Ackerman, R. A., Douven, I., Elqayam, S., & Teodorescu, K. (2020). Satisficing, meta-reasoning, and the rationality of further deliberation. In S. Elqayam, I. Douven, J. St. B. T. Evans, & N. Cruz (Eds.), Logic and uncertainty in the human mind: A tribute to David Over. London, UK: Routledge.

Satisficing, Meta-reasoning, and the Rationality of Further Deliberation

Rakefet Ackerman[1], Igor Douven[2], Shira Elqayam[3], and Kinneret Teodorescu[1]

[1]Faculty of Industrial Engineering and Management, Technion

[2]SND / CNRS / Sorbonne University

[3]Faculty of Health and Life Sciences, De Montfort University

Author Note

All authors contributed equally to the paper and are listed alphabetically.

Abstract

David Over has made seminal contributions to the study of human rationality, most memorably in the now-classic distinction, made in collaboration with Jonathan Evans, between normative and instrumental rationality. In this chapter, we discuss an under-explored aspect born of the tension between the two: the rationality of searching for further choice options. We review several candidates for a model of further deliberation, including classic Bayesian decision theory, bounded rationality, and the research domain concerned with the "secretary problem." We demonstrate how none of these can satisfy the twin adequacy criteria calling for an account of processing as well as an account of rationality. We then present a model based on a proposal made in Douven (2002) and the psychological model of meta-reasoning. We conclude with implications for human rationality.

Satisficing, Meta-reasoning, and the Rationality of Further Deliberation

Introduction

Imagine the following: Kaushika wants to buy a house. She sees an advert in the local paper, and contacts the seller to arrange a viewing. She views the house. She has not viewed any other houses, so at this point in time she has a singleton set of one house to choose from. At this stage, she stops the search and makes an offer for the house.

When is it rational to put a stop to one's search for alternatives? Is terminating a search after considering a set of one premature? In this chapter, we will outline some prerequisites for a theory of further deliberation. Classic Bayesian decision theory faces a problem: its definition of rationality hinges on selecting the action which maximizes utility out of a given set. However, little consideration is given to the question of how this set is constructed (although cf. Baron, 2008, p. 62 f). In any given situation, there might be myriad options for action. How the agent searches for further options and when she stops is of paramount significance for the psychology and philosophy of rationality alike. According to one theoretical proposal (Douven, 2002), if an initial set of options exists, in deciding when to search for further options and when to stop the search, the agent attends to (a) how satisfactory the existing options are; (b) how likely it is that she will find a better option; and (c) how costly it is to search further.

In this chapter, we reconsider several candidate models for the rationality of further deliberation in light of a set of adequacy criteria. We then develop a theoretical model based on updating the philosophical insights from Douven's (2002) proposal with recent developments in the psychology of meta-reasoning, the field dedicated to exploring the way that people monitor, regulate, and invest cognitive effort. We will draw on Ackerman and Thompson's (2015; 2017a; 2017b) recently proposed model of meta-reasoning to identify the meta-reasoning processes involved. For example, initial Feeling of Rightness, one's fast intuitive judgment of how good their initial response is, can account for (a); and judgment of solvability ("can I solve this?") or intermediate confidence may account for (b). While there is no direct parallel for (c), cost can be construed as mostly related to time investment, which is implicit in the meta-reasoning model. We will also argue

that it might be computed as an inverse of intermediate confidence in the response. Alternatively, there is space to consider enriching the meta-reasoning framework with monitoring and control factors that directly measure the cost of deliberation. It is also possible to flesh out Douven's model with further research questions taken from the meta-reasoning framework, for instance, concerning the way that people readjust their solution criterion downwards as they deliberate for longer (Ackerman, 2014). This integration can benefit both models of rationality and empirical research questions on the way that people search the alternatives space.

The topic of rational choice, especially as seen from a psychological viewpoint, seems a highly appropriate choice for a chapter written for a Festschrift volume for David Over. In a classic volume on reasoning and rationality (Evans & Over, 1996), David combined forces with Jonathan Evans (see Manktelow & Evans, this volume) to elaborate on what has by now become a seminal distinction between two types of rationality: rationality$_1$, or pragmatic or instrumental rationality, the rationality defined by the agent achieving her goals; and rationality$_2$, or normative rationality, the type of rationality defined by conforming to a normative standard such as logic or the probability calculus. One of the earliest chapters in the monograph was entitled "Problems with normative decision theory," discussing, among other issues, attempts to make decision theory more responsive to human psychology, specifically, to human cognitive limitations.

We submit that the issue of how a set of possibilities is constructed, and when people stop searching, is of paramount importance for human rationality, but under-explored in philosophy and the psychology of reasoning and decision-making alike. However, a new theoretical development in the psychology of reasoning now directs the theoretical spotlight exactly at this issue and related ones: the nascent research area of meta-reasoning, mentioned above. In this chapter, we explore how psychology of meta-reasoning can shed light on the search for possibilities, focusing in particular on the model from Douven (2002).

The rest of this chapter is structured as follows. We start with stating adequacy criteria for a theory of further deliberation. We then examine the candidate models available in

the literature in light of these criteria, including identifying issues with standard decision theory, as stated in Douven's proposal. We then proceed to presenting in greater detail the model of meta-reasoning suggested by Ackerman and Thompson. Finally, we explain how meta-reasoning can help to address the issue of further deliberation.

