# Age-Related Inhibition and Learning Effects: Evidence from Transitive Performance

**Joanna Bryson (j.j.bryson@bath.ac.uk)**
Department of Computer Science, University of Bath
Bath, BA2 7AY, United Kingdom
The Konrad Lorenz Institute for Evolution and Cognition Research
Adolf Lorenz Gasse 2, A-3422 Altenberg, Austria

**Abstract**

In the absence of sufficient cognitive stimulation, intelligence — and with it, a capacity for learning new things — tends to degrade with age. In this paper I explore a novel hypothesis: that this may be an adaptive solution, since learning is facilitated by temporarily inhibiting action selection and thus introducing temporal delays. An older animal that is not being cognitively challenged may be in a sufficiently stable environment that reducing the delay before action at the cost of also reducing learning capacity may be a sensible tradeoff. I support parts of this model by matching a simulation of it to known reaction-time results, and by providing an account for previously unexplained task-learning results in adult and aged macaques

**Keywords:** inhibition ; learning; action selection; transitive inference; primate task learning; consciousness.

## Introduction

In the absence of sufficient cognitive stimulation, intelligence — and with it, a capacity for learning new things — tends to degrade with age (Schaie et al., 2004). In contemporary culture, we face the prospect of reduced intelligence with some trepidation, and for good cause. In humans, intelligence is correlated with both health and longevity (Deary et al., 2004). We can take comfort, though, in the *lack of sufficient stimulation* aspect of these findings. For many individuals, intelligence *increases* over the course of their life (See Figure 1). "[I]nvolvement in a complex and intellectually stimulating environment" (Schaie et al., 2004, p. 311), whether vocation-based or as a consequence of being partnered with a "high cognitive status" spouse (Gruber-Baldini et al., 1995) is one of the main correlates of maintaining one's IQ into old age.

Why should intelligence be a "use it or lose it" type of trait? There may be a simple metabolic explanation as in muscle tone. Neurons are highly plastic so may be open to metabolic optimisation, if the *capacity* for intelligence (as well as thought itself) has metabolic cost. In this paper I concentrate on a hypothetical second, complimentary cost: thinking and learning take time. Time is costly for two reasons:

1. Competition with others. For social organisms, individuals compete with other agents with very nearly equivalent needs and capabilities (e.g. Schaik & Noordwijk, 1988). Organisms also face competition for resources with other species, the ultimate example being predation.

2. Time is itself a finite resource. An organism has a finite lifespan and daily tasks associated with metabolic demands. The consequence of this is an evolved time budget for activities that are mutually exclusive (Dunbar, 1992; Bryson & Tanguy, 2009).

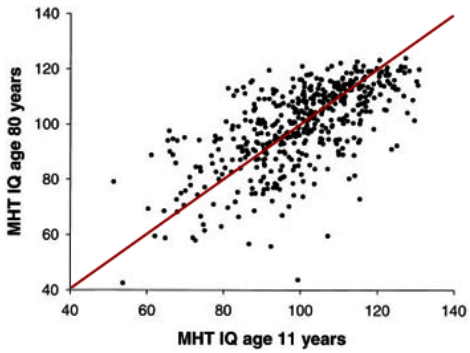


Figure 1: Age-corrected IQ correlations (Moray House Test) from the 1921 birth cohort from Lothian county, Scotland. After Deary et al. (2004, Fig. 3, p. 135)

Since cognition takes time, in the absence of advantage from that cognition it is adaptive to think as little as possible.

Put another way, the quality of an animal's existing model should affect its propensity to explore rather than exploit. An animal that has developed reliable skills and/or a good, predictive model of the environment should be more inclined to exploit its knowledge than one that has not. Indeed, learning happens at a higher rate the more unpredictable an animal finds a reward regime (Waelti et al., 2001; Holroyd & Coles, 2002; Belavkin & Ritter, 2003).

We already know that animals that are uncertain also take longer to decide and act (Palmer et al., 2005; Huk & Shadlen, 2005; Bogacz et al., 2006). In these experiments, uncertainty is often generated by weak sensory stimuli or limited prior experience of categories to be discriminated or actions to be expressed. Here I hypothesise another, more general factor determining the amount of time at least some species spend thinking before they act.

In this paper I focus primarily on a study of the effects of aging in rhesus macaques. Macaques were until very recently in evolutionary time the most successful genus of primates. Like hominids, macaques are highly social, highly cognitive, and found in diverse ecosystems ranging from tropical to temperate. The study, by Rapp et al. (1996), shows that aged rhesus macaques have two peculiarities in their task-learning performance. First, they do not exhibit a reaction time (RT)

effect traditionally attributed to computation the task requires, yet their performance is identical to younger animals that do. Second, they do not learn new behaviour when their reward schedule changes, unlike younger animals.

