
	3	1:20	5:200	3.91 (1.86)	0.833 (0.080)	46.9 (13.8)
	4	1:10	20:200	3.51 (1.42)	0.835 (0.071)	63.7 (30.2)

To analyze whether participants adapted their sample size to the ratios, we fitted a linear mixed model with sample size as dependent variable, ratio (coded from 1 to 4 to represent the four ratios 1:10 to 1:40), trial, and their interaction as fixed effects and a random intercept and random slopes for ratio and trial that varied by participant. Likelihood ratio tests showed that the interaction did not improve model fit, ($\chi^2(1) = 0.98$, $p = .323$). Ratio did improve model fit alongside trial compared to a model that included only trial as fixed effect ($\chi^2(1) = 29.98$, $p < .001$). Trial, however, did not improve model fit compared to a model with only random effects ($\chi^2(1) = 1.86$, $p = .172$). Hence, participants changed their sample size with the ratio, by an estimated 0.39 samples with each coded unit increase of ratio ($\beta = 0.39$, $SE = 0.08$, $t = 4.93$). Sample size did not change over trials at any ratio.

Optimality comparison. Table 5 summarizes means and standard deviations for all dependent variables, both actual and optimal and their difference, by ratio. Figure 6 provides a clearer overview by illustrating the mean optimal and actual sample size by ratio, with individual participants' values included as grey circles. To assess the potential influence of ratio and the difference between actual and optimal sample size, we fitted a linear mixed model to sample size as target variable and the dummy 'optimality' that indicated whether the sample size was a participant's optimal or actual sample size as fixed effect. Ratio and the interaction of optimality and ratio were added as additional fixed effects. A random intercept and a random slope for optimality that both varied by participant were added. An additional random slope for ratio led to a singular fit, so we did not include it.

Table 5

Descriptive statistics of actual and optimal dependent variables and their differences

Ratio (points)	*n opt* (*SD*)	*n* (*SD*)	*n diff* (*SD*)	*a opt* (*SD*)	*a* (*SD*)	*a diff* (*SD*)	*bal* opt (*SD*)	*bal* (*SD*)	*bal* diff (SD)
1:10 (20:200)	1 (0)	3.52 (1.36)	-2.52 (1.36)	.738 (.08)	.838 (.07)	-.100 (.09)	75.4 (32.9)	64.9 (30.2)	10.5 (34.9)
1:20 (5:100)	3.63 (0.55)	3.95 (1.67)	-0.32 (1.67)	.862 (.05)	.835 (.09)	.027 (.08)	54.2 (11.0)	47.3 (15.4)	6.9 (12.0)
1:40 (20:800)	6.78 (0.95)	4.46 (2.15)	2.37 (2.22)	.911 (.05)	.854 (.07)	.057 (.08)	522.0 (79.0)	476.5 (106.4)	45.5 (97.1)
1:80 (5:800)	10.44 (1.12)	4.78 (12.66)	5.67 (2.70)	.928 (.05)	.854 (.08)	.074 (.08)	290.1 (39.4)	259.0 (62.5)	31.1 (57.0)

Note. Means with standard deviations in brackets for optimal and actual sample size (n), accuracy (a), and balance (bal) and the respective differences (*diff*) between optimal and actual values based on data aggregated by participant and ratio.

A comparison to a reduced model that included only ratio as fixed effect showed that the inclusion of optimality and their interaction provided a better fit, as indicated by a likelihood ratio test ($\chi^2(2) = 2750.8$, $p < .001$). An additional comparison of the full model and a reduced model that did not include ratio as fixed effect showed that it also provided a better fit ($\chi^2(2) = 6647.7$, $p < .001$). Hence, the sample size difference between actual and optimal samples was about -5.8 when ratio was at hypothetical 0 ($\beta = -5.51$, $SE = 0.24$, $t = -23.08$), and increased by 2.7 with every unit increase of ratio, so that they were about equal for Ratio 2 and the optimal sample size was increasingly higher for Ratios 3 (1:40) and 4 (1:80), as shown by the interaction term ($\beta = 2.72$, $SE = 0.05$, $t = 54.07$). Participants' sample size deviated significantly from the optimal sample size, and this deviation changed with the ratio. This indicates that participants did not have a generally too low or too high sample size, but that they changed their strategy far less than the optimal strategy changed. As a consequence, participants exhibit both too small a sample size and too big a mean sample size, depending on the ratio.

Figure 6

Mean optimal and actual sample size by ratio

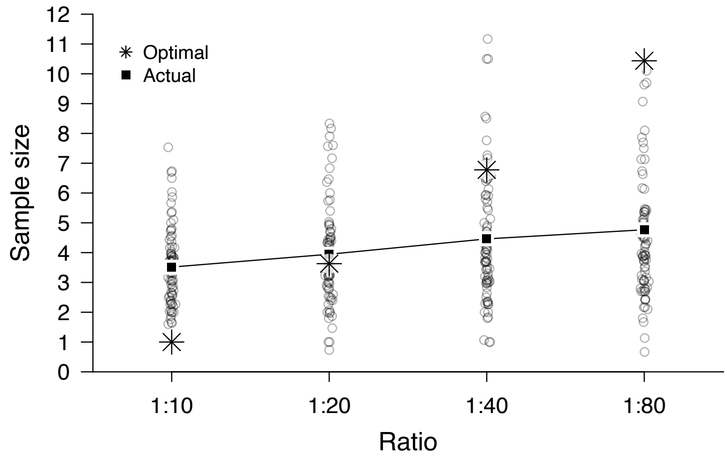

Note. Mean actual sample size is denoted by black squares and mean optimal sample size by asterisks, with individual participants represented by grey circles.

Adaptive changes. A more exact picture of the extent to which basic adaptivity was achieved is conveyed by the analysis of adaptive changes. Table 6 shows the proportion of participants who made a change in the same direction as the change in the optimal sample size. As we expected, these numbers were lower than for Experiment 1. In addition, a high proportion of the changes in the opposite direction implied changes away from the optimal sample size, so one cannot assume that the a-priori strategy change was then generally overridden by the experience with the ratio leading to a more efficient use of information and coming closer to the optimal sample size. This may have been the case in a few exceptions, one of which saw more participants change in the opposite direction to the optimal strategy, leading to greater proximity to the optimal strategy, than in line with the optimal-strategy change. Perhaps certain combination of ratios or points facilitated behavior change based on experience, but this may also have been accidental.