Adequacy criteria for the rationality of further deliberation

A theory of further deliberation needs to fulfill two criteria. To be descriptively adequate (and psychologically viable), it needs to specify a representational and processing mechanism. In Marr's terms (1982), a theory of further deliberation needs to have an algorithmic explanation, an explanation of *how* further deliberation is carried out. Another, no less important, criterion concerns rationality, specifically, the theory should explain when is it rational to keep searching and when is it rational to stop. We will examine our candidates in view of these two criteria. We will show that no extant theory satisfies both, but that our model has the potential to do so.

Candidate 1: Bayesian decision theory

Our first candidate for a theory of further deliberation is standard Bayesian decision theory (henceforth BDT). This is perhaps the weakest candidate: it offers neither a processing account nor a rationality one. In fact, the whole question is entirely outside its remit. BDT is concerned with situations in which the options we can choose from are already laid out before us, and the question the theory then aims to answer is which of those options we ought to choose. The answer it gives presumes that we can measure how desirable each of the options is under each of the various possible ways the world may turn out to be, as far as we are aware. Specifically, the theory implores us to choose the option with greatest probability-weighted average desirability, where the probabilities are those attached to the various possibilities that we consider to be open. In more formal terminology, the answer is that we ought to choose the option which maximizes expected utility. If there is no unique option satisfying this condition, then we should choose (arbitrarily) one among those with maximum expected utility.

Sometimes in life, decision-making situations can be adequately modeled using BDT, but these situations are few and far between. Moreover, the options are not always as if presented to us on a silver plate, and finding out what options we have, and when to stop looking for further options, are integral parts of the decision problem we are facing. Consider, for instance, what by now has become a relatively common practice: a patient consulting another doctor for a second opinion. A second doctor, possibly aware of some publications or developments her own doctor may have missed, may see some treatment options that her own doctor did not bring up during the consultation. By the same token, it might appear rational to consult a third doctor, who might be able to conceive of further treatment options still. While this process may have diminishing returns, as well as practical procedural and financial considerations, it does not seem to have an a priori rational stopping point. It seems, then, that the patient is facing another decision problem before she can decide upon her treatment, to wit, the problem of deciding how much time and effort to invest in finding out about possible treatment options. But then, in trying to make *that* decision, the problem of how to delineate what is going to count as the given options appears anew.

The way we have presented this "problem of further deliberation" already suggests that treating it as a kind of higher-order decision-making problem and then invoking BDT at a meta-level (or meta-meta-level, and so on) to solve it is not going to work, at least not in general. It is in fact known from the history of analytic philosophy that by "going higher-order" one is typically just kicking the can down the road: one solves a problem at order n only at the expense of seeing it reappear at order $n + 1$.

Another approach is to pretend that the problem does not exist. In the above case, for instance, we might make the idealizing assumption that patients always choose from the treatment options presented to them by their own doctor. While this may look like cheating, idealization is a common strategy in science, and one that often enough has proved valuable (Kuipers, 2001). Idealizing assumptions about human decision making may still allow us to make accurate predictions at the aggregate level. And that economists and sociologists are generally quite happy with BDT, idealizations (e.g.,

about the options always being given to us) included, is precisely because, indeed, at the group level these idealizations arguably have a negligible influence, so that, at that level, BDT may still help us arrive at accurate predictions. For instance, by (probably falsely) assuming that all patients with a certain condition will consider the same treatment options, a health insurance provider may be able to fairly accurately estimate the cost of offering a specific plan.

However, even economists are no longer entirely happy with BDT; witness the Nobel Prizes in economics given to Herbert Simon (1978), Daniel Kahneman (2002), and, most recently, Richard Thaler (2017)—all for contributions flying in the face of classic BDT. As psychologists, our main focus is on behavior, and idealization fails the processing criterion: it does not specify *how* people arrive at their decisions. Moreover, it is not really a viable rationality option, since it makes unrealistic assumptions about human rationality at the level of the individual, and at this level, idealizations are likely to diminish the chances of obtaining an empirically accurate model. After all, we are not dealing with idealized people, and unlike at the aggregate level, there is nothing by which the effects of idealizations might somehow cancel out. At this level, then, and at least from a psychologist's perspective, BDT might serve as a starting point, but one can hardly expect it to be a realistic model of human decision making as it stands. Indeed, Evans and Over (1996, pp. 29–36) severely criticized BDT for being at the same time too weak and too strong to capture our notions of pragmatic rationality; for example, it requires absolute knowledge of the consequences of behavior as well as total coherence, while at the same time eschewing a critique of subjective values.

Candidate 2: Bounded rationality

Simon (1979; 1982) famously observed that, often, an agent has two intertwined decisions to make rather than one: which option to choose, and when to stop searching for a solution. As he further pointed out, as biological computational machines, humans have limited resources; for instance, we do not have unlimited working memory and attentional resources. Thus, human rationality is necessarily *bounded*: we do not exhaustively search

for solutions to find the *best* one, we just look for a solution that is *good enough*. In Simon's nomenclature, we *satisfice* rather than *optimize*.

Bounded rationality can be seen as an alternative candidate for a theory of further deliberation. However, it falls short for both psychological and philosophical reasons. It fails our first adequacy criterion, since it offers no processing model. On the philosophical angle Douven (2002, p. 322), according to bounded rationality accounts, it is rational to choose the first option that is good enough, that is, the first option an agent comes across whose expected utility meets her aspiration level. But this can yield counter-intuitive results. For suppose that, as in our example of flat search, an agent who does come across a good enough option is, at the same time, highly confident that deliberating further for just a short while is almost certain to result in her conceiving of a much better option. Also suppose there is no impediment for her to spending a little longer deliberating. Then choosing the first good-enough option would make her seem rash. By contrast, the intuition is easily accounted for by Douven's (2002) proposal, given that whether the choice is rational depends also on how likely it is, according to the agent, that she can find a better alternative.

