

COMMENTARIES

Understanding the A-not-B Error: Working memory vs. reinforced response, or active trace vs. latent trace

Adele Diamond

Center for Developmental Cognitive Neuroscience, Waltham, USA

In many ways, what Munakata has done here is a *tour de force*. She attempts to account for diverse findings and offers testable hypotheses about predicted behaviors. Her modelling work is theoretically driven and proposes explanations not readily apparent from simply observing the behaviors. I have some reservations about the work, however: (a) Munakata's theory provides less that is new than Munakata implies, (b) there are errors of fact in the manuscript, (c) the characterization of my theory of the A-not-B error is incorrect, and (d) although Munakata acknowledges some instances where her theory has difficulty accounting for the data, her theory has a bit more difficulty in this regard than Munakata acknowledges.

What is old and what is new in Munakata's formulation? The interpretation of the A-not-B error that Munakata offers shares much in common with other interpretations, some of which have been around for many years. However, there are also aspects of Munakata's formulation that are genuinely new. There are at least three ways in which she departs from the interpretations offered by myself or Goldman-Rakic.

Reading the initial pages of Munakata's manuscript, I had difficulty discerning the difference between her theory of the A-not-B error and my own. For example, on page 164 Munakata writes that "the A-not-B error arises based on a competition between 'latent' memory traces for A and 'active memory' traces for B". I have long argued that "A-not-B sets up a competition between the ability to use short-term recall to guide behavior and a conditioned behavioral tendency to repeat a rewarded response" (Diamond, 1985: 880). If Munakata's "active" memory trace was equivalent to "short-term recall" or "information held in mind" (whether one calls this working memory or sustained attention) (Diamond, 1985; 1991a; 1996) and her "latent memory trace" was

equivalent to "a conditioned tendency" or "the tendency to repeat a rewarded response" (Diamond, 1985; 1991b; 1996) then there would be no difference between Munakata's new formulation and my old one.

It turns out that while Munakata's active memory appears to correspond with what Goldman-Rakic and I have each called working memory, her conception of the latent memory trace is new; it is not the same as a conditioned tendency. On page 164, Munakata talks about latent memory traces in terms of "changes in firing thresholds or synapses that affect neurons' subsequent firing to a stimulus". In the course of processing stimuli, monkeys lay down latent memory traces for them, resulting in facilitated processing (i.e., reduced firing) when they are repeated (p. 164). This sounds akin to priming, and the neural region Munakata talks about in this regard is the inferotemporal cortex (Area IT). No one so far as I know has previously proposed this as the behavioral tendency that must compete with the memory of where the reward has been hidden in the A-not-B task. I have proposed that the competing behavioral tendency is a conditioned response, built up on the basis of reinforcement experience, subserved by subcortical systems, not IT (e.g., Diamond 1985; 1991a). Munakata shares, I think, my notion that the latent trace represents a response bias (Diamond, 1991a: 162) and that this represents a kind of procedural or implicit memory, but the difference in our formulations is interesting.

As I read on, Munakata's account of the A-not-B error began to sound like the formulations put forward by Goldman-Rakic (1987) and Kimberg and Farah (submitted). For example, Goldman-Rakic has argued that the A-not-B error reflects an immature ability to maintain active representations on the "stage" of one's mind (working memory), and that apparent failures of

Address for correspondence: Adele Diamond, Director, Center for Developmental Cognitive Neuroscience, Eunice Kennedy Shriver Center, 200 Trapelo Road, Waltham, MA 02154, USA.

© Blackwell Publishers Ltd. 1998, 108 Cowley Road, Oxford OX4 1JF, UK and 350 Main Street, Malden, MA 02148, USA.

inhibitory control result simply from memory failures; subjects emit their prepotent or conditioned behavioral tendency when their ability to hold the correct response in mind has been exceeded (i.e., when they forget). This appears to be echoed in Munakata's suggestion that the critical development "is the gradual improvement in this ability to maintain active representations" (p. 166). "Apparent improvements in infants' abilities to solve problems or to inhibit inappropriate responses may result from their increasing maintenance abilities" (p. 173).

A new wrinkle in Munakata's formulation is the notion that the representation of the hiding at *B* "fades more rapidly than in prior trials" to *A* (p. 171). I, and others I believe, have assumed that representations fade at a constant rate over trials, unless baseline conditions change. However, neither I, nor anyone else I know of until Munakata, has proposed that the representation of the hiding at *B* fades more rapidly than the representation of the hiding at *A*. I confess that it confuses me that Munakata adds that the representation of the hiding at *B* fades more rapidly "due to the influence of latent traces" (p. 171). Here, the latent traces appear to be affecting the active traces, whereas I had thought that Munakata was agreeing with Goldman-Rakic – the latent traces are simply what one sees when the active trace is too faint to guide behavior.

