

Automatic activation of attribute knowledge in heuristic inference from memory

Patrick H. Khader · Thorsten Pachur · Kerstin Jost

Published online: 7 November 2012
© Psychonomic Society, Inc. 2012

Abstract In memory-based decision making, people often rely on simple heuristics such as take-the-best (TTB; Gigerenzer & Goldstein, *Psychological Review*, 103, 650–669, 1996), which processes information about the alternatives sequentially and stops processing as soon as a decision can be made. In this article, we examine the memory processes associated with TTB—in particular, to what degree the selective memory retrieval of relevant information required by TTB is accompanied by automatic activation of associated but irrelevant information. To address this question, we studied the fan effect (Anderson, *Cognitive Psychology*, 6, 451–474, 1974), which is assumed to arise from automatic spread of activation, in inferences from memory. Participants were instructed to use TTB when making decisions about objects on the basis of previously memorized attribute information. Both the number of attributes required by TTB and the number of attributes associated with an object (i.e., fan level) were

manipulated. As it turned out, response times and the correct execution of TTB were a function not only of the number of required attributes, but also of the number of associated attributes. This suggests that information that TTB “ignores” is nevertheless activated in memory.

Keywords Spreading activation · Decision making · Long-term memory · Heuristics · Fan effect

One of the fundamental tenets in cognitive psychology is that the capacity for controlled information processing is limited (e.g., Broadbent, 1958; Miller, 1956). As a consequence, much of cognition relies on “boundedly rational” mechanisms (Simon, 1990). In decision making, for instance, a prominent example of a boundedly rational decision strategy is the take-the-best (TTB; Gigerenzer & Goldstein, 1996) heuristic. When asked whether, say, Barcelona or Madrid has more inhabitants, TTB predicts that you start by comparing the cities on the attribute¹ that is most predictive of city size—say, whether the city is the national capital. Because this attribute discriminates (i.e., Madrid is the capital), search is stopped, and Madrid is judged to be larger. If the most valid attribute does not discriminate (e.g., when Barcelona and Seville are compared), search will extend to the second-most valid attribute (e.g., whether the city has an international airport) until a discriminating attribute is found (for a critical discussion, see Dougherty, Franco-Watkins, & Thomas, 2008).

Due to its stopping rule, TTB often inspects only part of the attribute information. It has therefore been hypothesized that people rely on the heuristic, in particular, when information costs are high, such as in memory-based decisions

This work was supported by Grant KH235/1-1 of the German Research Foundation (DFG). We thank Katharina Dobs for programming the stimulus presentation, Martin Vogel, Katharina Schnabel, and Lilian A. E. Weber for data acquisition, Michael Dougherty for helpful comments, and Laura Wiles for editing the manuscript.

P. H. Khader (✉)
Department of Psychology,
Ludwig Maximilians University Munich,
Munich, Germany
e-mail: Khader@lmu.de

T. Pachur
Cognitive and Decision Sciences,
University of Basel,
Basel, Switzerland
e-mail: thorsten.pachur@unibas.ch

K. Jost (✉)
Institute of Psychology, RWTH Aachen University,
Aachen, Germany
e-mail: Jost@psych.rwth-aachen.de

¹ In the context of inference tasks, the term “cue” is sometimes preferred. However, given that “cue” has a specific and different meaning in memory research, here we use the term “attribute” to avoid confusion.

(Gigerenzer & Todd, 1999). To test this hypothesis, Bröder and Schiffer (2003) developed a paradigm in which participants are first taught to associate objects (e.g., city names) with various attribute values (e.g., whether the city has an international airport or whether it has a successful soccer team). Subsequently, they are asked to make inferences about the objects (e.g., which of two cities is larger) on the basis of the attribute values stored in long-term memory (LTM). Results show that people indeed use TTB more when attribute information has to be retrieved from LTM (*inference from memory*), as compared with when it is provided on a computer screen (*inference from givens*).