Table 6

Proportion of sample size changes in line with optimal per condition and block

Cond.	Blocks and ratios	Adaptive *n* change	Other changes that…
1 (*N*=38)	B1 to B2 1:20 to 1:40	.58 (22)	…decreased distance from optimal strategy: .08 (3) …increased distance from optimal strategy: .24 (9) No change: .11 (4) (closer to *opt n* in B1 for 3/4)
	B2 to B3 1:40 to 1:80	.42 (16)	…decreased distance from optimal strategy: .32 (12) …increased distance from optimal strategy: .26 (10) No change: 0

	B3 to B4 1:80 to 1:10	.53 (20)	…decreased distance from optimal strategy: 0 …increased distance from optimal strategy: .47 (18) No change: 0
2 (*N*=41)	B1 to B2 1:80 to 1:40	.59 (24)	…decreased distance from optimal strategy: .34 (14) …increased distance from optimal strategy: .07 (3) No change: 0
	B2 to B3 1:40 to 1:20	.34 (14)	…decreased distance from optimal strategy: .59 (24) …increased distance from optimal strategy: .04 (2) No change: .02 (1) (closer to *opt n* in B2)
	B3 to B4 1:20 to 1:10	.68 (28)	…decreased distance from optimal strategy: 0 …increased distance from optimal strategy: .32 (13) No change: 0

Note. The table displays the proportion of participants who changed their sample size in line with the direction of the optimal strategy change per block and condition (cond.), denoted adaptive *n* change. The remaining participants are categorized into changes that decrease or increase the distance from the optimal strategy, meaning that the sample size in the current block is closer to or further from the optimal sample size of the current block than the sample size in the previous block from the current block's optimal sample size. Proportions are based on all participants in the condition. Numbers in brackets indicate the number of participants.

Discussion

Experiment 2 showed that although the block ratios were harder to compare because both components changed at the same, participants' sample size still changed in response to the ratio change. However, the proportions of adaptive changes demonstrate that they were having greater difficulties, indicating that the a-priori planning presumably contributed to adaptive changes in Experiment 1 to a large extent. Despite greater difficulties, it was presumably still the a-priori process contributing most to adaptivity in this experiment. The complementary fine-tuning process based on the experience of the ratios should have been visible in approximations to the optimal strategy. There was no evidence that this was the case across trials, since there was no sample size change across trials. The proportions of changes indicated that it may have played a role for a few select ratio changes, but other interpretations are also possible.

Compared to the most efficient resource use of the optimal strategy, participants changed their sample size to a lesser degree. This is also in line with the explanation that the fine-tuning process was contributing less than it could. For while the a-priori planning process is well suited to effecting the qualitative change in the 'right' direction, the adjustments necessary to reach a suitable change quantitatively can only be brought about by some type of experience-based fine-tuning process. To more closely examine the contribution of this fine-tuning process, we designed another experiment that varied whether participants received feedback or not.

Experiment 3

Since the two previous experiments suggested that participants may be relying more on a-priori planning of change in response to the change announcement than on the experience of the new ratio, we designed another experiment to investigate the role played by the fine-tuning process in adaptation. Adaptation based on experiencing the ratio heavily relies on feedback, which provides information about the accuracy rate, but also makes the information cost more salient. Hence, we implemented one condition that received the usual trial-by-trial feedback, and another that only received feedback in the first of three blocks. As an additional factor, the ratio change was either implemented by changing the information-cost component or by changing the payoff component. That way, we could additionally investigate which had a greater influence on participants' strategy.

Hypotheses

We expected participants who received feedback to change their sample size more than participants who did not receive feedback, both across ratios and trials, with this greater change resulting in closer proximity to the optimal strategy. In the conditions without feedback we expect the payoff change to have a greater impact than the information cost change, since in the announcement, payoff changes seem more relevant, because the bigger of the two values is changing. In the conditions that received feedback, the influence of the information-cost change should be equal to the influence of the payoff change, since these participants experienced that the relation between the two components is more important than either of the components. Again, since accuracy and balance depend on sample size to such a large extent, we only analyzed sample size, since analyses of the other variables have little added value.

Methods

Participants. We recruited 147 participants from the study platform Prolific for an online experiment. 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 to be students, 60.1% were not, and for 31.8% data on student status was unavailable.

Design & Procedure. The experiment had a 2x2-between-participant design, with feedback (present or absent) and ratio change (via payoff or information cost) as orthogonal factors. As additional within-participant factor, we implemented three blocks with two different ratios. All participants started with a ratio of 1:10, which increased to a ratio of 1:40 in the second block, and decreased again to 1:10, the same ratio as in the first block. The ratio of 1:10 was always instantiated with the same points, 20:200. The instantiation of the ratio of 1:40 depended on the condition. The information-cost-change condition went through a change of information cost to a ratio of 5:200, while the payoff-change-condition faced a ratio of 20:800. 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 clearer.

The blocks consisted of 40 trials to give participants more time to adjust to each new ratio. The task was identical to the other experiments, except that the cover story was concerned with products and whether their reviews were predominantly positive or negative, and the maximum sample size was 16. Participants who did not receive feedback also did not see their balance displayed.

Results

Exclusions. Since binomial tests showed that their accuracy rate was not different from guessing, 18 participants had to be excluded from the analysis. Five additional participants were excluded because their mean sample size was below 0.75, which did not leave enough data to reliably assess their behavior, leaving a total sample size of 126 for analysis. None had to be excluded because of slow response times.

Influence of ratio. Table 7 shows the means and standard deviations of sample size, accuracy, and balance separately for block and the four conditions. Since the ratio changes across the three blocks consisted of an increase and an equivalent decrease back to the original ratio, the sample size cannot be expected to have a linear relation with the blocks. To be able to analyze the influences on sample size nonetheless, we decided to run two separate linear mixed models, one for comparing the first and second block, a second comparing the second and third block. To correct for the inflated probability of Type-1 errors, we applied a Bonferroni correction by decreasing our alpha level to .05/2= .025, which was used in the analyses below. Figure 7 visualizes the sample size across quarter blocks, separately for conditions and in comparison with the optimal strategy.