An exciting new idea is Munakata's explanation for "why non-perseverative responding appears earlier in development in gaze/expectation than in reach" (p. 171). Munakata argues that this is "based entirely on the lower frequency of reaching, which results in fewer opportunities for the reaching system to update based on a recently presented location" (p. 171). No one had previously proposed this, and I find it a very promising notion. In contrast, for example, I have talked about the reach being reinforced, but not the gaze, as the explanation for why infants can sometimes look to the correct well (*B*) even as they reach back to *A*.

Errors of fact

To support her statement that in "violation-of-expectation variants of the A-not-B task, 8–12 month old infants look longer when a toy hidden at *B* is revealed at *A* than when it is revealed at *B* [correct performance because it is taken as an indication of surprise that the toy is revealed at a place other than where it had been hidden] following delays at which they would nevertheless search perseveratively at *A*" (p. 163), Munakata cites Baillargeon and Graber (1988) and Baillargeon,

DeVos, and Graber (1989). However, the work reported in those two papers did not investigate that question and so never demonstrated what Munakata asserts. What was reported in those papers is that infants are correct in their looking behavior on the *A* trials, but then infants are also correct in their reaching behavior on the *A* trials. No trials at *B* have ever been reported in work from Baillargeon's lab. They found that infants look longer when a toy hidden at *A* is revealed at *B* than when it is revealed at *A*. This point is rather important because these two papers are often incorrectly cited as demonstrating that while the A-not-B error is seen in reaching, it is much rarer in looking, and that infants can withstand far longer delays in the A-not-B task in violation-of-expectation variants than in the standard reaching version. The crucial trials in the A-not-B task, however, are the reversal trials (the trials to *B*) and those trials were never administered by Baillargeon and her colleagues.

It is true, however, that when investigators have introduced delays as brief as 3–7 sec to infants of 9 or 10 months, they sometimes find errors even on the trials at *A* (Diamond, 1985; Sophian & Wellman, 1983), whereas Baillargeon *et al.* (1989) found good performance by infants of 8 months on the trials at *A* in their violation-of-expectation paradigm even with delays of 30 or 70 sec. It would thus appear that infants can tolerate longer delays on the *A* trials in the latter paradigm than in the former one. There may be two reasons for this other than the one Munakata entertains:

(1) To reach correctly in the standard A-not-B paradigm, infants must in a sense predict where the toy is going to be found before they see it; they must recall where the toy was hidden. To look longer in a violation-of-expectation paradigm, however, infants need only recognize that something is amiss. Years ago, Clifton (1974) studied heart rate conditioning by exposing newborn infants to a 300 Hz square wave that predicted the delivery of glucose. She replicated the failure to find normal heart rate conditioning in newborns reported by others; she failed to find a heart rate change to the tone in anticipation of the glucose. However, she did find that on those rare trials where the glucose was omitted, newborns' heart rates decelerated when the expected glucose did not appear when it should have. This surprise shown by Clifton's newborns seems analogous to me to the surprise shown in violation-of-expectation paradigms, and the failure to predict when the glucose would appear seems analogous to me to the failure to predict where the toy will be found in the A-not-B task. Thus, the difference may be in the cognitive demands of the Baillargeon and A-not-B tasks, rather than that one assesses gaze and the other reaching.

(2) The other reason may be the different criteria used to determine good performance in the two paradigms. In the A-not-B task, we judge performance on the trials at A against an expectation of 100% correct performance. In violation-of-expectation paradigms, performance is measured against chance (50% correct).

The assertion in Munakata's paper that practice trials are typically provided at the start of an A-not-B experiment to induce infants to reach to A (p. 166) is incorrect. This incorrect assumption led Munakata to include in her model four initial pre-trials at A before what she calls the first trial at A. Although some labs have presented the A-not-B task with practice trials at A (e.g., Bremner, 1978; Butterworth, 1977; Cummings & Bjork, 1983; Frye, 1980; Harris, 1974), most labs have not (e.g., Benson & Uzgiris, 1985; Bower & Patterson, 1972; Diamond, 1985; Evans, 1973; Fox, Kagan, & Weiskopf, 1979; Horobin & Acredolo, 1986; Schubert, Werner, & Lipsitt, 1978; Sophian & Yengo, 1985; Willatts, 1979).