Candidate 3: Optimal stopping problems

Yet another prominent candidate for further deliberation is what has been dubbed "optimal stopping problems" in mathematics. The most famous example has been named "the secretary problem." Imagine that you are a CEO looking to hire the best secretary in town (your board of directors specifically instructed you to make no compromises). All qualified applicants are invited for an interview. Applicants are randomly interviewed one after the other, but you can only hire a candidate following the interview; previous applicants become unavailable once you proceed to interview someone else (e.g., because they find a job elsewhere). You must fill the vacancy so you will proceed until you hire someone. The challenge is when to stop the search. As it happens, under a few simplifying assumptions, there is a mathematically derived solution of when it is optimal to stop (Ferguson, 1989; Freeman, 1983): Interview N/e candidates, where N is the

known number of candidates and e is Euler's number, and reject them. Then, hire the first candidate after N/e who is better than all previous ones.

Prima facie, this seems to be our best candidate for a theory of further deliberation. It is tailor-made for further deliberation, and it offers a mathematical model of optimal behavior. Yet, the elegant solution described above is too limited in scope. We note that a mathematical model is not the same as a rationality model. While it offers a possible norm, there is much more to rationality than following a normative model, especially when the model is as restricted as this one. First, the optimal cutoff directly implies that the decision maker should have full knowledge about the total number of applicants/options, which is rarely the case in real life situations (e.g., when searching for a marriage partner). Second, this optimal solution holds only when previous applicants become unavailable. This assumption too is violated in many real life scenarios. For example, when searching for a laptop, one can review options and decide to buy one of the first laptops she examined. Last, the solution is optimal only for selection of the best candidate. That is, only the very best choice "wins"—making a good enough choice is not taken into account. In Simon's terms, it is an optimizing, rather than satisficing reward structure. Small wonder, then, that research shows a large normative–descriptive gap: people typically terminate the search earlier than that mathematically defined optimum (e.g., Bearden, Murphy, & Rapoport, 2006; Rapoport & Tversky, 1970; Seale & Rapoport, 1997; 2000).

The classic secretary problem was introduced in the early 1950s. Since then, many efforts have been made to derive the optimal solution of the problem's variants in which one or more of its artificial underlying assumptions is relaxed (e.g., Bartoszynski & Govindarajulu, 1978; Bearden & Murphy, 2007; Bruss, 1984; Choe & Bai, 1983; Gilbert & Mosteller, 1966, reprinted 2006; Lee, Gregory, & Welsh, 2005; Mucci, 1973; Pressman & Sonin 1972; Seal & Rapoport, 2000; Smith, 1975; Stewart, 1981; Yang, 1974; Zwick, Rapoport, Lo, & Muthukrishnan, 2003). However, such solutions are often less elegant than the solution to the classic problem.

Optimal solutions are descriptively inadequate. Real people terminate the search earlier

than norm; individual differences are not accounted for (Bearden & Murphy, 2007). Several search strategies were suggested over the years to account for actual behavior in secretary problem tasks, some of them based on satisficing-type heuristics (Lee et al., 2005; Seal / Rapoport, 2000). Moreover, cognitive mechanisms are not well-specified in this literature. For example, the single-fixed-threshold heuristic (Lee et al., 2005) only describes the overall process (searching for options above a specific threshold) and is silent about where this threshold comes from. The work on stopping rules started as a mathematical problem, and so psychological work on processing, of which there are promising starts, is in its infancy. As sometimes happens in psychology, there are interesting overlaps between this work and meta-reasoning, although work on either research area has been done in isolation. We will take this up again in the sections on meta-reasoning.

Philosophical foundations of further deliberation

In order to arrive at a more realistic model, Douven (2002) started from BDT, but then tried to "concretize" (i.e., "de-idealize") it, in the sense of Nowak (1980), hoping thereby to obtain a decision-making model more predictive of actual human decision making. At the core of Douven's proposal is the notion of *value of further deliberation*. To describe it requires some stage-setting. First, Douven assumes that, for each decision-making problem, we can identify a time limit, that is, a time before which a decision *must* be made. This is related to what Cherniak (1986) evocatively dubbed "the finitary predicament." A limitless search for options is beyond us, and often counterproductive. Also in Douven's proposal, a decision-making problem that an agent faces at time t is associated with a set of options, S_t, which the agent considers choosing from at that time. Furthermore, "c_t" designates the agent's estimate at t of the cost of further deliberation about other possible options until the time limit for the decision-making problem at hand. We note that cost and time limit might be related, as we elaborate in the next section. And finally, where o is an option in S_t, p_t^o is the agent's degree of belief at t that by deliberating further about possible options till the time limit she will come to think of

some option o^* whose expected utility exceeds not just that of o, but exceeds the sum of o's expected utility and the cost of further deliberation c_t. Then the value of further deliberation at t is defined as $p_t^o/EU(o)$, for all $o \in S_t$; we use the notation $V_t(o)$ for this value. And the *overall* value of further deliberation at t, $V_t(S_t)$, is defined to equal the least value of further deliberation for any of the options in S_t, that is, $V_t(S_t) = V_t(o)$ for $o \in S_t$ such that $V_t(o) \leqslant V_t(o^*)$, for all $o^* \in S_t$.

