Cognition 170 (2018) 102–122

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Original Articles

Cognitive costs of decision-making strategies: A resource demand decomposition analysis with a cognitive architecture

Hanna B. Fechner[a,b,*], Lael J. Schooler[d,a], Thorsten Pachur[c]

[a] Max Planck Institute for Human Development, Center for Adaptive Behavior and Cognition, Lentzeallee 94, 14195 Berlin, Germany
[b] University of Zürich, Department of Psychology, Binzmühlestrasse 14/22, 8050 Zürich, Switzerland
[c] Max Planck Institute for Human Development, Center for Adaptive Rationality, Lentzeallee 94, 14195 Berlin, Germany
[d] Syracuse University, Department of Psychology, 509 Huntington Hall, Syracuse, NY 13244, United States

ARTICLE INFO

Keywords:
Cognitive costs
Strategy
Decision making
Working memory
ACT-R
Computational modeling

ABSTRACT

Several theories of cognition distinguish between strategies that differ in the mental effort that their use requires. But how can the effort—or cognitive costs—associated with a strategy be conceptualized and measured? We propose an approach that decomposes the effort a strategy requires into the time costs associated with the demands for using specific cognitive resources. We refer to this approach as *resource demand decomposition analysis* (RDDA) and instantiate it in the cognitive architecture Adaptive Control of Thought–Rational (ACT-R). ACT-R provides the means to develop computer simulations of the strategies. These simulations take into account how strategies interact with quantitative implementations of cognitive resources and incorporate the possibility of parallel processing. Using this approach, we quantified, decomposed, and compared the time costs of two prominent strategies for decision making, take-the-best and tallying. Because take-the-best often ignores information and foregoes information integration, it has been considered simpler than strategies like tallying. However, in both ACT-R simulations and an empirical study we found that under increasing cognitive demands the response times (i.e., time costs) of take-the-best sometimes exceeded those of tallying. The RDDA suggested that this pattern is driven by greater requirements for working memory updates, memory retrievals, and the coordination of mental actions when using take-the-best compared to tallying. The results illustrate that assessing the relative simplicity of strategies requires consideration of the overall cognitive system in which the strategies are embedded.

1. Introduction

It is often assumed that people have different cognitive strategies at their disposal for achieving the tasks they face (e.g., Brown, 1995; Lemaire & Siegler, 1995; Reder, 1987; Taatgen & Anderson, 2002; van Rijn, van Someren, & van der Maas, 2003). This notion has been particularly influential in research on judgment and decision making (e.g., Beach & Mitchell, 1978; Gigerenzer, Hertwig, & Pachur, 2011; Payne, Bettman, & Johnson, 1993). Specifically, it has been proposed that decision strategies differ in terms of the cognitive costs that they impose on the decision maker and that people adjust their strategy selection accordingly (e.g., Gigerenzer, Todd, & the ABC Research Group, 1999; Johnson & Payne, 1985; Payne et al., 1993; Shah & Oppenheimer, 2008; Shugan, 1980; see also Gray, Sims, Fu, & Schoelles, 2006). Consistent with this view, strategies that are presumed to be simpler are more frequently used under conditions in which cognitive capacities are constrained, even when they can be associated with lower accuracy (e.g., Ford, Schmitt, Schechtman, Hults, & Doherty, 1989; Horn, Pachur, & Mata, 2015; Pachur & Hertwig, 2006; Payne, Bettman, & Johnson, 1988; Rieskamp & Hoffrage, 2008).