Finally, in my opinion, Munakata incorrectly summarizes the results of studies of A-not-B involving multiple hiding wells. Her summary is that infants perform better on the B trials in the A-not-B paradigm if there are multiple wells than if there are only two wells. She notes that "Diamond, Cruttenden, & Neiderman (1994) argued that improvements in performance with additional hiding locations are an artifact of procedural differences between two-well and multiple-well variants of the A-not-B task" (p. 176). However, we did not just argue this; we demonstrated it empirically: When all wells in the multiple-well variant were covered simultaneously, as is normally done in the 2-well variant, we found that performance was poor with multiple wells. Munakata needs to account for those findings. Munakata cites Bjork and Cummings (1984) as demonstrating the opposite; they left the A well covered in both the two-well and multiple-well variants, uncovering only the B well to hide the toy and then recovering it; they found that on the first trial at B about 50% of the infants were correct whether two or five wells were used. This is curious because when Harris (1973) covered the A-well first and then the B-well (covering the B-well last as Bjork and Cummings did) they found that all infants (100%) succeeded. Not only are the Bjork and Cummings' findings discrepant from those of Harris, but they are also discrepant from our own findings: We found far better performance with multiple wells when only the correct choice was uncovered and recovered versus when all the wells were covered simultaneously; Bjork and Cummings found comparable performance regardless of the order of covering.

Setting the record straight: my theory of why infants make the A-not-B error

In her manuscript, Munakata has written that "Diamond (1985) proposed that infants' inability to inhibit a conditioned reaching response to A causes the A-not-B error" (p. 163) and she has characterized my interpretation of the A-not-B error as an "inhibition theory" (p. 163). This summary and characterization are wrong. Consistently, since I started writing about the A-not-B error through the present, I have proposed that success on the A-not-B task requires both memory and inhibition. For example, in the 1985 paper Munakata cites I wrote: "On the one hand, short-term recall memory appears to be one of the abilities required by the A-not-B task.... On the other hand, it is clear that memory cannot fully explain the A-not-B error.... The factor, then, in addition to memory, required for success on A-not-B, is the ability to resist the conditioned tendency to reach back to A" (p. 880). "Improved performance on A-not-B with age ... depends on both recall memory and the ability to resist or inhibit prepotent response tendencies" (p. 882). Moreover, the Diamond (1985) paper that Munakata cites provides some of the strongest evidence for the role of memory, or the ability to hold active representations on line, in the A-not-B literature. Evidence was presented that infants' performance on the A-not-B task is exquisitely sensitive to the delay used between hiding and retrieval: Infants who were making the A-not-B error at a given delay, ceased to err if the delay was reduced by only 2–3 sec, and made the A-not-B error again when the longer delay was reinstated. Thus, varying the delay, holding everything else constant, significantly affected whether the A-not-B error occurred or not. It was also reported that as infants grew older, progressively longer delays were required to produce the A-not-B error. Indeed, "recall" appears in the title of the paper but "inhibition" does not.

I have never wavered from this memory + inhibition interpretation. Only last year, I re-stated my summary of many of the A-not-B results in essentially a 2×2 table (memory demands: high or low along one axis and inhibitory demands: high or low along the other) first put forward in the 1985 paper "Some errors can be elicited simply by taxing working memory or sustained attention, e.g., by using a long delay at the first hiding location (e.g., Sophian & Wellman 1983). Similarly, some errors can be elicited simply by taxing inhibitory control, e.g., some infants err on the reversal trials to B even when the covers are transparent (e.g., Butterworth 1977; Willatts 1985). However, most errors by far occur when subjects must both hold information in mind and also exercise inhibitory control over their behavior;

i.e., on reversal trials to *B* when the covers are opaque and a delay is imposed” (Diamond, 1996: 1485).

Munakata suggests that if my account of the A-not-B error were correct, “infants should become less reactive and more planful given more time, and so should make fewer perseverative errors with longer delays, but instead they make more” (p. 163). Yet, my account predicts exactly what is observed. The memory of where the toy was last hidden is fragile and short-lived; it fades rapidly over time. However, the conditioned predisposition to repeat a rewarded response is robust and long-lasting. Hence, given more time, the conditioned predisposition is likely to win out over the active representation “It may seem contradictory to argue that infants have difficulty remembering where the toy was hidden a few seconds ago, and yet can remember where they last found the toy on previous trials (which happened perhaps minutes ago). This is not contradictory because two different kinds of memory are involved, which rely on different neural systems. The kind of memory that shows up as a response bias is the kind of memory that has traditionally been assessed using conditioning paradigms. Studies that have used conditioning to assess memory in infants... have typically found quite long memory in very young infants” (1991a: 162).