Then Douven's (2002) proposal is to replace BDT with the theory consisting of these axioms:

Axiom 1 *The rationality of choosing option o at time t is a linearly decreasing function of $V_t(o)$.*

Axiom 2 *The rationality of choosing to deliberate further at t is a linearly increasing function of $V_t(S_t)$.*

Note that this account makes rational choice a graded concept: it can be more or less rational to choose an option, as the value of further deliberation and the chance of identifying a better option can be smaller or larger. Note further that, on the above account, expected utility is still crucially connected to the concept of rational choice. The account is also still broadly Bayesian. In fact, as Douven (2002, p. 320) shows, we can think of BDT as a kind of limiting case of the new account. The main difference with BDT is that rational choice is no longer *only* a matter of expected utility but also depends on whether we should be satisfied with the set of options before us or rather should deliberate further about other options that may be open to us, a factor that pre-theoretically matters as well to our verdict about the rationality of the choice that is made in the end, but that BDT pretends not to matter, by way of idealizing assumption. Typically, by making a theory more realistic it also becomes somewhat messier, which in turn often makes it harder to obtain formal results. Nonetheless, Douven (2002, Sect. 3) derives a number of theorems from Axioms 1 and 2 above that give the new account of rational choice clear empirical content. For example, it is shown that, given these axioms, the rationality of further deliberation at time t is inversely proportional to the rationality of choosing at t the or a best option in S_t, and also that for options o, o^* in S_t, if

$EU(o) = EU(o^*)$, then choosing o at t is no more nor less rational than choosing o^* at t. Whether such consequences of the theory are in accordance with people's rationality judgments and, more importantly, whether they are in accordance with people's choice behavior are open empirical questions.

To end this section, we mention that, on Douven's (2002) account of rational choice, behavior can be rational even if it does not select the optimal choice; it just will not be *fully* rational in that case. This is one reason why the account should at least appear somewhat attractive to the advocates of theories of bounded rationality.

Meta-reasoning: A psychological foundation of further deliberation

Like BDT, research in metacognition is silent regarding judgments and effort regulation when searching for alternative options. Classic metacognitive research has dealt with effort regulation involved in learning and retrieval from memory (see Bjork, Dunlosky, & Kornell, 2013, for a review). Ackerman and Thompson (2017) suggested a meta-reasoning framework to guide research regarding effort regulation when solving problems and reasoning challenges. Answering knowledge questions and solving problems involves search behaviors: a respondent should decide whether further search for a better solution is worthwhile. In both tasks, when there is certainty about the best response one could possibly provide, and this best response is available, search behavior stops with little deliberation. However, in many cases there is uncertainty whether a currently considered response is the best possible and whether it is good enough to be provided. This uncertainty has been associated before with search behaviors (see Mehlhorn et al., 2015, for a review). A common finding in the decision-making literature is that people do not search optimally (Zwick, Rapoport, Lo, & Muthukrishnan, 2003). For instance, people tend to rely on smaller samples than they should for finding the option that maximizes the potential utility (e.g., Bearden, Rapoport, & Murphy, 2006; Plonsky, Teodorescu, & Erev, 2015; see Wulff, Mergenthaler-Canseco, & Hertwig, in press, for a meta-analysis). Moreover, such samples are often biased due to metacognitive blindness to their origins (Fiedler, 2000). Reliance on small or biased samples is yet another manifestation of

unsatisfactory search behavior. The metacognitive conceptualization provides a tool box for delving into potential biases in subjective confidence regarding each considered response and when constructing alternatives.

Ackerman and Thompson (2017) presented Figure 1 and six research questions which are all relevant for considering how people generate relevant alternatives. In the current context, the first column in Figure 1 represents the cognitive search itself, which results in the alternatives to consider. The second column, of metacognitive monitoring, represents assessment whether the list of alternatives is satisfactory. The third column, of metacognitive control, represents decisions whether to continue searching for additional alternatives or cease searching.

The six research questions are these:

1. Reasoning and problem-solving processes extend over a period of time. How are these processes monitored? The first metacognitive aspect to consider when considering search for alternatives is the timing of monitoring, or judgments of chances that the list of alternatives includes the best option one could produce. As in other reasoning tasks, relevant judgments can take place before, during, and at the conclusion of alternative generation. Before starting, one may judge whether there are any alternatives which will be satisfactory, and maybe assess also how many alternatives can be relevant or worth considering. During alternative generation, one should consider the quality of the alternatives considered so far, and whether there is still a chance to find better alternatives. When deciding to stop searching for additional alternatives, it is still relevant to consider whether the current list is exhaustive or satisfactory and whether external help (e.g., searching the Internet or asking a peer) can yield additional potentially good alternatives.