My interpretation of these results is that the RT associated with this task is completely independent of the task performance, but rather indicative of a general-purpose learning system which allocates time-consuming cognitive resources based on uncertainty in action selection. I argue that this general-purpose system may be attenuated by aging, and speculate that it may be correlated with conscious attention.

## Proposed Global Model for Allocating Cognitive Resources

In this section I briefly review the adaptive tradeoffs influencing aging, then I review an established model explaining reaction-time results in a cognitive vision task. From this I propose a general definition and description of cognition, and propose a model of how cognitive resources are allocated. The following sections pertain to data supporting this model.

### Evolution and Aging

Life is often defined in terms of a capacity to replicate. Darwin (1859) proposed an explanation for the diversity of life we see by showing how a process of generating an oversupply of diverse offspring which then compete for finite resources leads to systematic sorts of changes in number and characteristics of species. This explanation of biological evolution has been refined to show that competition between individuals is essentially a proxy for competition between what truly replicates (Dawkins, 1976). Genes form coalitions expressed as organisms, but the organisms compete not necessarily for their own survival, but rather for strategies most likely to propagate as many of their genes as possible into the future.

Since both evolution and other non-biological environmental forces generate constant change in an ecosystem, it is generally not in a replicator's best interests to inhabit one rigidly-defined organism. Thus many organisms express life histories where generations turn over at a high rate. Offspring carry newer genomes than their parents with incrementally more and newer input from selection. However, other organisms exploit niches where benefits of long life tradeoff more strongly against the benefits of evolutionary approaches to learning. For example, trees exploit size to increase access to light and water and to escape some forms of predation. Acquiring size though takes time.

Some animal species also exploit size and its associated longevity, probably initially for similar reasons as some plants do. However, animals also evolved a capacity for individual learning which provides further benefit, though also associated costs (Barrickman et al., 2008). For example, many social species transmit behaviour not only genetically but through social learning, but reliably acquiring behaviour this way takes time (Čače & Bryson, 2007).

There are two reasons then that it might be adaptive for older animals to reduce their propensity for learning. First, an animal that has been successful for a long time in a fixed environment is likely to have already acquired an adequate behaviour repertoire. Secondly, in a condition where an environment is changing in ways an experienced animal is unable to anticipate, the tradeoff concerning individual vs. genetic learning may shift. It may become more important to favour genetic evolution to the disadvantage of older animals, even if they are still able to reproduce.

### Cognition, Learning, Action and Time

I now return to the question of the relationship between time and cognitive behaviour. First, I propose a definition:

> **Cognition** is the process of search an agent goes through when its next action is not readily apparent.

Further, I propose this process consists of:

1. fixing attention on an aspect of the present behaviour context, and
2. searching the expected consequences of potential actions that are primed by this context.

At an explicit level, we consciously experience processes similar to this. The search for a next action persists until we find a solution that passes some threshold criteria for "correct" or "good enough", *or* until we "give up". But why and when do we give up? And how much of this process is really under conscious control? My suggestion is that the priming of potential actions is automatic once attention is fixed on the task, and that in absence of explicit success criteria, the cognitive resources we devote to this task are inversely correlated with our confidence about the current leading candidate action.

To rephrase my proposal, the less certain we are about our next action, the longer we inhibit our performance of that action. While our selection is inhibited and our attention is on the task, an automated process searches for other candidate actions, attending to one at a time. Each such discovered action is tested against an internal model derived from experience (expectations). If an action is found about which we are more certain of a positive outcome, our behaviour expression becomes proportionally less inhibited. When the inhibition times out, the current best candidate action is expressed and its outcome observed.

This model is not entirely novel; it derives from two established ones. The first is the well-established model described by Wolfe et al. (2000) of conjunctive visual search. Vision research has identified a number of properties, such as colour and shape, which are called "pop-out properties" (Treisman & Gelade, 1980). A subject conducting visual search for such a property takes the same amount of time to find an exemplar regardless of the number of distracting items. However, if one must search for an item that combines two pop-out properties, time becomes dependent on the number of distracters. For example, one can easily find the only green **T** in a field of **T**s,

or the only **L** in a field of **T**s, but finding the only green **L** in a field composed of brown and green **T**s and brown **L**s takes time proportional either to the number of green objects or the number of **L**s. Wolfe et al. (2000) propose that these reaction-time results indicate that the task is solved by randomly sampling from a pool of candidates determined by pop-out for one property, then testing whether the sample candidate also has the second property.