Problems for Munakata’s theory

In Munakata’s model “additional trials [at *A*] primarily serve to influence the strength of the networks prepotent response, such that additional *A* trials lead to a stronger A-not-B error” (p. 181). However, infants are no more likely to make the A-not-B error, and no more likely to repeat that error over successive trials at *B*, if they receive 2 vs. 3, 1 vs. 3, or 3 vs. 5 *A* trials (Butterworth, 1977; Diamond, 1983; Evans, 1973). Even administering 8 or 10 trials at *A*, rather than only 2, does not affect the likelihood that infants will make the A-not-B error, although it does affect over how many trials they continue to reach back to *A* if they do make the A-not-B error (i.e., it affects the length of the perseverative error strings; Landers, 1971). To be fair, this is also a problem for my formulation. I have acknowledged the problem (Diamond, 1991b). Munakata needs to face it more squarely than to say, “It is unclear whether infants make more A-not-B errors following more *A* trials” (p. 181). Evidently, one trial is sufficient to establish the bias to reach to *A*; providing a few more trials at *A* does not seem to make this bias any stronger; although a lot more *A* trials (8 or more) does seem to strengthen the bias.

I do not see how Munakata’s model can account for infants’ A-not-B errors when the toy is visible, as when transparent covers are used. If errors are due to an inadequate ability to sustain the active representation of where the toy has been hidden (i.e., if errors are due to forgetting where the toy went), then infants should not err if they do not need to sustain an active representation because the toy is visible. My theory has no problem accounting for this because sometimes the prepotent response to reach back to *A* wins even when the infant knows where the toy is located. My theory can also account for why fewer errors are made with transparent versus opaque covers – errors are most likely when both working memory and inhibitory control are taxed.

Monkeys with damage to the hippocampal formation (the H+ lesion) can hold information in mind for 1–10 sec, but not for 30 sec (e.g., Zola-Morgan, Squire, & Amaral, 1989). If the A-not-B error is due solely to an inadequate ability to hold information in mind, as I think Munakata is suggesting, why then do monkeys with H+ lesions not show the A-not-B error at delays of 30 sec? They err at that delay; but they are no more likely to err on a *B* trial than on an *A* trial (Diamond, Zola-Morgan, & Squire, 1989).

Similarly, if errors occur because infants forget (the active representation fades), then one might expect infants to err randomly. However, we have demonstrated that when multiple wells are used infants reach consistently in the direction of the *A* well, rather than randomly to either side of *B* (Diamond, Cruttenden, & Neiderman, 1994). Such behavior is easily accommodated by my formulation, and indeed was predicted based on that formulation (Diamond, 1985). Perhaps Munakata can account for this by the similarity of her latent traces to my prepotent response tendency.

While Munakata’s formulation is imperfect, and overlaps a good deal with older formulations, her achievement here is considerable nevertheless, spanning as it does developmental psychology, cognitive neuroscience, and computational modelling, presenting interesting new hypotheses, and testable predictions.

References

- Albert, M., Diamond, A., Fitch, H., Neville, H., Rapp, P., & Tallal, P. (in press). Development Chapter 54 in F.E. Bloom, S.C. Landis, J.L. Roberts, L.R. Squire & M.J. Zigmond editors. *Fundamental Neuroscience*. San Diego: Academic Press.
- Baillargeon, R., DeVos, J., & Graber, M. (1989). Location memory in 8-month-old infants in a non-search AB task: Further evidence. *Cognitive Development*, 4, 345–367.