2. What determines whether to continue, switch strategies, or terminate thinking about a problem? As response to each of the above judgments, people should decide how to continue. In both the decision-making and the reasoning literature, there are discussions of confidence as the guide by which people decide whether to stop investing effort (Lee, Newell, & Vandekerckhove, 2014). While classic models referred to people's stopping

criterion as a constant level of confidence regardless of the invested time, recent models have shown that in many cases, people are willing to compromise on their level of confidence as they invest longer in the task. However, decision-making models often assume that there are two given alternatives to choose from (e.g., Hawkins et al., 2015). Most models of reasoning do not consider the number or particulars of the considered alternatives. A leading model—the dual-process theory—discusses the conditions under which people provide the first solution that comes to mind (type 1 processes) versus those in which they engage in more effortful thinking (type 2 processes; see Evans & Stanovich, 2013, for a review). A metacognitive view on the decision whether to stop or continue to invest efforts, suggests the Feeling of Rightness (see Fig. 1) as a judgment that bridges the two processing types (Thompson, Prowse Turner, & Pennycook, 2011). Another model, the Diminishing Criterion Model (Ackerman, 2014), deals with stopping rules for the task as a whole as well. It is unique in including no limit on the number of considered options and two stopping criteria: a confidence threshold, which drops as a person invests longer in the task, and a time limit.

3. What cues do we rely on to monitor our reasoning? Overall, it is now well established that metacognitive monitoring of all types is based on heuristics cues that are generally reliable, but sometimes misleading (Koriat, 1997). For instance, fluency, the ease with which a response comes to mind, is a ubiquitous cue for many judgments, which underlies many biases (see Unkelbach & Greifeneder, 2013, for a review). Such cues may lead people to stop generating alternatives too early, either because they think they already found the best alternative or that they cannot find it, which both may be reliable or unreliable (over- or under-confidence).

4. How does understanding meta-reasoning contribute to understanding the processes that mediate reasoning and problem solving? Considering alternative generation from a metacognitive perspective can explain non-optimal search behaviors. This potential contribution is relevant for phenomena such as reliance on small samples, as mentioned above, but also learned helplessness (Teodorescu & Erev, 2014), the challenging balance between exploration and exploitation (Mehlhorn et al., 2015), the secretary problem

(Zwick et al., 2003) outlined above, and more.

5. How do individuals differ in their ability to assess their performance? We know very little regarding individual differences in reasoning, nonetheless for search for alternatives. Individual differences should be considered in terms of at least three different levels:

(i) general cognitive ability—variables such as intelligence, working memory, and attention span;

(ii) thinking dispositions—personality characteristics such as open mindedness, creativity, uncertainty orientation, and need for cognition;

(iii) cultural differences—the role of general social tendency such as conformism and individualism.

6. Can reasoning be improved by insights from meta-reasoning research? We know too little at this point about directions for improvement. The aim to improve alternative generation should be considered the target of the research in this domain.

Meta-reasoning and further deliberation

We are now in a position to draw the threads together and suggest an outline of a theory of further deliberation. As a reminder, here are the potential parameters that affect the search for alternatives, according to Douven (2002):

(1) How satisfactory does the agent find the existing option(s)?

(2) What is the agent's estimate, or prediction, of her chance of success in finding a better alternative?

(3) What is the cost of further search?

(4) What is the time limit for the search?

We find it striking that all four parameters reflect, in essence, meta-reasoning research questions, and can be operationalized as meta-reasoning processes—sometimes with additional help from the work on optimal stopping rules. As we will see, meta-reasoning can also be used to raise further research questions, and to lead to a rationality model. It should be evident at this stage that meta-reasoning monitoring and control process must be involved in all four of Douven's parameters. In the case of parameter (1), satisfaction

with the current options, it can be translated as a Feeling of Rightness (FOR) (Thompson et al., 2011). The higher the FOR, the less likely is further deliberation. Parameter (2), prediction of success, might be akin to judgment of solvability (JoS) in Ackerman and Thompson's (2017) model.

It is an open empirical question how these monitoring processes differ. We assume that different processes—mainly heuristic cues—may underlie the various judgments. Recent research shows that judgment of solvability and confidence differ at least in some underlying cues (e.g., Ackerman & Beller, 2017). Potential cues include conceptual fluency and positive affect: we might stop with a familiar option (a laptop by a known brand); or have nostalgic preference for, say, a house in our childhood's neighborhood.

The cues for monitoring the generation of alternatives are unknown. In order to find them, researchers should design manipulations that affect differently the quality of the alternatives and the judgments regarding their quality. Such dissociation between actual chance for success and judgments suggests a cue that underlies the judgment even when it generates a bias. For an example of such experimental dissociation, see Metcalfe and Finn (2008); for a methodological review, see Ackerman (2019).

Perhaps the most helpful insight coming from metareasoning is the understanding that Douven's parameters (3) and (4)—the cost of further search and the time limit—are inter-dependent rather than entirely seperate parameters. This is a monitoring issue related to the "stop?" question. Ackerman's (2014) Diminishing Criterion Model (DCM) is highly significant here. As mentioned above, the model includes a time limit. According to DCM, as time goes by, reasoners become increasingly willing to settle for less ambitious goals; in other words, they increasingly satisfice. We can say that the subjective cost of further deliberation is at least in part a function of the passage of time.

In the present context, the DCM model suggests that people come up with alternatives and judge the chance that they include the best available option. As they invest more effort in generating the options, they are willing to stop with lower confidence that the list is adequate. Moreover, it suggests that people set a maximum time they are willing to invest in generating the alternatives, in line with Douven's model. When this time

elapses, they stop producing alternatives, even if their confidence that the considered alternatives include the best option is low. If the Diminishing Criterion Model indeed generalizes to alternative generation, it may explain several non-optimal behaviors (see question 4 in the previous section).