But what exactly are the memory processes accompanying the use of TTB in inferences from memory? According to Gigerenzer and Goldstein (1996), with TTB, “search extends through only a portion of the total knowledge in memory” (p. 653), and TTB thus “reduces search in memory considerably” (p. 658). One possibility is, therefore, that information ignored by the heuristic is not activated in memory at all. Supporting this possibility, people using TTB have been found to make decisions faster the fewer the number of attributes that have to be inspected (Khader et al., 2011) and to be faster than users of a strategy that always considers all attributes (Bröder & Gaissmaier, 2007). On the other hand, TTB’s notion of sequential and, thus, controlled memory retrieval appears inconsistent with existing memory theories. The well-known ACT model of memory proposed by Anderson and colleagues (e.g., Anderson & Reder, 1999), for instance, assumes that the presentation of a stimulus leads to an automatic spread of activation to all information associated with it. Accordingly, another possibility is that having learned, say, that a city has an international airport as well as a successful soccer team, presentation of the city activates values on both learned attributes—even if only one of them is required to make a decision. The goal of this article is to examine the occurrence of such automatic LTM activation of attribute knowledge when TTB is used.

To do so, we connect the inference-from-memory paradigm (Bröder & Schiffer, 2003) used in decision research with the fan effect paradigm (Anderson, 1974) used in memory research. To study activation processes during memory retrieval, Anderson (1974) developed a task in which participants had to study sentences stating that a specific person was in a specific location (e.g., “The hippie is in the park”) and manipulated the level of *associative fan* for each person and location, defined as the number of locations that were associated with a person and vice versa. For instance, the person “hippie” would have a fan level of 2 if participants learned not only that the hippie was in the park, but also that the hippie was in the church. After the study phase, participants were presented with test sentences and had to determine as quickly as possible whether or not each had been part of the study set. When judging a previously studied sentence, people took longer and made more

errors the greater the number of associations that had been learned about the respective person or location. This phenomenon is called the *fan effect* and is assumed to arise from increased retrieval interference due to automatic spread of activation. Specifically, the larger the number of facts that are associated with a concept, the less activation is spread to each fact (Anderson & Reder, 1999), decreasing their accessibility and, thus, producing longer response times (RTs), as well as less accurate responses.

In this study, we examined the occurrence of the fan effect when people make memory-based decisions with TTB. Participants first learned to associate objects with various attribute information and were then instructed to use TTB² to make inferences about the objects on the basis of the memorized attributes—as in the inference-from-memory paradigm. As in the fan effect paradigm, we manipulated the number of attributes associated with an object. In addition, we manipulated the number of attributes that were required (according to TTB) to make a decision. We expected that, as had been observed in previous studies (Bröder & Gaissmaier, 2007; Khader et al., 2011), RTs would be faster the fewer the number of attributes TTB had to inspect. Our rationale concerning the manipulation of fan level was as follows: If participants take longer to make a decision about objects that have a higher number of associations than about objects that have only a few (holding the number of required attributes constant), then this is an indication that activation spreads also to those attributes that are irrelevant for the decision. By connecting research on memory-based decision making to the fan effect, we contribute to bridging the often-lamented conceptual and methodological gap between decision research and cognitive psychology in general (e.g., Weber, Goldstein, & Barlas, 1995). Note that we do not question the status of TTB as a model of memory-based decision making. Rather, our goal is to shed light on the underlying memory processes when TTB is used and, thus, inform possible model extensions of TTB.

Method

Participants

Twenty-seven students at the University of Marburg participated and received either money or course credits for compensation. Because we were interested in the application of TTB, we excluded 4 participants whose responses in the

² We examined instructed, rather than spontaneous, use of TTB because, without instruction, people may switch between different strategies. Since we were interested in knowledge activation when TTB is used, we chose investigating TTB in its purest form.

decision task (see below) deviated from the decisions predicted by TTB on more than 30 % of the trials. The final sample consisted of 23 participants (16 female; mean age = 21.7 years, $SD = 2.2$).

Materials

We used the same stimuli as in Khader et al. (2011). Participants first had to memorize attribute information about 24 companies (represented as pronounceable five- or six-letter nonwords, such as GNINT, NARCH, and CLEEF) and subsequently decide which of two companies would be more successful in the next year. Four binary attributes were used: the geographical location of the company, the manager of the company, the product the company produced, and the color of the products the company produced (for details, see Khader et al., 2011; example stimuli are shown in Fig. 1).