My account of action selection extends from this model by assuming the pool of candidate actions are primed from memory rather than gathered from the visual field. It conforms in that the candidates' evaluation is necessarily sequential, presumably due to a cognitive constraint. For example, th comparisons required could require hippocampal processing (Van Elzakker et al., 2003), which we know has task-dependent representations (Kobayashi et al., 1997). To the extent that my account proposes attention being directed by uncertainty in action selection, it is also similar to the account of Norman & Shallice (1986).

## Transitive Performance: A Model Task

Transitive performance is a standard cognitive task used in experimental psychology. Transitive inference (TI) formally refers to the process of reasoning whereby one infers that if, for some quality, $A > B$ and $B > C$, then $A > C$. In some domains, such as integers or heights, this property will hold for any *A*, *B* or *C*, though for others it does not. Though classically described as an example of concrete operational thought (Piaget, 1954), it has now been demonstrated in a variety of animals and young children not normally considered to be capable of advanced cognition (Grosenick et al., 2007). The behaviour of choosing *A* from *AC* without training after having previously been trained to select *A* from *AB* and *B* from *BC* is consequently referred to as "transitive *performance*" (TP).

### Representations Underlying TP

A variety of models have been proposed to to explain pre-/non-operational transitive performance. Each implies its own processes and representations (see Fig. 2.)

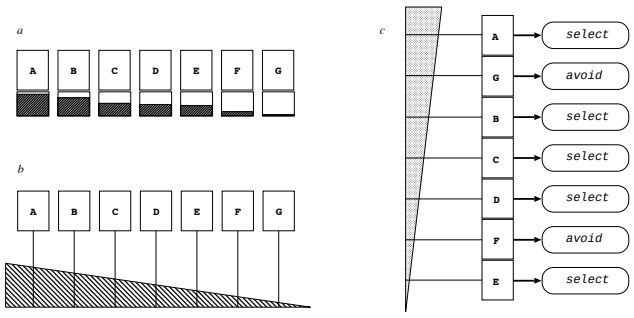


Figure 2: Three sorts of representations that may underlie expressed transitive inference. See text for details. *a*, weights associated directly with stimuli on the basis of the stimuli's association with reward. *b* stimuli ordered within a single index. *c* stimuli both ordered and associated with an action.

- **a** *simple association:* a reinforcement weight is paired with each individual stimulus; when more than one stimuli is present, the most reinforced one is selected. This representation is discriminated from others listed here in that items in the middle of the series may have no significant difference in weight (Frank et al., 2003). This representation is most closely associated with the value-transfer model, and is well-supported for explaining TP in birds (von Fersen et al., 1991; Zentall & Sherburne, 1994)

- **b** *innately sequential:* each value is associated with an ordinal index in a series or a value selected from a gradient; two items are compared by using their indexical values. This representation can be discriminated by generalisations (or mistakes) across series, e.g. the 3*rd* item of one series may be inserted as the 3*rd* item of another series. This representation has been convincingly modelled in a related literature on human series learning (Henson, 1998) and has been favoured by many modellers (e.g. Frank et al., 2003). A problem with this model is explaining how subjects determine that TP stimuli form a sequence.

- **c** *prioritised associations:* an ordered set of paired associations between stimuli and actions; each stimuli is associated with an action (e.g. select the stimuli). If there are actions associated with more than one present stimuli, whichever stimulus has the highest priority takes the agent's attention and determines its behaviour. This representation can be discriminated from the others by the fact that it fails systematically when *three* stimuli are presented, even though TP is observed for any pair. It is the only model fully consistent with human and other primate behavioural data on three-item performance (McGonigle & Chalmers, 1977; Chalmers & McGonigle, 1984; Bryson & Leong, 2007). It is also parsimonious in that it does not assume the recognition of a sequence. Rather, it assumes prioritising associations is a standard part of task-learning, an assumption quite similar to that made by the well-established ACT-R model of human learning (Wood et al., 2004).

### The Two-Tier Model

In previous work we have simulated the prioritised associative model of transitive performance (Bryson & Leong, 2007; Wood et al., 2004). We called our model *the two-tier model*, to emphasise the fact that we were proposing two concurrent disjoint learning tasks. One tier learns the of paired association between stimuli and one of two actions: *select* or *inhibit selection* (Harris & McGonigle, 1994). Another tier learns a prioritisation of these rules. Of course, because this learning system is modular it can easily be trapped in local minima. Though apparently a disadvantage, the difficulties our simulation has learning the original adjacent pairs models the difficulty of live subjects. Both our simulation and real animals need orderly training to reliably master the task, and even so some agents fail to meet criteria. Further, the

errors made by our system are consistent with those made by live subjects. The modules we propose are consistent with one of the leading theories of hippocampal functioning (Baxter & Murray, 2001; Heckers et al., 2004; Alvarado & Bachevalier, 2005).Modularity is adaptive because intelligent solutions are only useful if they are computationally tractable (Bryson, 2002).