There are surprising parallels here to the work on optimal stopping rules. As mentioned, current models describe three stages: data collection, generation of initial threshold, and threshold updating (Teodorescu, Sang, & Todd, 2018). For example, when the data collection stage ends could also be dependent on confidence thresholds. When the environment appears more difficult—for instance, because the variance between the options is large—the threshold of when to terminate the data collection stage might be higher, resulting in further data collection; however, it also means that the cost of data collection is higher, which should result in *less* data collection. Only empirical research can determine which way reasoners behave.

Another research question which Ackerman and Thompson (2017) highlight is the role of individual differences in meta-reasoning. One potentially relevant dimension of individual differences is one that derives from Simon's notion of bounded rationality: the tendency to satisfice versus optimize, or maximize. There are several such individual differences scales (Diab, Gillespi, & Highhouse, 2008; Nenkov, Morrin, Ward, Schwartz, & Hulland, 2008; Schwartz et al., 2002; Turner, Rim, Betz, & Nygren, 2012; Ďuriník, Procházka, & Cígler, 2018), but they all share a distinction between two different cognitive styles: Simon-style satisficing, as a cognitive style involving settling for just-good-enough choices, on the one side of the scale; and maximizing (which in this literature stands for optimizing, the painstaking search for the best choice), on the other side of the scale. Perhaps unsurprisingly, maximizers achieve more, but satisficers tend to be happier with what they achieved (Iyegar, Wells, & Schartz, 2006). The effect of having too many options is well documented (e.g., Iyengar & Lepper, 2000): people tend to be less happy in that kind of situation, and they are less inclined to act on the options. As an aside, this really begs the question. Not all optimizing is normative, but normative search does require more optimizing; that is, optimizing is necessary, if not sufficient, for normative

further deliberation. But apparently, this does not guarantee happiness; on the contrary, it might actually undermine it. So if normatively required search does not make you happier, what is the point?

The implications for rationality

Much of the rationality debate traditionally focused on BDT, and therefore on the choice between members of a given set. This is an instance of what Elqayam and Evans (2011) called "clear norms bias": the tendency to focus on research questions where there is a clear normative system, to the neglect of research questions where normative standards are less obvious. This might be why the rationality of further deliberation has received a lot less attention in the literature than the rationality of choice, even though the former is prerequisite for the latter—without selecting the set in the first place, there is nothing to choose *from*. Nevertheless, there is much we can learn about rationality from considering further deliberation.

The first, and highly important point, is that rationality can no longer be perceived as an all-or-nothing issue, as it is in standard BDT. Acts are more or less rational, as the value of further deliberation can be smaller or greater. This fits well with notions of rationality that allow a measure of relativism (e.g., Achourioti, Fugard, & Stenning, 2014; Skovgaard-Olsen, Kellen, Hahn, & Klauer, 2019; Stupple & Ball, 2014; Douven, 2019). We favor the framework of *grounded rationality* (Elqayam, 2012), a descriptivist framework which involves no recourse to normative standards, and moderate relativism. Grounded rationality takes over where bounded rationality leaves off: it adds contextualized, relative constraints, whilst sticking to Simon's idea that the benefits of computation should outweigh its costs. It fits nicely with the gradual notion of rationality emerging from our analysis of further deliberation. Our analysis shows that the working definition proposed in Elqayam (2012) requires a slight revision. We propose that *Behavior B is rational for agent A, in epistemic context E, to the extent that B facilitates achievement of A's goals within the computational constraints of E*.[1] Note that behavior

[1] The working definition proposed in Elqayam (2012) was: Behavior B is rational for agent A, in

is always relative to the agent and the context, within this framework.

This brings us to the importance of individual differences. Since different agents have different goals, rationality can no longer be considered a one-size-fits-all matter (see also Skovgaard-Olsen et al., 2019). Note that optimizers and satisficers have different second- and third-order goals, even when first-order goals seem the same. These different goals lead to different stopping rules. Furthermore, given these different goals, what is a rational stopping rule for an optimizer could be an irrational one for a satisficer, and vice versa.

In other words, further deliberation can be argued to be rational to the extent that it is sensitive to meta-reasoning parameters of monitoring and control, and individual differences. Meta-reasoning monitoring parameters include how satisfactory the current option(s) are, and what is the subjective prospect of success in finding better ones; meta-reasoning control parameters include first- and second-order goals, costs of further search, and time limit. Optimizing versus satisficing is one dimension of individual differences which is directly relevant to further deliberation, but there can be other relevant dimensions, such as thinking dispositions (Stanovich, 1999), and second- and third-order preferences (Stanovich, 2008).

Lastly and importantly: we outlined potential parameters of a processing model, with a particular emphasis on meta-reasoning processes. What we do not have yet, though, is a specific, detailed processing model that can predict behavior of further deliberation and can be used to derive experiments. Future work can include a range of methodologies, such as a combination of the secretary problem with meta-reasoning methods; and agent-based simulations in different environments.

In conclusion: How and when it is rational to postpone one's decision and first look for further alternatives is only partially covered by extant theories of rational decision making, including classic BDT, proposed solutions to the secretary problem, and theories of bounded rationality. In this chapter, we presented a possible theoretical model based on Douven's proposal and developments in work on the secretary problem and

---

epistemic context E, iff B facilitates achievement of A's goals within the computational constraints of E.

meta-reasoning. We support a gradual and context-sensitive model of grounded rationality. Further search is rational to the extent that it is sensitive to meta-reasoning monitoring and control parameters as well as relevant cognitive variability factors. These parameters can be used to extend empirical work and develop a processing model to predict behavior.