But what is a simple strategy—or more generally, what cognitive costs are associated with the use of a strategy, and how can they be determined? A prominent proposal to quantify the complexity of decision strategies was provided by Payne, Bettman, and Johnson (e.g., 1993): Their approach has been to decompose strategies into cognitive operations or processing steps, called *elementary information processes* (EIPs; e.g., Bettman, Johnson, & Payne, 1990). The central idea is to represent a strategy as a sequence of mental events—such as reading a piece of information into short-term memory, multiplying a probability and a payoff, or comparing the values of two alternatives on an attribute—and to count (and weight) the basic processing steps that a strategy requires. The more EIPs and the more frequently they are used, the greater the cognitive demands that a strategy imposes on the decision maker. Measures of strategy complexity obtained with this

* Corresponding author at: Max Planck Institute for Human Development, Center for Adaptive Behavior and Cognition, Lentzeallee 94, 14195 Berlin, Germany.
E-mail addresses: hfechner@mpib-berlin.mpg.de (H.B. Fechner), lschoole@syr.edu (L.J. Schooler), pachur@mpib-berlin.mpg.de (T. Pachur).

http://dx.doi.org/10.1016/j.cognition.2017.09.003
Received 8 December 2016; Received in revised form 6 September 2017; Accepted 7 September 2017

approach have been shown to be consistent with differences between strategies in response times and subjective reports of cognitive effort (Bettman et al., 1990), as well as the strategies people select in different decision environments (Payne et al., 1988).

Despite the success and explanatory power of the EIP framework, the question remains whether counting the number of operations (that are often arithmetic rather than cognitive in nature) provides an appropriate model for a psychological concept of cognitive costs and complexity. Without doubt, the number of steps necessary to acquire and evaluate information bears on the complexity of a decision strategy. Yet the cognitive effort required is also a function of the constraints of the cognitive resources that are involved in the execution of that strategy. For instance, the cognitive system often has to create and update representations of decision-relevant information, but there are limits to how much information can be stored and manipulated at one time and how fast this can be accomplished. Moreover, to evaluate cognitive effort it is important to know how the different cognitive resources interact; while some processing operations can be accomplished only serially, others can be executed in parallel. Finally, in many real-world situations decisions take place in the context of other cognitive activities (such as simultaneously driving a car and talking to another person); therefore, it is also necessary to consider how the cognitive resources are affected in multitasking situations, where some operations might be processed serially, whereas others can be processed in parallel. However, due to the strict sequential nature of its processing steps the EIP framework cannot handle parallel processing.

In this article, we propose what we call *resource demand decomposition analysis* (RDDA) to quantify the complexity of decision strategies. RDDA extends previous approaches to decompose decision strategies into cognitive operations, such as the EIP approach, in two main respects. First, it integrates strategies with theory-based implementations of cognitive resources (i.e., visual and auditory perception, working memory updating and memory retrieval, motor responses, and the coordination of the underlying mental actions) by grounding them in the cognitive architecture ACT-R (Adaptive Control of Thought–Rational; Anderson, 2007). This makes it possible to attribute cognitive costs incurred by a strategy to the specific underlying cognitive resources. Second, ACT-R is explicit about how cognitive resources interact with each other, that is, which cognitive processing operations can be executed in parallel and which require serial processing. Thus, this framework can, for instance, handle multimodal sensory input that is processed in parallel, but it also highlights the processing bottlenecks that can strongly impact the cognitive costs of behavior. Moreover, ACT-R has been used as the foundation for the threaded cognition theory of multitasking (Salvucci & Taatgen, 2008). Thus, implementing decision strategies within the ACT-R framework makes it possible to quantify the cognitive costs of strategies while taking into account serial and parallel processing—in both single- and multitasking situations.