Table 7
Descriptive statistics by condition and block

Factors		Block	Ratio	*n (SD)*	*a (SD)*	*bal (SD*
Control	Cost	1	1:10 (20:200)	3.54 (1.82)	0.782 (.11)	42.2(38.0)
		2	1:40 (5:200)	5.03 (2.44)	0.840 (.099)	111.0 (32.2)
		3	1:10 (20:200)	4.26 (1.92)	0.836 (.092)	49.2 (34.2)
	Payoff	1	1:10 (20:200)	3.56 (1.78)	0.791 (.09)	45.0 (33.1)
		2	1:40 (20:800)	5.29 (2.94)	0.857 (.073)	464.8 (92.5)
		3	1:10 (20:200)	4.64 (2.68)	0.826 (.066)	37.4 (42.9)
Feedback	Cost	1	1:10 (20:200)	3.27 (1.69)	0.774 (.107)	44.2 (36.9)
		2	1:40 (5:200)	4.96 (2.37)	0.819 (.094)	102.8 (30.75)
		3	1:10 (20:200)	3.83 (1.85)	0.802 (.110)	44.4 (35.6)
	Payoff	1	1:10 (20:200)	2.06 (1.36)	0.735 (.124)	52.6 (38.0)
		2	1:40 (20:800)	4.08 (2.46)	0.837 (.086)	457.5 (105.8)
		3	1:10 (20:200)	2.72 (1.41)	0.782 (.133)	58.5 (34.4)

Note. Mean sample size *n*, mean accuracy *a*, mean balance *bal*, and standard deviation in brackets, by feedback condition, change condition and block

We fitted two linear mixed models to the data for Block 1 and 2, and the data for Block 2 and 3 with sample size as target variable. We included ratio (1:10 coded as -1 and 1:40 coded as 1), trial number (centered on 0), feedback (coded -1 for no feedback and 1 for feedback) and change (coded -1 for information-cost change and 1 for payoff change) as fixed effects. We also included a random intercept, a random slope for ratio, and a random slope for optimality that all varied by participant. Including more random slopes led to a failure to converge, so we chose to include ratio and optimality because we presumed that the most variance between individuals would be caused by their different responses to different ratios and their individual deviation from the optimal sample size. We used the R package afex for this analysis with the Satterthwaite' approximation to calculate the p-values (Singmann et al., 2023).

Figure 7

Sample size displayed by block quarters

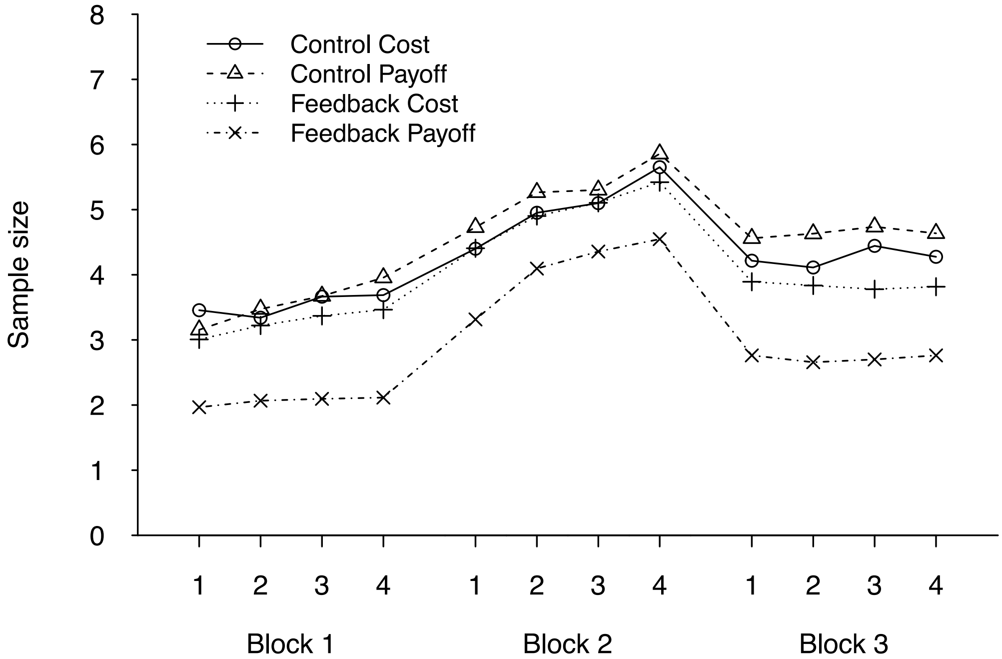

Note. Each data point is aggregated across one quarter of a block, i.e. ten trials.

In both models, ratio and trial number had a significant effect on sample size (see Table X for all test statistics of model comparisons). Overall, participants implemented different sample sizes for the different ratios, with higher sample sizes for the higher ratio ($\beta_1 = 1.64$, $SE_1 = 0.05$, $\beta_2 = 1.46$, $SE_2 = 0.05$). Averaged across ratios, they exhibited different sample sizes with increasing trial number ($\beta_1 = 0.01$, $SE_1 = 0.002$; $\beta_2 = 0.01$, $SE_2 = 0.002$). This is in line with our hypotheses. In Model 2, the interaction between ratio and trial was also significant, meaning that the change of sample size with trial number differed for the two ratios ($\beta_2 = 0.01$, $SE_2 = 0.002$). Hence, participants adjusted their sample size slightly less during the final block than during the previous block. This may be because the final block was a familiar ratio, or because that ratio required smaller optimal sample size to which participants were closer to. The equivalent interaction was not significant for Model 1, so there was no indication of a different sample size change between Block 1 and Block 2.

Influence of conditions. However, against our expectation, there was no interaction between the ratio and change condition in Model 1, and no main effect of change condition in Model 2. So whether participants experienced a payoff change or an information-cost change between ratios did not influence their sample size.

Regarding the effect of feedback, in Model 1 there was no interaction between feedback condition and ratio, and no interaction between feedback, ratio, and trial. The first would have indicated that the feedback condition had a different influence on sample size for the different ratios, which we expected, since all conditions received feedback in the first block and the manipulation entailed by the condition did not take effect until Block 2. The second interaction would have demonstrated that the feedback condition influenced the progression of sample size within the block differently for the two blocks. We expected participants with feedback to adjust their sample size more than participants without feedback, for which there is no evidence in Blocks 1 and 2. In Model 2, there was a significant interaction between ratio and feedback condition, indicating that the feedback condition had a different effect on sample size in the two ratios, with the sample size for the feedback conditions being closer to the no-feedback conditions in Block 2 than Block 3 ($\beta_2 = 0.116$, $SE_2 = 0.045$). What this difference in adaptation means will be determined in comparison with the optimal strategy, which will be examined next.