Acknowledgements

We thank Jonathan Evans for helpful comments on an earlier draft. We thank David Over for his major impact on the field of reasoning; for setting the stage, early on, for many of the subsequent debates in the field; and for many productive collaborations and interesting discussions over the years.

References

Ackerman, R. (2014). The diminishing criterion model for metacognitive regulation of time investment. *Journal of Experimental Psychology: General, 143*, 1349–1368.

Ackerman, R. (2019). Heuristic cues for meta-reasoning judgments: Review and methodology. *Psihologijske Teme, 28*(1), 1–20.

Ackerman, R., & Thompson, V. A. (2015). Meta-reasoning: What we can learn from meta-memory. In A. Feeney & V. A. Thompson (eds.), *Reasoning as memory* (pp. 164–178). Hove UK: Psychology Press.

Ackerman, R., & Thompson, V. A. (2017a). Meta-reasoning: Shedding metacognitive light on reasoning research. In L. Ball & V. A. Thompson (eds.), *International handbook of thinking and reasoning* (pp. 1–15). Hove UK: Psychology Press.

Ackerman, R., & Thompson, V. A. (2017b). Meta-reasoning: Monitoring and control of thinking and reasoning. *Trends in Cognitive Sciences, 21*, 607–617.

Achourioti, T., Fugard, A. J. B., & Stenning, K. (2014). The empirical study of norms is just what we are missing. *Frontiers in Psychology, 5*.

Baron, J. (2008). *Thinking and deciding* (4th ed.). Cambridge: Cambridge University Press.

Bearden, J. N., Rapoport, A., & Murphy, R. O. (2006). Sequential observation and selection with rank-dependent payoffs: An experimental study. *Management Science, 52*, 1437–1449.

Bearden, J. N., & Murphy, R. O. (2007). On generalized secretary problems. In M. Abdellaoui, R. D. Luce, M. J. Machina, & B. Munier (eds.), *Uncertainty and risk* (pp. 187–205). Berlin/Heidelberg: Springer.

Bjork, R. A., Dunlosky, J., & Kornell, N. (2013). Self-regulated learning: Beliefs, techniques, and illusions. *Annual Review of Psychology, 64*, 417–444.

Bruss, F. T. (1984). A unified approach to a class of best choice problems with an unknown number of options. *Annals of Probability, 12*(3), 882–889.

Cherniak, C. (1986). *Minimal rationality.* Cambridge MA: MIT Press.

Choe, K. I., & Bai, D. S. (1983). A secretary problem with backward solicitation and uncertain employment. Journal of Applied Probability, 20(4), 891–896.

Douven, I. (2002). The rationality of further deliberation. *Economics & Philosophy, 18*, 303–328.

Douven, I. (2019). The ecological rationality of explanatory reasoning. *Studies in History and Philosophy of Science*, in press.

Elqayam, S. (2012). Grounded rationality: Descriptivism in epistemic context. *Synthese, 189*, 39–49.

Elqayam, S., & Evans, J. St. B. T. (2011). Subtracting "ought" from "is": Descriptivism versus normativism in the study of human thinking. *Behavioral and Brain Sciences, 34*(5), 233–248.

Evans, J. St. B. T., & Over, D. E. (1996). *Rationality and reasoning.* Hove UK: Psychology Press.

Evans, J. St. B. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition advancing the debate. *Perspectives on Psychological Science, 8*, 223–241.

Fiedler, K. (2000). Beware of samples! A cognitive–ecological sampling approach to judgment biases. *Psychological Review, 107*(4), 659–676.

Gilbert, J. P., & Mosteller, F. (2006). Recognizing the maximum of a sequence. In

Selected Papers of Frederick Mosteller (pp. 355–398). New York: Springer.

Hawkins, G. E., Forstmann, B. U., Wagenmakers, E.-J., Ratcliff, R., & Brown, S. D. (2015). Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *Journal of Neuroscience, 35*, 2476–2484.

Iyengar, S. S., & Lepper, M. R. (2000). When choice is demotivating: Can one desire too much of a good thing?. *Journal of Personality and Social Psychology, 79*(6), 995–1006.

Koriat, A. (1997). Monitoring one's own knowledge during study: A cue-utilization approach to judgments of learning. *Journal of Experimental Psychology: General, 126*, 349–370.

Kuipers, T. A. F. (2001). *Structures in science: Heuristic patterns based on cognitive structures*. Dordrecht: Kluwer.

Lee, M. D., Newell, B. R., & Vandekerckhove, J. (2014). Modeling the adaptation of search termination in human decision making. *Decision, 1*, 223–251.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: Freeman.

Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., Gonzalez, C. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision, 2*, 191–215.

Metcalfe, J., & Finn, B. (2008). Evidence that judgments of learning are causally related to study choice. *Psychonomic Bulletin & Review, 15*, 174–179.

Mucci, A. G. (1973). On a class of secretary problems. *Annals of Probability, 1*(3), 417–427.

Nowak, L. (1980). *The structure of idealization*. Dordrecht: Reidel.

Presman, E. L., & Sonin, I. M. (1973). The best choice problem for a random number of objects. Theory of Probability & Its Applications, 17(4), 657–668.

Plonsky, O., Teodorescu, K., & Erev, I. (2015). Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychological Review, 122*, 621–647.

Simon, H. A. (1979). *Models of thought*. New Haven CT: Yale University Press.

Simon, H. A. (1982). *Models of bounded rationality.* Cambridge MA: MIT Press.