Procedure

Participants were first presented with several learning phases in which they were taught (1) associations between company names and attributes, (2) use of TTB, (3) the attribute hierarchy (i.e., how important the different attributes were for predicting a company's success), and (4) the attribute direction (i.e., which of the two locations, managers, objects, or colors, respectively, indicated higher success of a company).

For each participant, only three attributes were used, randomly drawn from the four attributes. In order to create different levels of associative fan (*fan level*), we manipulated the number of attributes that were associated with a













company (one, two, or three attributes). Specifically, eight companies were associated with information on the most important attribute only (fan level 1), another eight companies were associated with information on the most and the second-most important attributes (fan level 2), and another eight companies were associated with information on all three attributes (fan level 3). The mapping of the company names to the attribute patterns, as well as the constellation of the attribute patterns, was random for each participant. Assignment of attribute values to companies was pseudo-randomized, with the restriction that, for each fan level, every attribute value and combination of attribute values occurred equally often. For example, assume that the three attributes are “products,” “colors,” and “managers” and that the “products” attribute is the most important one. Then, for half of the companies with only one association (i.e., fan level 1), an association with cups would be learned, and for the other half, an association with plates would be learned. The same procedure was used to create the attribute patterns for companies of fan levels 2 and 3 (i.e., for which two and three attributes were learned, respectively). As a consequence, the combinations of attributes that were assigned to the companies were completely balanced.

Learning the associations between companies and attributes

To associate the company names with attribute information, participants were presented with a company name and the two picture stimuli of a specific attribute (e.g., two faces) and had to indicate the correct stimulus representing the attribute value of the company (followed by feedback). After all attributes of a single company (i.e., one, two, or

Fig. 1 Overview of the experimental conditions, with examples of companies and stimuli used as attributes in the decision experiment.

Participants learned to associate each company with one, two, or three out of four possible attributes. The attributes could be a spatial position (the location of the company), a face (the manager of the company), an object (the product the company made), or a color (the color of the product). In the shown example, the most important attribute is “object,” the second-most important attribute is “location,” and the least important attribute is “face”

		Number of required attributes		
		1	2	3
Fan Level	3	  <div> <div>X</div> <div>GNINT + NARCH</div> <div>X</div> </div>	  <div> <div>X</div> <div>GNINT + CLEEF</div> <div>X</div> </div>	  <div> <div>X</div> <div>GNINT + BLOOR</div> <div>X</div> </div>
	2	  <div> <div>X</div> <div>HAXOR + FILZEC</div> <div>X</div> </div>	  <div> <div>X</div> <div>HAXOR + SLARB</div> <div>X</div> </div>	
	1	  <div> <div>HELTE + TINKS</div> </div>		

three attributes, depending on the fan level) had been learned (correct responses given twice in a row), the next company was presented, and so on. After all companies had been presented, the cycle started again (with the companies and attributes presented in a new random order) until perfect performance was achieved twice in a row.

Learning TTB

Participants returned to the lab on the following day. First, they freshened up the learned attribute knowledge until again reaching perfect performance. Next, participants were trained to make decisions using TTB in a fictitious applicant-selection scenario, in which they were instructed to indicate (on the basis of three attributes) which candidate, according to TTB, was more suitable (followed by feedback; for details, see Khader et al., 2011).

Learning the attribute hierarchy and the attribute direction

Next, participants learned how important the different attributes were for deciding which of two companies would be more successful. Each of the three attributes (on which the participant had learned information about the companies) was presented separately, and participants had to indicate its importance (10, 9, or 8, with higher numbers indicating higher importance), followed by feedback. This procedure was repeated until correct responses had been given to each of the three attributes three times in a row. The attribute hierarchy was varied randomly across participants. Next, participants learned the attribute direction—that is, which value of a given attribute indicated higher success of a company. For this purpose, participants were presented with the two stimuli of each attribute (e.g., a cup and a plate for the “products” attribute) and had to select which stimulus indicated success (followed by feedback). The complete procedure was repeated until participants had given the correct response three times in a row.