Optimality. To examine what the sample size changes above implied for participants' efficiency, we report the influence of the predictor optimality and its interactions next. First of all, there was a main effect of optimality in both models, meaning that overall the optimal sample size was larger than the actual sample size, ($\beta_1 = 1.23$, $SE_1 = 0.09$, $\beta_2 = 1.023$, $SE_2 = 0.10$). Optimality also interacted with ratio and trial in both models, indicating first, that an increase in ratio further increased the distance between optimal and actual sample size ($\beta_1 = 0.77$, $SE_1 = 0.03$, $\beta_2 = 0.97$, $SE_2 = 0.03$), and second, that an increase in trial number was associated with a decrease in the distance between the optimal and actual sample size ($\beta_1 = -0.01$, $SE_1 = 0.002$, $\beta_2 = -0.008$, $SE_2 = 0.002$). So participants managed to decrease the distance between their own and the optimal sample size over the course of the blocks, on average. This is in line with our expectation. The fact that the distance between optimal and actual sample size was larger for Block 2 is a reflection of the higher optimal strategy that the ratio implied, which participants did not come as close to as the lower optimal strategy. There was also a significant interaction between optimality, ratio, and trial, indicating that the improvement across trials depended on the ratio. For Block 2, the improvement across trials was even higher than for Block 1 and Block 3 ($\beta_1 = -0.001$, $SE_1 = 0.002$; $\beta_2 = -0.010$, $SE_2 = 0.00$).

Although there was no main effect of feedback on sample size, and no interactions with ratio or trial number, there was an interaction with optimality in the first model. This means that the difference between actual and sample size was moderated by the feedback conditions, however, with feedback conditions exhibiting a larger distance to the optimal sample size than no-feedback conditions ($\beta_1 = 0.22$, $SE_1 = 0.09$). This was unexpected. However, inspection of the individual condition shows that this is mainly driven by the feedback-payoff condition, which already displays a lower sample size in Block 1, before the feedback manipulation takes effect. This might be a false positive because the difference between conditions can only have been the result of chance.

There was no significant interaction between optimality, ratio, and feedback to complement the significant interaction between ratio and feedback reported above. Hence, the fact that the sample size difference between the feedback and the no-feedback conditions was smaller in Block 3 was not reflected in a different deviation from the optimal sample size.

In the second model, that is, for Blocks 2 and 3, there was a four-way interaction between optimality, ratio, feedback condition, and change condition. The difference between the optimal and actual sample size and how it differed for the two ratios was different among the four conditions, without a difference across both types of change for either type of feedback. More specifically, the increase in deviation from optimal strategy with the ratio was slightly smaller for the no-feedback payoff condition (β_2 = -0.06, SE_2 = 0.03). Again, this might be a chance finding related to the low initial sample size of the feedback-payoff condition which carried over to the other blocks. All other fixed effects were not significant (see Table 8 for all test statistics).The descriptive statistics for the optimality comparison of sample size, accuracy and balance across all conditions are summarized in Table 9.

Table 8

Model Comparisons for Fixed Effects for Blocks 1 and 2, and Blocks 2 and 3

		Model 1: Block 1 and 2			Model 2: Block 2 and 3		
Fixed Effect	df1	df2	*F*	*p*	df2	*F*	*p*
opt	1	122	170.41	<.001	122	97.841	<.001
ratio	1	121.99	1106.24	<.001	121.99	1043.98	<.001
trial	1	19756	31.21	<.001	19754.06	21.56	<.001
feedback (fb)	1	122	2.57	0.112	122	4.45	.037
change (ch)	1	122	1.28	0.26	122	0.82	.368
opt:ratio	1	19755.94	883.23	<.001	19754.02	1441.11	<.001
opt:trial	1	19755.94	28.19	<.001	19753.99	13.76	<.001
ratio:trial	1	19756	3.94	0.047	19754.06	9.18	.002
opt:fb	1	122	5.56	0.020	122	4.46	.037
ratio:fb	1	121.99	0.69	0.407	121.99	6.53	.012
trial:fb	1	19756	0.27	0.601	19754.06	1.10	.295
opt :pay	1	122	1.40	0.240	122	0.42	.517
ratio:pay	1	121.99	1.23	0.270	121.99	0.86	.355
trial:pay	1	19756	0.34	0.562	19754.06	0.12	.731
fb:pay	1	122	2.91	0.091	122	2.78	.098
opt:ratio:trial	1	19755.94	7.96	0.005	19753.99	20.55	<.001
opt:ratio:fb	1	19755.94	0.66	0.416	19754.02	0.51	.477
opt:trial:fb	1	19755.94	0.21	0.643	19753.99	0.08	.777
ratio:trial:fb	1	19756	0.05	0.816	19754.06	0.08	.777
opt:ratio:pay	1	19755.94	0.37	0.542	19754.02	1.18	.278
opt:trial:pay	1	19755.94	1.19	0.275	19753.99	0.63	.429
ratio:trial:pay	1	19756	0.30	0.583	19754.06	0.64	.423
opt:fb:pay	1	122	1.79	0.184	122	1.90	.171
ratio:fb:pay	1	121.99	0.37	0.546	121.99	0.11	.744
trial:fb:pay	1	19756	0.27	0.600	19754.06	0.02	.902

opt:ratio:trial:fb	1	19755.94	1.26	0.262	19753.99	0.15	.699
opt:ratio:trial:pay	1	19755.94	0.44	0.507	19753.99	0.97	.325
opt:ratio:fb:pay	1	19755.94	2.50	0.114	19754.02	5.12	.024
opt:trial:fb:pay	1	19755.94	0.29	0.588	19753.99	0.51	.476
ratio:trial:fb:payc	1	19756	0.04	0.839	19754.06	0.04	.839
opt:ratio:trial:fb:payc	1	19755.94	2.43	0.119	19753.99	0.10	.754

Note. This table displays the model comparison tests for the linear mixed model for Blocks 1 and 2 (Model 1) and for Blocks 2 and 3 (Model 2). Numerator degrees of freedom (df1) apply to both models; numerator degrees of freedom (df2), *F*-values, and *p*-values are listed separately for the two models for all fixed effects. Rows with gray background highlight significant effects.