### Explaining the Symbolic Distance Effect

I now return to testing the action-selection model described in Section . The two-tier model has been shown to account for all the known characteristics of TP with one exception, the Symbolic Distance Effect (SDE). The SDE is a characteristic reaction time (RT) effect; one that first drew interest to TP as a learning task (Bryant & Trabasso, 1971). When subjects execute a transitive comparison, they operate *faster* the further away two items are in the implied sequence. For example, a correct decision on *BD* would be slower than one on *BE*, even if *E* is not the last item in the sequence[1]

If TP were performed by simple inference, then items further apart would be expected to take *longer*, because more inferences have to be performed. That they are in fact faster helped motivate the *innately sequential* theories described earlier, as researchers tried to conceive an internal representation where further-removed stimuli were easier to discriminate (Bryant & Trabasso, 1971). However, the SDE is not a reliable individual effect, only an aggregate one (McGonigle & Chalmers, 1992). This is another factor in favour of prioritised association models, since they allow individual differences in the order of representation (Harris & McGonigle, 1994; Wood et al., 2004).

In the two-tier model, priorities of the associated stimulus / action pairs are normalised such that each pair had a value between 0 and 1, and all the priority values add up to 1. This paper's action-selection hypothesis assumes RT should be correlated with certainty in action choice. Here I propose "certainty" be modelled as the distance between the priorities of any two stimuli. I examined these values for a population of 144 agents with a wide range of individual learning parameters, the same range reported in Bryson & Leong (2007). These were trained on adjacent pairs drawn from seven stimuli, *A*–*G*. The agents for the present experiment were never exposed to transitive test pairs, so these could not affect their priority values. Although there is a large variety of individual orderings of priority on stimuli, in aggregate the average ratio of

$$\frac{|priority(B) - priority(F)|}{|priority(C) - priority(E)|} \quad (1)$$

is 2.034, with ninety-five percent confidence intervals in the range $[1.98, 2.09]$ ($SD = 0.33, SE = .03$). In other words, if the reaction time for the two-tier model is set proportional to the absolute difference between the priorities for the stimuli, then two-tier simulations do display the SDE in aggregate though not for every individual. On this qualitative level, the results match those reported by McGonigle & Chalmers (1992) for live subjects.

[1]End items are by far the easiest stimuli in TP, because unlike intervening items they are uniformly rewarded. Thus TP studies generally exclude end items from study.

## Age, Action and Learning

As further support for my hypothesis, I now propose a novel explanation of the results of Rapp et al. (1996). It has generally been assumed that RT results are due to *processing* complexity; for example that more sequential steps must be taken to perform a task. But the model I am proposing here assumes an any-time process (Dean & Boddy, 1988). An any-time algorithm is one which can always generate some solution, but that improves the solution the longer it has to compute. In my proposal, the length of time given is determined by the agent's certainty about its next action, and controlled by frontal-lobe inhibition. During the period of delay, the agent visually and intellectually attends to the problem space, mentally exploring possible next acts until it becomes disinhibited either because the inhibition fades naturally or because the level is reset when a superior solution is discovered. As I proposed in the introduction, the period of inhibition may also be determined by age, growing shorter as an adult ages.

In the case of TP, the agent is unlikely to find another solution because the task space is severely impoverished, as is typical in artificial experimental settings. However, we are still able to view the inhibitory aspect of this process, which results in the SDE.

### TP in Aged Monkeys (Rapp et al., 1996)

Support for this theory can be found in the work of Rapp et al. (1996), on TP in aged monkeys. First, Rapp et al. show that rhesus macaques of an advanced age (20-24 years) learn the initial, adjacent pairs for TP in approximately the same time as younger animals. Further, those that pass criteria in learning the pairs perform just as well as younger animals on the transitive pairs. However, these aged animals do not exhibit the SDE, but rather complete *all* transitive tests in the same time, which is much faster than their younger conspecifics. In general, RTs tend to increase with age, so this finding is particularly remarkable. It also indicates that the SDE does not reflect computational processing necessary for TP.