Decision task

In the decision task, participants were shown pairs of company names (e.g., TIRCH vs. SNILM) and were asked to decide, using TTB, which company would be more successful in the next year by pressing the respective key. Note that for giving a correct response, participants had to recall the relevant attribute information about the companies learned in the initial learning phase. Participants were instructed to decide as quickly and as accurately as possible. The pair comparisons were constructed such that only companies of the same fan level (i.e., ones that were associated with the same number of attributes) were compared and for which TTB led to an unambiguous decision (i.e., where one of the

attributes discriminated). In the experimental design, two factors were thus partially crossed: “fan level” and “number of required attributes” (see Fig. 1). For pairs of companies of fan level 1 (i.e., for which only one association—namely, the most important attribute—was learned), 16 pairs could be constructed (by pairing the four companies with one attribute value with the four companies with the other attribute value). For pairs of companies of fan level 2 (i.e., for which two associations—namely, the most and the second-most important attribute—were learned), 16 pairs could be constructed for which only the most important attribute had to be considered (i.e., it discriminated); for another 8 pairs, the second-most important attribute also had to be considered. For pairs of companies of fan level 3, 16, 8, and 4 pairs could be constructed in which the most important attribute only, the two most important attributes, or all three attributes had to be considered, respectively. In order to keep the number of trials comparable across the different levels of “number of required attributes”, company pairs that required the retrieval of two attributes were presented twice, and company pairs that required the retrieval of three attributes were presented three times. Note that trial numbers were constant across the fan levels for each level of the “number of required attributes” factor. Overall, participants were presented with 92 trials, which were presented in four blocks of 23 trials each. After each block, participants could take a short break if necessary.

Results

The RTs in the decision task are shown in the upper panel of Fig. 2. For the RT analysis, incorrect decisions (i.e., where participants did not make the decision as predicted by TTB) were excluded (on average, 5.34 %, $SD = 5.48$). All reported t tests were one-tailed (testing for directional research hypotheses); the significance level was set to $\alpha = .05$.

Replicating earlier studies (Bröder & Gaissmaier, 2007; Khader et al., 2011), participants responded faster the fewer the attributes required by TTB to make a decision. A contrast analysis indicated a significant linear trend across the three levels of number of required attributes for companies of fan level 3 (i.e., the light-gray bars in Fig. 2), $F(1, 22) = 127.10$, $p < .001$, η_p^2 (partial eta squared; i.e., $SS_{\text{between}}/SS_{\text{total}} + SS_{\text{error}}$) = .85 [a quadratic trend was not significant, $F(1, 22) = 2.88$, $p = .104$, $\eta_p^2 = .12$]. Participants responded faster when TTB required the retrieval of only one rather than two attributes, $t(22) = 7.07$, $p < .001$, (Cohen’s) $d = 1.43$, and when TTB required the retrieval of two rather than three attributes, $t(22) = 4.74$, $p < .001$, $d = 0.65$. A corresponding pattern was found for companies of fan level 2, where participants responded faster when TTB required the

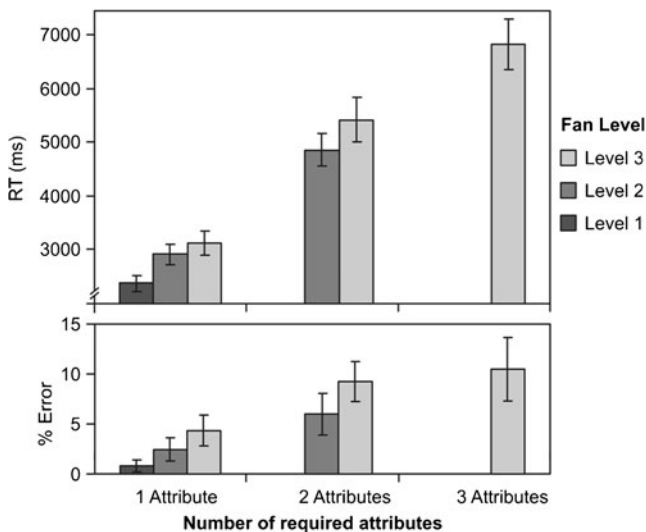


Fig. 2 Response times (RTs) and error rates in the decision task (the error bars are standard errors), showing that both measures increase not only with the number of required attributes, but also with the number of attributes associated to a company (i.e., fan level)

retrieval of only one rather than two attributes (i.e., the medium-gray bars in Fig. 2), $t(22) = 11.36$, $p < .001$, $d = 1.58$.