Table 9

Descriptive statistics for actual and optimal variables and their difference by block

Block	*n*	*n opt*	*n diff*	*a*	*a opt*	*a diff*	*bal*	*bal opt*	*bal diff*
1	3.21	4.01	0.80	.774	.856	.082	45.4	62.2	16.8
	(1.76)	(0.54)	(1.86)	(.106)	(.058)	(.106)	(35.9)	(24.8)	(28.5)
2	4.92	8.86	3.94	.837	.931	.093	272.7	309	36.3
	(2.58)	(0.91)	(2.55)	(.088)	(.370)	(.090)	(191.6)	(200.0)	(72.0)
3	3.98	3.97	-0.01	.813	.862	.049	45.7	65.5	19.81
	(2.17)	(0.55)	(2.15)	(.101)	(.055)	(.093)	(37.8)	(24.2)	(31.4)

Note. Means with *SD* in brackets for actual and optimal sample size *n (opt)*, their difference *n diff*, actual and optimal accuracy *a (opt)*, their difference *a diff*, actual and optimal balance *bal (opt)*, their difference *bal diff.* Difference scores were created at the participant level and averaged across participants.

Discussion

The results showed that participants adapt to changes in ratio not just by changing at the start of each block, but also by changing within the blocks. The changes within the blocks also brought them closer to the optimal sample size. However, this change did not differ between conditions that received feedback and those that did not. All participants changed their sample size across trial number approximately the same way. This is astounding. Feedback provides the necessary information to evaluate one's performance and to improve it. At least in the conditions without feedback the within-block change occurred without this information. Instead, participants in the no-feedback conditions may have been simulating 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) also 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. However, since the conditions that did receive feedback had a much better informational basis to improve their efficiency, it is surprising that they did not fare better than the no-feedback conditions. Since processing feedback and inferring what type of strategy is more efficient is not a straightforward process, it is possible that the feedback was not as instructive as we made it out to be. We designed a final experiment to examine the effects of a feedback and experiment structure that should make it easier to improve.

Experiment 4

Hypotheses

This experiment featured an improved type of feedback. In addition to a condition who received the usual trial-by-trial feedback about the different aspects of a single task, we implemented a condition with block feedback, provided after every ten trials, summarizing all aspects across those trials. Summarizing the feedback should take some cognitive strain off the participants since they did not have to integrate their accuracy and information costs across multiple trials themselves. In addition, the blocks should make the comparison of strategies easier since they provided a clear structure for when and how long a particular strategy should be explored and provided summarized feedback for precisely that unit. The condition that received block-feedback did not receive trial-by-trial feedback. We expected this to make it easier to pursue a strategy consistently, because participants would be less tempted to implement changes every time they were informed of an incorrect or correct decision. Hence, participants were provided with excellent conditions for improvement and only had to make use of the blocks to try out different strategies and keep in mind the final feedback to be able to compare the strategies. Because the focus was on the fine-tuning process of adapting within a ratio based on experience, we only implemented one ratio across all blocks.

We hypothesized that participants who received block feedback would change their sample size more across blocks than those who received trial-by-trial-feedback. We expected this change to lead to an approximation to the optimal sample size. Accuracy and balance will only be reported descriptively, because they depend so highly on sample size.

Methods

Participants. A total of 176 participants completed the study via Prolific. 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.[27] 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.

Design and Procedure. The experiment had two between-participant conditions that differed in the feedback participants received. Participants in the block-feedback condition received summarized feedback at the end of each block, while participants in the trial-

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

feedback condition received feedback after each trial. Each block consisted of 10 trials and participants worked through 10 blocks in total.[28]

The maximum sample size per decision was 32. To facilitate the processing of the trial-by-trial feedback, the components appeared one by one at a rate of 750 ms: whether the decision was correct and the number of points gained or lost, the points spent on the information cost, the resulting total payoff for that trial, and the updated cumulative balance that resulted in. After an additional 1500 ms, the next trial or block started automatically.

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 and how much they incorporated the different elements of the feedback (decision correctness, decision outcomes (gains and losses), points spent on reviews, and the total balance) into their strategy by clicking on a continuous slider anchored on ‘very much’ and ‘not at all’ at the ends.

Results

Exclusions. Four participants were excluded because they had mean response times of more than 3 *SD* above the mean for the first arrow. The response time for the first arrow is of particular importance, because the previous feedback page redirected to the sampling page automatically. Not paying attention to the feedback would likely lead to missing the start of the next trial and to longer response times for the first arrow. Since paying attention to the feedback was very important for the experiment, participants who may have not been paying attention to it needed to be excluded.

For participants whose mean response time for the first arrow was between one and two *SD* above the mean, we excluded them if they had more than three interruptions of more than 15 seconds, which applied to another 4 participants. 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.

Changes across block and trial. The mean differences provided in Table 10 give a first overview of the data and how close the actual and optimal sample size and accuracy are across blocks. What is more informative with regard to our hypotheses is the development of sample size across the ten blocks of the experiment, illustrated in Figure 8, and whether the different feedback types had an influence on that. Sample size was not only of interest in itself, but especially in comparison with the optimal sample size, which allowed us to assess whether changes in sample size approximated the optimal sample size more or deviated further from it.

Hence, to examine what happened over the course of the experiment, we fitted a linear mixed model with sample size as target variable (using the R package afex, Singmann et al., 2023). We included condition (coded as -1 for the trial-feedback condition and +1 for the block-feedback condition), optimality (-1 for actual sample size and +1 for optimal sample size), block (centered on 0, from -4.5 to +4.5), trial (centered on 0, from -4.5 to +4.5) as fixed

[28] Due to a programming error, the trial-by-trial feedback conditions actually worked through 12 blocks; the last two block was discarded to have equal amounts of data for the two conditions.

effects. We also included a random intercept and random slopes for block and optimality that all varied by participant.