- Baillargeon, R., & Graber, M. (1988). Evidence of location memory in 8-month-old infants in a non-search AB task. *Developmental Psychology*, **24**, 502–511.
- Bjork, E.L., & Cummings, E.M. (1984). Infant search errors: Stage of concept development or stage of memory development. *Memory and Cognition*, **12**, 1–19.
- Benson, J.B., & Uzgiris, I.C. (1985). Effect of self-initiated locomotion on infant search activity. *Developmental Psychology*, **21**, 923–931.
- Bower, T.G.R., & Paterson, J.G. (1972). Stages in the development of the object concept. *Cognition*, **1**, 47–55.
- Bremner, J.G. (1978). Spatial errors made by infants: Inadequate spatial cues or evidence of egocentrism? *British Journal of Psychology*, **69**, 77–84.
- Butterworth, G. (1975). Spatial location codes in determining search errors. *Child Development*, **46**, 866–870.
- Butterworth, G. (1977). Object disappearance and error in Piaget's stage IV task. *Journal of Experimental Child Psychology*, **23**, 391–401.
- Clifton, R.K. (1974). Heart rate conditioning in the newborn infant. *Journal of Experimental Child Psychology*, **18**, 9–21.
- Cummings, E.M., & Bjork, E.L. (1983). Perseveration and search on a five-choice visible displacement hiding task. *The Journal of Genetic Psychology*, **142**, 283–291.
- Diamond, A. (1983). Behavior changes between 6–12 months of age: What can they tell us about how the mind of the infant is changing? *Unpublished doctoral dissertation, Harvard University*.
- Diamond, A. (1985). The development of the ability to use recall to guide action, as indicated by infants' performance on AB. *Child Development*, **56**, 868–883.
- Diamond, A. (1991, a). Frontal lobe involvement in cognitive changes during the first year of life. In K. R. Gibson & A. C. Petersen (Eds.), *Brain maturation and cognitive development: Comparative and cross-cultural perspectives* (pp. 127–180). NY: Aldine de Gruyter.
- Diamond, A. (1991, b). Neuropsychological insights into the meaning of object concept development. In S. Carey & R. Gelman (Eds.), *The epigenesis of mind: Essays on biology and knowledge* (pp. 67–110). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Diamond, A. (1996). Evidence for the importance of dopamine for prefrontal cortex functions early in life. *Philosophical Transactions of the Royal Society (London) Series B*, **351**, 1483–1494.
- Diamond, A. (submitted). The role of motivation in performance of the A-not-B task.
- Diamond, A., Cruttenden, L., & Neiderman, D. (1994). AB with multiple wells: I. Why multiple wells are sometimes easier than two wells. II. Memory or memory + inhibition? *Developmental Psychology*, **30**, 192–205.
- Diamond, A., Zola-Morgan, S., & Squire, L.R. (1989). Successful performance by monkeys with lesions of the hippocampal formation on AB and object retrieval, two tasks that mark developmental changes in human infants. *Behavioral Neuroscience*, **103**, 526–537.
- Evans, W.F. (1973). The stage IV error in Piaget's theory of object concept development: An investigation of the rise of activity. *Unpublished doctoral dissertation, University of Houston*.
- Fox, N., Kagan, J., & Weiskopf, S. (1979). The growth of memory during infancy. *Genetic Psychology Monographs*, **99**, 91–130.
- Frye, D. (1980). Stages of development: The stage IV error. *Infant Behavior and Development*, **3**, 115–126.
- Goldman-Rakic, P.S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. *Handbook of Physiology*, **5**, 373–417.
- Harris, P.L. (1974). Perseverative search at a visibly empty place by young infants. *Journal of Experimental Child Psychology*, **18**, 535–542.
- Horobin, K., & Acredolo, L. (1986). The role of attentiveness, mobility history, and separation of hiding sites on stage IV search behavior. *Journal of Experimental Child Psychology*, **41**, 114–127.
- Kimberg, D., & Farah, M. (submitted). Are there inhibitory processes in the prefrontal cortex? Working memory and the mechanisms underlying higher cognitive processes.
- Landers, W.F. (1971). Effects of differential experience on infants' performance in a Piagetian stage IV object-concept task. *Developmental Psychology*, **5**, 48–54.
- Schuberth, R.E., Werner, J.S., & Lipsitt, L.P. (1978). The stage IV error in Piaget's theory of object concept development: A reconsideration of the spatial localization hypothesis. *Child Development*, **49**, 744–748.
- Sophian, C., & Wellman, H. M. (1983). Selective information use and perseveration in the search behavior of infants and young children. *Journal of Experimental Child Psychology*, **35**, 369–390.
- Sophian, C. & Yengo, L. (1985). Infants' search for visible objects: Implications for the interpretation of early search errors. *Journal of Experimental Child Psychology*, **40**, 260–278.
- Willatts, P. (1985). Adjustment of means-ends coordination and the representation of spatial relations in the production of search errors by infants. *British Journal of Developmental Psychology*, **3**, 259–272.
- Zola-Morgan, S., Squire, L.R., & Amaral, D.G. (1989). Lesions of the hippocampal formation but not lesions of the fornix or mammillary nuclei produce long lasting memory impairment in monkeys. *Journal of Neuroscience*, **9**, 897–912.