More important, RTs were also influenced by fan level. That is, the larger the number of attributes that were associated with the two companies in a pair, the longer it took participants to indicate which company was more successful. For comparisons requiring the retrieval of only one attribute, planned contrasts indicated a highly significant linear trend from fan level 1 to fan level 3, $F(1, 22) = 16.18$, $p = .001$, $\eta_p^2 = .42$ [a quadratic trend was not significant, $F(1, 22) = 2.34$, $p = .140$, $\eta_p^2 = .10$]. Participants responded more slowly for companies of fan level 3 than for companies of fan level 2, $t(22) = 1.77$, $p = .046$, $d = 0.21$, for which, in turn, they responded more slowly than for companies of fan level 1, $t(22) = 3.24$, $p = .002$, $d = 0.68$. For trials requiring the retrieval of two attributes, participants responded more slowly for companies of fan level 3 than for companies of fan level 2, $t(22) = 2.04$, $p = .027$, $d = 0.31$.

The lower panel of Fig. 2 shows a similar pattern for the error rates. A contrast analysis indicated a marginally significant linear trend across the three levels of number of required attributes (computed for the companies of fan level 3), $F(1, 22) = 3.16$, $p = .089$, $\eta_p^2 = .13$; a test for a quadratic trend was not significant, $F(1, 22) = 1.16$, $p = .293$, $\eta_p^2 = .05$. Participants made fewer errors in comparisons requiring only one rather than two attributes, $t(22) = 2.05$, $p = .027$, $d = 0.57$, but there was no difference between comparisons requiring two versus three attributes, $t(22) < 1$. Participants also made fewer errors in comparisons requiring one rather than

two attributes for companies of fan level 2, $t(22) = 2.42$, $p = .012$, $d = 0.44$. In addition, error rates were affected by fan level. In comparisons for which one attribute was required, there was a marginally significant linear trend, $F(1, 22) = 4.08$, $p = .056$, $\eta_p^2 = .16$, but no quadratic trend, $F(1, 22) < 1$. Post hoc comparisons between fan levels showed no significant differences. However, in comparisons for which two attributes were required, participants made significantly fewer errors for companies of fan level 2 than for companies of fan level 3, $t(22) = 1.74$, $p = .049$, $d = 0.33$. Importantly, note that these analyses of the error rates also rule out the possibility that the fan effect resulted from a speed–accuracy trade-off: participants did not make more—but rather fewer—errors the faster they responded.

Discussion

We used the fan effect as an indicator of automatic activation in memory to examine whether the selective memory retrieval required by TTB is accompanied by automatic activation of associated but irrelevant information. Replicating previous studies on instructed (Khader et al., 2011) and spontaneous (cf. Bröder & Gaissmaier, 2007) use of TTB, people decided faster the fewer the number of attributes that had to be retrieved. Most important, we also found a fan effect; that is, RTs and error rates were also affected by the number of attributes associated with the companies (i.e., fan level). For instance, even when people had to inspect only one attribute to make a decision, they took longer and made more errors using TTB when the companies were associated with, say, two attributes than when the companies were associated with only one attribute. As was proposed by Anderson and Reder (1999), these results can be explained by assuming that a larger number of associated attributes leads to decreased activation spreading to each attribute, creating retrieval interference by making each attribute less accessible (or less discriminable from other attributes). Our results indicate that the use of TTB is accompanied by an automatic activation of irrelevant attributes.³ We were able to fully replicate the present results in an independent study (using functional magnetic resonance imaging; Khader et al., unpublished data), suggesting that the automatic activation of irrelevant attributes is a robust regularity in memory-based decision making with TTB.

³ Alternatively, it is possible that, always, all attributes are automatically retrieved and the differences between fan levels reflect the suppression of irrelevant attributes. However, Anderson and Reder (1999) found no evidence that the fan effect is caused by suppression.

Our demonstration of a fan effect when TTB is used shows that the memory processes associated with TTB may not operate in the controlled and sequential fashion suggested by the process description of the heuristic. Although the decision maker is clearly able to focus on only a subset of attribute information, due to the architecture of the cognitive system, irrelevant attribute information is activated as well and can exert an influence on the time it takes to make a decision.