Table 10

Descriptive statistics for optimal and actual dependent variables and their difference

Cond	***n***	***n opt***	***n diff***	***a***	***a opt***	***a diff***	***bal***	***bal opt***	***bal diff***
Block	5.62	4.5	-0.77	.850	.862	.019	83.8	99.7	15.36
fb	(3.27)	(0.54)	(0.65)	(.082)	(.036)	(.036)	(22.22)	(15.37)	(15.44)
Trial	4.92	4.26	-1.01	.835	.856	.013	84.9	99.6	15.23
fb	(2.27)	(0.44)	(0.55)	(.086)	(.036)	(.035)	(26.39)	(14.50)	(14.59)

Note. Means with *SD* in brackets for actual and optimal sample size *n (opt)*, their difference *n diff*, actual and optimal accuracy *a (opt)*, their difference *a diff*, actual and optimal balance *bal (opt)*, their difference *bal diff.* Difference scores were created at the participant level and averaged across participants.

The tests of the fixed effects and the estimates and standard errors of the estimate are displayed in Table 11 and show that there was neither a main effect of block, nor of trial on sample size. The interaction of trial and block also did not influence sample size. There was a small effect of feedback that was significant, so what type of feedback participants received did have a very small effect on their sample size, with the block-feedback condition displaying slightly higher sample sizes. However, considering the lack of other significant relations, this effect needs be interpreted cautiously. Without an interaction between feedback type and block (or trial, or both), it is unlikely that the block feedback actually caused a change in sample size across the experiment. In line with this lack of change, there is also no evidence that the feedback type caused an approximation of the sample size to the optimal sample size, since the interaction between block, feedback, and optimality is not significant (and neither is the interaction between trial, feedback, and optimality or the interaction between block, trial, feedback, and optimality). Hence, instead of providing a better information base to improve performance, the feedback type may rather have had an immediate influence that resulted in larger samples. Perhaps the absence of the trial-by-trial feedback encouraged participants to sample more, potentially because they did not see their balance evolving and misjudged the cost of information. There was no evidence that this higher sample size also led to a greater distance from the optimal strategy, as the interaction between optimality and feedback was not significant.

Overall, participants sampled more than the optimal strategy by a about half an observation, on average. The block number actually moderated this relation, indicating that as the blocks progressed participants' sample size deviated further from optimality, although this effect was very small.

Table 11

Statistics of model comparisons including estimates of fixed effects

Fixed effect	Num. *df*	Den. *df*	*F-value*	*p-value*	β	*SE*
block	1	158	1.11	.294	0.02	0.02
trial	1	31508	3.15	.076	0.01	0.01
feedback (fb)	1	158	4.21	.042	0.23	0.11
optimality (opt)	1	158	14.38	<.001	-0.45	0.12
block:trial	1	31508	0.02	.880	0.00	0.00
block:feedback	1	158	0.10	.753	0.00	0.02
trial:feedback	1	31508	0.48	0.488	-0.01	0.01
block:opt	1	31508	6.35	0.012	-0.02	0.01
trial:opt	1	31508	2.95	.086	-0.01	0.01
feedback:opt	1	158	0.96	.329	-0.12	0.12
block:trial:feedback	1	31508	0.08	.778	0.00	0.00
block:trial:opt	1	31508	0.07	.789	0.00	0.00
block:feedback:opt	1	31508	0.29	.589	0.00	0.01
trial:feedback:opt	1	31508	0.04	.834	0.00	0.01
block:trial:feedback:opt	1	31508	0.28	.594	0.00	0.00

Note. Displaying the numerator and denominator degrees of freedom, *F*-value, and *p*-value for the comparison of the model fit with compared to without the fixed effect, and additional estimate and standard error for the respective fixed effects displayed as rows, where fixed effects listed together separated by a colon indicate interactions

Figure 8

Actual and optimal sample size across blocks separately for condition

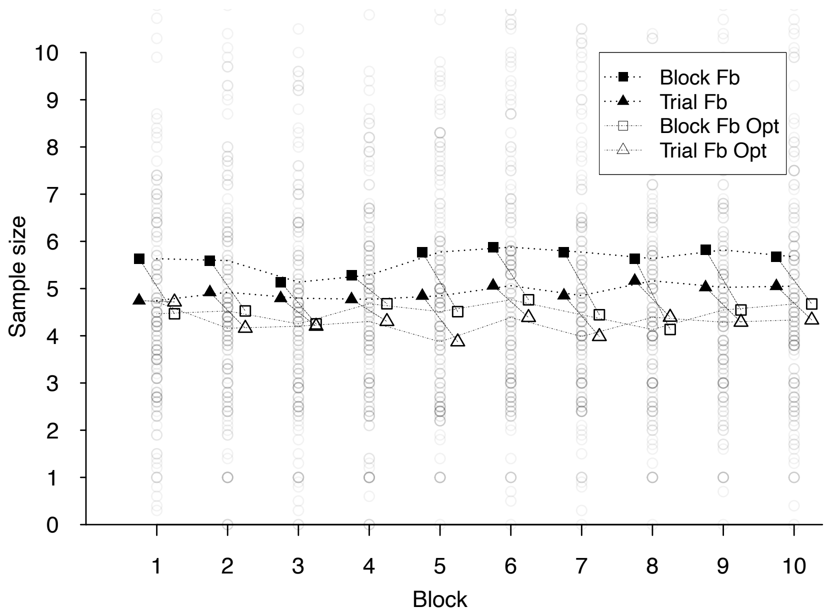

Note. Sample size across the ten blocks for the two conditions represented by black symbols and their respective optimal sample size represented by the same symbol in white. Individual participants' block averages as light grey circles.

Questionnaire. In the trial-feedback condition, 61.11% of participants indicated that they changed their strategy during the experiment, while 40.23% did in the block-feedback condition. This difference was not significant, as a test for equality of proportions showed, ($\chi^2(1) = 6.06$, $p = .014$). The mean ratings of how much participants incorporated the elements of the feedback into their strategy, displayed in Table 12, only differed significantly between conditions for the feedback about decision outcomes, that is, gains and losses. Overall, however, participants did indicate that they incorporated the feedback into their strategy, so the lack of change was presumably not just the result of participants not paying attention to the feedback.

Table 12

Descriptive and inferential statistics for questionnaire results

Feedback aspects	Trial-feedback Mean (*SD*)	Block feedback Mean (SD)	*df*	*t*-value	*p*-value
correctness	719.65 (260.61)	699.40 (270.53)	157	0.48	0.633
outcomes	751.01 (212.68)	656.87 (270.59)	157	2.40	0.018
points spent on reviews	621.53 (290.49)	562.47 (306.43)	157	1.24	0.217
balance	660.31 (283.24)	585.12 (307.87)	157	1.59	0.114

Note. Participants' mean ratings of how much they incorporated the feedback about different aspects into their strategy, test-statistics of two-sample t-tests between the two conditions (scale 0-1000).