We focus on the analysis of the cognitive costs of two paradigmatic examples of decision strategies that are assumed to differ in terms of in their complexity. The compensatory *tallying* strategy (Gigerenzer & Goldstein, 1996) involves examining all attributes and integrating those in favor of each decision alternative before making a decision; thereby one attribute can compensate for another. The noncompensatory *take-the-best* strategy (Gigerenzer & Goldstein, 1996) involves inspecting attributes sequentially and making a decision on the basis of a single attribute while ignoring the rest of the attributes; those ignored attributes cannot compensate for attributes that triggered a decision. Take-the-best has properties that are assumed to reduce cognitive effort (Shah & Oppenheimer, 2008): It typically examines fewer attributes of decision alternatives than compensatory strategies, and it also foregoes integrating information across attributes. As a consequence, take-the-best has been considered to be cognitively less effortful than compensatory strategies (e.g., Gigerenzer et al., 1999; Mata, Schooler, & Rieskamp, 2007; Rieskamp & Hoffrage, 2008; see also Payne et al., 1993). To compare take-the-best and tallying in terms of their cognitive costs, we tested the strategies in the context of concurrent activities that imposed different levels of cognitive demand. This allowed us to identify the circumstances under which take-the-best is simpler than tallying, but also to identify instances when it incurs cognitive costs that exceed those of tallying. Implementing the strategies in ACT-R thereby also addresses concerns that previous, less comprehensive descriptions of take-the-best might ignore hidden costs in executing the strategy (e.g., Dougherty, Franco-Watkins, & Thomas, 2008; Newell, 2005). For instance, the RDDA makes explicit the costs for selecting and acquiring attribute information during the decision process.

In the following, we describe tallying and take-the-best and discuss empirical evidence hinting at the cognitive costs associated with their use. We then give an overview of how we implemented these strategies and a concurrent task as computational models in ACT-R. We present simulation results from these models for response times and accuracy in strategy execution and concurrent activities. To foreshadow a main result, the simulations indicate that under increasing cognitive demands, the response times of take-the-best can exceed those of tallying; a similar pattern also emerged in an empirical study. RDDA is then used to attribute response time differences of the strategies to differences in the use of underlying cognitive resources. We conclude by discussing methodological and theoretical implications of our findings and highlight future directions for research with the ACT-R-based RDDA approach.

2. Compensatory vs. noncompensatory decision strategies

A prominent idea for how people make decisions between alternatives has been that the values on all attributes are weighted and summed for each alternative and the resulting scores are then compared (e.g., Keeney & Raiffa, 1993; Payne et al., 1993). Such a weighted additive strategy often leads to high decision accuracy (e.g., Payne et al., 1993) but it also requires a substantial number of cognitive operations. Simplifications have been proposed that, for instance, only involve summing the attribute values without weighting (equal-weight linear models; e.g., Dawes, 1979; Einhorn & Hogarth, 1975). In this article, we focus on the strategy tallying (Gigerenzer & Goldstein, 1996). Tallying involves the examination of all attribute values in whichever order they are available, and only integrates those values that favor each alternative. The amount of mental computation is therefore reduced relative to a weighted additive or equal-weight strategy. Several studies have found that tallying provides a good description of people's decisions (e.g., Bröder & Gaissmaier, 2007; Pachur & Aebi-Forrer, 2013; Pachur & Marinello, 2013; Platzer & Bröder, 2012).

In contrast to compensatory strategies such as tallying, noncompensatory strategies typically allow one to ignore part of the information and often to make a decision based on a single attribute. In this article we focus on the strategy take-the-best (Gigerenzer & Goldstein, 1996). Take-the-best involves inspecting attributes sequentially in the order of their importance and comparing alternatives on the respective attributes. As soon as an attribute discriminates between the alternatives, search is stopped and no further attributes are inspected. As a consequence, with take-the-best one has to inspect different numbers of attributes, depending on the attribute patterns of the alternatives. Search can be very restricted (if an attribute of high importance discriminates) or more extensive (if only an attribute of lower importance discriminates; e.g., Bröder & Gaissmaier, 2007; Khader et al., 2011).

Take-the-best and tallying share central characteristics with other noncompensatory and compensatory strategies, respectively, that have been proposed in the literature (e.g., Brandstätter, Gigerenzer, & Hertwig, 2006; Pachur, Hertwig, & Rieskamp, 2013; Payne et al., 1988; Thorngate, 1980; Tversky, 1972). The conclusions of our analyses are thus likely to hold, to some extent, beyond the specific