Smith, M. H. (1975). A secretary problem with uncertain employment. *Journal of Applied Probability, 12*(3), 620–624.

Skovgaard-Olsen, N., Kellen, D., Hahn, U., & Klauer, K. C. (2019). Norm conflicts and conditionals. *Psychological Review*, in press.

Stewart, T. J. (1981). The secretary problem with an unknown number of options. *Operations Research, 29*(1), 130–145.

Stanovich, K. E. (1999). *Who is rational? Studies of individual differences in reasoning.* Mahway NJ: Lawrence Erlbaum Associates.

Stanovich, K. E. (2008). Higher order preference and the master rationality motive. *Thinking & Reasoning, 14*, 111–127.

Stupple, E. J. N., & Ball, L. J. (2014). The intersection between descriptivism and meliorism in reasoning research: Further proposals in support of soft normativism. *Frontiers in Psychology, 5*.

Teodorescu, K., & Erev, I. (2014). Learned helplessness and learned prevalence: Exploring the causal relations among perceived controllability, reward prevalence, and exploration. *Psychological Science, 25*, 1861–1869.

Teodorescu, K., Sang, K., & Todd, P. M. (2018). Post-decision search in repeated and variable environments. *Judgment & Decision Making, 13*(5), 484–500.

Thompson, V. A., Prowse Turner, J. A., & Pennycook, G. (2011). Intuition, reason, and metacognition. *Cognitive Psychology, 63*, 107–140.

Unkelbach, C., & Greifeneder, R. (2013). A general model of [U+FB02]uency effects in judgment and decision making. In C. Unkelbach & R. Greifeneder (eds.), *The experience of thinking: How the fluency of mental processes influences cognition and behaviour* (pp. 11–32). Hove UK: Psychology Press.

Wulff, D. U., Mergenthaler-Canseco, M., & Hertwig, R. (2019). A meta-analytic review of two modes of learning and the description-experience gap. *Psychological Bulletin & Review*, in press.

Yang, M. C. (1974). Recognizing the maximum of a random sequence based on relative

rank with backward solicitation. *Journal of Applied Probability, 11*(3), 504–512.

Zwick, R., Rapoport, A., Lo, A. K. C., & Muthukrishnan, A. (2003). Consumer sequential search: Not enough or too much? *Marketing Science, 22*, 503–519.

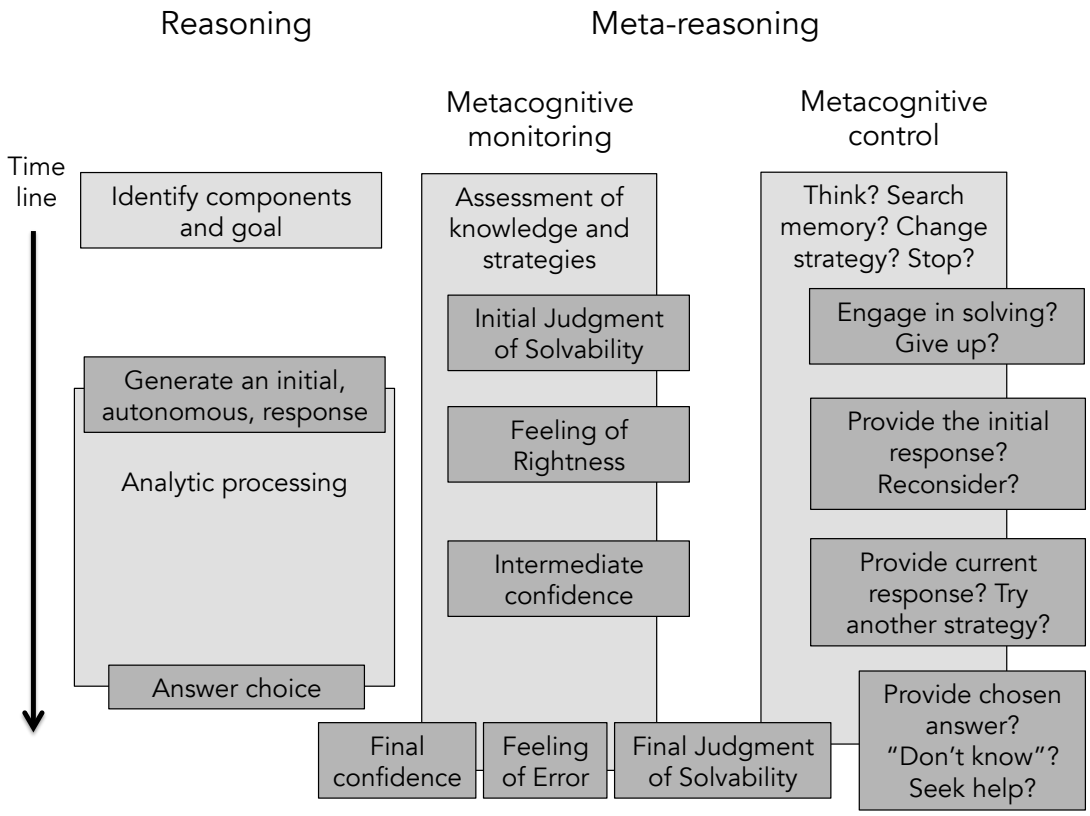

Figure 1. The meta-reasoning framework as put forward by Ackerman and Thompson (2017) for describing metacognitive processes that underlie effort regulation while solving problems or reasoning challenges.