Discussion

Although the optimal strategy represents the most efficient use of resources, which is a core element of adaptivity, the deviation from the optimal sample size is not problematic in itself, since participants were not familiar with the task and are not required to just know how to use resources most efficiently. What is concerning, however, is that participants did not display any systematic changes in sample size across the experiment, regardless of whether they were provided with feedback after every trial or whether they received summarized feedback at the end of each block. If anything, their sample size deviated further from optimality as the blocks progressed. This was not in line with our expectations. We expected that 100 trials of a fairly simple task would supply participants with ample opportunity to improve at least slightly. 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. However, there may have been an issue well before the feedback became relevant. If participants did not change their sample size systematically, it is plausible that they did not even try out different strategies, leaving nothing to monitor and evaluate comparatively to arrive at a better strategy.

This does not necessarily mean that participants are not adaptive. It does mean that the conditions provided by us are not sufficient for participants to engage in the fine-tuning of their strategy that would lead to improved efficiency, which is interesting in itself. It indicates that the conditions required for adaptivity may be a lot more specific than we previously assumed.

General Discussion

We assumed that people use two different approaches to adapt to new decision situations. One is an a-priori planning of change in response to explicit announcements of context changes, which we saw in Experiment 1 and which we made more difficult in Experiment 2. The second assumed process is what we called fine-tuning and describes further improvement based on one's experience with the particular information cost-benefit trade-off. However, we were not able to observe this process, either at all or not in the way we expected. Experiment 3 demonstrated sample-size adaptation across trials – but this adaptation did not rely on feedback. So either there is a process that can occur without feedback just by estimating one's accuracy and simulating the effects of different information costs and payoffs to arrive at a good substitute for real feedback. More likely, however, the adaptation was a delayed implementation of the a-priori strategy change planned in response to the announcement. This would be in line with the findings in Experiment 4, where we did not observe any change in sample size at all.

Hence, with regard to the first process, people can definitely be considered adaptive. However, with this process alone they did not achieve particularly efficient resource use, which is the core of adaptivity. The second, fine-tuning process is required additionally, for which the jury is still out. It is unlikely that people are completely unable to engage in this process, since it has been demonstrated in other paradigms (e.g. Evans and Brown, 2017) and is essentially a basic learning process without which we would probably be unable to function as well as we do. More likely, we have not yet succeeded in providing participants with the necessary assistance or structure. This may also have been the case in speed-accuracy trade-off experiments investigating similar questions, in which participants also failed to improve (Fiedler et al., 2021).

From a metacognitive perspective, one aspect that might be missing is a particular combination of control and monitoring functions (Nelson & Narens, 1990). While participants reported to engage with the feedback in Experiment 4, indicating monitoring of their performance, there may not have been any changes in strategy to monitor and compare in the first place. The control function has to step in and actually implement different strategies for the monitoring to function can yield meaningful results. Although we expected the short-block structure to naturally invite this kind of exploration with the ideal feedback provided at the end, it might not have been enough. Participants might not have been aware how much they could improve because with a fairly high accuracy rate, one tends to receive many gains and have a steadily increasing cumulative balance. On the other hand, they may have been even more adaptive than we ever thought to consider. Perhaps they can estimate much more they could have gained and also known that it is not much more than they were already receiving. Since their balance was not much worse than optimal, one could argue that the effort and the potential losses that the exploration of different strategies implies was not the worth the benefits of eventually executing a better strategy. Investigating this further would be of great value to further understanding adaptivity.

An indication of how much better participants could perform or how others are performing might be enough to prompt exploration and evaluation. This would be similar but also crucially different to the types of feedback that Evans and Brown (2017) implemented. Their high feedback did not only provide participants with the information of how much better they could be, but also straight up told them 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. We would be more interested in prompting participants to explore themselves, in other words to find the minimum assistance for them to improve efficiency.

There may be other issues contributing to a supposed lack of adaptivity in our paradigm. Perhaps people are better adapted to a dynamic environment (Navarro, Newell, & Schulze, 2016), or might have been searching for patterns where there are none (Plonsky, Teodorescu, & Erev, 2015) instead of accepting the stochastic nature of the task and focussing 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 world instead of the lab. Hence, we cannot and do not want to conclude that people are not adaptive. The reason we implemented a somewhat artificial task that was not intended to represent the environments' informational structure (Dhami, Hertwig, & Hoffrage, 2004), is because we are not just interested in determining whether people are adaptive in close-to real-world decision situations. We wanted to find potential processes that allow people to become adaptive by confronting them with an unfamiliar task that only emulated some real-world features, such as the sampling of observations. Hence, the artificial task structure is both the greatest strength and limitation of the experiments presented.

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 the fine-tuning process. 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 cost 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 considering the resources required for improving it (e.g., Lieder et al., 2017).

Conclusion

Overall, the experiments demonstrated that participants strive to adapt to changes in decision environments through an a-priori planning process. They also showed that the conditions for further improvement of efficiency through a secondary fine-tuning process are not straightforward. We offer a few ideas of what might have been missing for successful adaptation from a metacognitive perspective. Further investigating these questions would allow us to better understand human information search and decision making and how it is and develops to be adaptive. This more complete theory of adaptivity might eventually allow us to assist both decision making directly – where even an adaptive response is bound to fall short – but also assist the development of adaptive decision making for environments that may not naturally offer the required conditions.

Acknowledgements: -

Competing interests: The authors declare none.

Financial Support statement: This work was supported by the Deutsche Forschungsgemeinschaft (DFG) under Grant FI 294/30-1.

References

Balci, F., Simen, P., Niyogi, R., Saxe, A., Hughes, J.A., Holmes, P., & Cohen, J.D. (2011). Acquisition of decision-making criteria: reward rate ultimately beats accuracy. *Attention, Perception, & Psychophysics*, *640*(2), 657.