We should point to an alternative explanation for the RTs for comparisons involving companies of fan level 1 and requiring one attribute being faster than those involving companies of fan levels 2 and 3.⁴ Since comparisons between companies of fan level 1 always discriminated on the first attribute, it is possible that (some) participants noticed that retrieving the attribute value of only one of the companies would be sufficient to make a correct decision. The apparent fan effect between fan level 1 and higher fan levels would then be due to the retrieval of fewer attribute values rather than less retrieval interference. However, this alternative account cannot explain the RT difference between fan levels 2 and 3 (which were replicated by Khader et al., in preparation). Moreover, even if knowledge about a company's fan level contributed to the effect of fan level, it is not implausible to assume that this knowledge also feeds on an automatic activation of these associations.

Our results have several theoretical and conceptual implications. First, they suggest that the principle of activation spread in memory networks—previously shown mainly in the context of a recognition task—also holds for decision making requiring recall. Second, the results highlight a possible boundary condition for an effective use of TTB. Specifically, executing TTB correctly will be more difficult (in terms of longer RTs and less accurate responses) in domains where the objects are embedded in a rich network of attribute knowledge. Note, however, that our evidence for automatic activation does not mean that TTB does not reduce retrieval efforts. When, according to TTB, fewer attributes were required, people responded more quickly and more accurately (Fig. 2), and the size of this effect was considerably larger than the size of the fan effect. Moreover, in a neuroimaging study, Khader et al. (2011) found that a boost in neural activation in the representation area of an attribute—arguably reflecting retrieval effort—occurred only when an attribute was relevant for a decision.

Our finding of a large RT effect of the number of required attributes and a smaller effect of fan level could be accommodated within recent models of working memory (Cowan,

1995; Oberauer, 2002). These models assume different activation levels of LTM contents: the currently activated portion of LTM and the subset of activated LTM that is the focus of attention. From this perspective, the complete attribute information associated with the decision options can be seen as the currently activated set of LTM representations, whereas only the attributes that are actually required for the decision are in the focus of attention. This important characteristic of the retrieval dynamics involved in TTB should be considered in studies and models of memory-based decision making.

References

- Anderson, J. R. (1974). Retrieval of propositional information from long-term memory. *Cognitive Psychology*, 6, 451–474.
- Anderson, J. R., & Reder, L. M. (1999). The fan effect: New results and new theories. *Journal of Experimental Psychology: General*, 128, 186–197.
- Broadbent, D. E. (1958). *Perception and communication*. London: Pergamon Press.
- Bröder, A., & Gaissmaier, W. (2007). Sequential processing of cues in memory-based multiattribute decisions. *Psychonomic Bulletin and Review*, 14, 895–900.
- Bröder, A., & Schiffer, S. (2003). Take The Best versus simultaneous feature matching: Probabilistic inferences from memory and effects of representation format. *Journal of Experimental Psychology: General*, 135, 207–236.
- Cowan, N. (1995). *Attention and memory: An integrated framework*. New York: Oxford University Press.
- Dougherty, M. R., Franco-Watkins, A. M., & Thomas, R. (2008). Psychological plausibility of the theory of probabilistic mental models and the fast and frugal heuristics. *Psychological Review*, 115, 199–213.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: models of bounded rationality. *Psychological Review*, 103, 650–669.
- Gigerenzer, G., & Todd, P. M. (1999). Fast and frugal heuristics: The adaptive toolbox. In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 3–34). New York: Oxford University Press.
- Khader, P. H., Pachur, T., Meier, S., Bien, S., Jost, K., & Rösler, F. (2011). Memory-based decision making with heuristics involves increased activation of decision-relevant memory representations. *Journal of Cognitive Neuroscience*, 23, 3540–3554.
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, 81–97.
- Oberauer, K. (2002). Access to information in working memory: Exploring the focus of attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 411–421.
- Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, 41, 1–19.
- Weber, E. U., Goldstein, W. M., & Barlas, S. (1995). And let us not forget memory: The role of memory processes and techniques in the study of judgment and choice. In J. R. Busemeyer, R. Hastie, & D. L. Medin (Eds.), *Decision making from a cognitive perspective: The psychology of learning and motivation*, 32 (pp. 33–82). New York: Academic Press.

⁴ We thank one anonymous reviewer for this suggestion.