Bates D., Mächler M., Bolker B., Walker S. (2015). "Fitting Linear Mixed-Effects Models Using lme4." *Journal of Statistical Software, 67*(1), 1–48. https://doi.org/10.18637/jss.v067.i01

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

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

Evans, N. J., Bennett, A. J., & Brown, S. D. (2019). Optimal or not; depends on the task. *Psychonomic Bulletin & Review*, *26*(3), 1027–1034. https://doi.org/10.3758/s13423-018-1536-4

Fiedler, K., Ackerman, R., & Scarampi, C. (2019). Metacognition: Monitoring and controlling one's own knowledge, reasoning and decisions. In R. J. Sternberg & J. Funke (Eds.). *The Psychology of Human Thought: An Introduction* (pp. 89-111). Heidelberg: Heidelberg University Publishing.

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

Fiedler, K., & Wänke, M. (2009). The cognitive-ecological approach to rationality in social psychology. *Social Cognition*, *27*(5), 699–732. https://doi.org/10.1521/soco.2009.27.5.699

Fried, L. S., & Peterson, C. R. (1969). Information seeking: Optional versus fixed stopping. *Journal of Experimental Psychology*, *80*(3 PART 1), 525–529. https://doi.org/10.1037/h0027484

Gigerenzer, G., Todd, P. M., & The ABC Research Group. (1999). *Simple heuristics that make us smart.* Oxford University Press.

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

Khodadadi, A., Fakhari, P., & Busemeyer, J. R. (2014). Learning to maximize reward rate: A model based on semi-Markov decision processes. *Frontiers in Neuroscience*, *8*. https://doi.org/10.3389/fnins.2014.00101

Lee, M. D., & Cummins, T. D. R. (2004). Evidence Accumulation in Decision Making: Comparing 'Take the Best' and 'Rational' Models. *Psychonomic Bulletin & Review*, *11*(2), 343–352. https://link.springer.com/article/10.3758/BF03196581

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., M. Huys, Q. J., & 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

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. *Psy- chology of Learning and Motivation*, *26*, 125–173. https://www.doi.org/10.1016/s0079-7421(08)60053-5

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

Simen, P., Contreras, D., Buck, C., Hu, P., Holmes, P., & Cohen, J. D. (2009). Reward rate optimization in two-alternative decision making: Empirical tests of theoretical predictions. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(6), 1865–1897. https://doi.org/10.1037/a0016926

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review, 63,* 129 –138. http://dx.doi.org/10.1037/ h0042769

Singmann H., Bolker B., Westfall J., Aust F., Ben-Shachar M. (2023). _afex: Analysis of Factorial Experiments_. R package version 1.2-1, https://CRAN.R-project.org/package=afex.

Slovic, P., & Lichtenstein, S. (1971). Comparison of Bayesian and regression approaches to the study of information processing in judgment. *Organizational Behavior and Human Performance*, *6*(6), 649–744. https://doi.org/10.1016/0030-5073(71)90033-X

Snapper, K. J., & Peterson, C. R. (1971). Information seeking and data diagnosticity. *Journal of Experimental Psychology*, *87*(3), 429–433. https://doi.org/10.1037/h0030557

Söllner, A., & Bröder, A. (2015). Toolbox or Adjustable Spanner ? A Critical Comparison of Two Metaphors for Adaptive Decision Making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *42*(2), 215–237. https://doi.org/10.1037/xlm0000162

Starns, J. J., & Ratcliff, R. (2010). The effects of aging on the speed–accuracy compromise: Boundary optimality in the diffusion model. *Psychology and Aging*, *25*(2), 377.

Todd, P. M., & Gigerenzer, G. (2007). Environments that make us smart: Ecological rationality. *Current Directions in Psychological Science*, *16*(3), 167–171. https://doi.org/10.1111/j.1467-8721.2007.00497.x

List of Submitted Scientific Publications

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
2. McCaughey, L., Prager, J. & Fiedler, K. (2023a). Rivals reloaded – Adapting sample-based speed-accuracy trade-offs through competitive pressure. Under review at *Journal of Experimental Psychology: Learning, Memory, and Cognition.*
3. McCaughey, L., Prager, J. & Fiedler, K. (2023b). The information cost-benefit trade-off as a sampling problem in information search. In Fiedler, K., Juslin, P., & Denrell, J. (Eds.). *Sampling in Judgment and Decision Making* (pp. 324-356). Cambridge University Press. *In press.*
4. McCaughey, L., Prager, J. & Fiedler, K. (2023c). Adapting to information search costs in sample-based decisions. Under review at *Judgment and Decision Making.*

UNIVERSITÄT HEIDELBERG
ZUKUNFT SEIT 1386

FAKULTÄT FÜR VERHALTENS- UND EMPIRISCHE KULTURWISSENSCHAFTEN

Promotionsausschuss der Fakultät für Verhaltens- und Empirische Kulturwissenschaften der Ruprecht-Karls-Universität Heidelberg / Doctoral Committee of the Faculty of Behavioural and Cultural Studies of Heidelberg University

Erklärung gemäß § 8 (1) c) der Promotionsordnung der Universität Heidelberg für die Fakultät für Verhaltens- und Empirische Kulturwissenschaften / Declaration in accordance to § 8 (1) c) of the doctoral degree regulation of Heidelberg University, Faculty of Behavioural and Cultural Studies

Ich erkläre, dass ich die vorgelegte Dissertation selbstständig angefertigt, nur die angegebenen Hilfsmittel benutzt und die Zitate gekennzeichnet habe. / I declare that I have made the submitted dissertation independently, using only the specified tools and have correctly marked all quotations.

Erklärung gemäß § 8 (1) d) der Promotionsordnung der Universität Heidelberg für die Fakultät für Verhaltens- und Empirische Kulturwissenschaften / Declaration in accordance to § 8 (1) d) of the doctoral degree regulation of Heidelberg University, Faculty of Behavioural and Cultural Studies

Ich erkläre, dass ich die vorgelegte Dissertation in dieser oder einer anderen Form nicht anderweitig als Prüfungsarbeit verwendet oder einer anderen Fakultät als Dissertation vorgelegt habe. / I declare that I did not use the submitted dissertation in this or any other form as an examination paper until now and that I did not submit it in another faculty.

Vorname Nachname / First name Family name	Linda McCaughey
Datum / Date	30.3.2023
Unterschrift / Signature	Dem Dekanat der Fakultät für Verhaltens- und Empirische Kulturwissenschaften liegt eine unterschriebene Version vom 30.03.2023 vor.