Second, due to an error in procedure, all monkeys received an unusual reward schedule for the only transitive pair they were initially tested on, *BD*. The error in procedure was as follows. Bryant & Trabasso (1971), who originally specified the protocol for pre-operational TP, specified that the test, transitive pairs should be "non-differentially rewarded". Since Bryant & Trabasso worked with children, the reward consisted of thanking the child for any choice, but not telling them whether they were right or wrong. This solution is not typically seen as available for non-human species. McGonigle & Chalmers (1977) reasoned that non-differential reward should simply be reward regardless of choice. Since well-trained subjects presumably expect reward for the choice they have made, this schedule is the least disruptive available.

Rapp et al. however interpreted *non-differential reward* to mean *random reward*. In the initial testing, they rewarded any choice on *BD* with 50% probability. The subjects were then trained on further adjacent pairs (*EF* and *FG*). Rapp et al. were surprised to subsequently find that the younger animals, while performing correctly on the newly-available transitive pairs such as *BE*, *CE*, and *DF*, were now at chance on *BD*. The aged monkeys however still performed correctly on *BD*.

The fact that the performance changed for the younger adult animals *after* the unusual reward schedule (not on test day) has subsequently been explained by Ellenbogen et al. (2007), who show the importance of both delay and sleep for learning TP implicitly. What is more relevant to the present theory is the correlation between the lack of SDE in the aged monkeys and the lack of effect of the random reward schedule. This sustains my argument that the purpose of the period of the SDE is in fact to explore alternative solutions which leads as a side effect to learning.

## Conclusion

In this article I have proposed two things:

1. A relatively novel explanation of reaction time effects in action selection. This extends from the Wolfe et al. (2000) account of reaction time for visual cognition.

2. A novel adaptationist explanation of why aged animals lose frontal-lobe inhibition.

The explanation of reaction time effects is that animals allocate time to be spent on search and learning based on how uncertain they are of the next action they should perform. This uncertainty triggers a level of inhibition which is also dependent on a second factor — the age of the animal. The adaptationist explanation for this latter dependency I have proposed is that in a stable environment aged individuals can be expected to already know appropriate action, and in a changing environments the optimal average reproductive age for a species may decrease.

Clearly at this point my hypotheses are speculative, as such they motivate a good deal of future research. For example, this model predicts that it may be possible to eliminate reaction time effects (such as the symbolic-distance effect for transitive performance) in human subjects, perhaps by distracting their attention with another task or by chemically altering their level of inhibition. My hypothesis predicts that the new RT will be directly correlated with the level of learning, but less so with the level of performance. In the present paper I have supported my proposals in two ways:

- By providing the first explanation of previous findings by Rapp et al. (1996), that aged monkeys express no SDE and also do not detect changes to their reward schedule, and

- to extend the Bryson & Leong (2007) two-tiered model of transitive performance to account for the symbolic-distance effect.

## Acknowledgements

Mark Baxter first drew my attention to Rapp et al. (1996) as an unexplained result in the area of both aging and transitive performance, and has generally served as an expert consultant, including reading drafts. Thanks also to Nick Priest for a useful discussion of evolution and aging.

## References

Alvarado, M. C., & Bachevalier, J. (2005). Comparison of the effects of damage to the perirhinal and parahippocampal cortex on transverse patterning and location memory in rhesus macaques. *Journal of Neuroscience*, *25*(6), 1599–1609.

Barrickman, N. L., Bastian, M. L., Isler, K., & Schaik, C. P. van. (2008, May). Life history costs and benefits of encephalization: A comparative test using data from long-term studies of primates in the wild. *Journal of Human Evolution*, *54*(5), 568–590.

Baxter, M. G., & Murray, E. A. (2001). Opposite relationship of hippocampal and rhinal cortex damage to delayed nonmatching-to-sample deficits in monkeys. *Hippocampus*, *11*(1), 61–71.

Belavkin, R. V., & Ritter, F. E. (2003, April). The use of entropy for analysis and control of cognitive models. In F. Detje, D. Dörner, & H. Schaub (Eds.), *Proceedings of the Fifth International Conference on Cognitive Modeling* (pp. 21–26). Bamberg, Germany: Universitäts–Verlag Bamberg. (ISBN 3-933463-15-7)

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*(4), 700.

Bryant, P. E., & Trabasso, T. (1971, August 13). Transitive inferences and memory in young children. *Nature*, *232*, 456–458.

Bryson, J. J. (2002, December). Language isn't quite *that* special. *Brain and Behavioral Sciences*, *25*(6), 679–680. (commentary on Carruthers,"The Cognitive Functions of Language", same volume)

Bryson, J. J., & Leong, J. C. S. (2007, January). Primate errors in transitive 'inference': A two-tier learning model. *Animal Cognition*, *10*(1), 1–15.

Bryson, J. J., & Tanguy, E. A. R. (2009, September). Simplifying the design of human-like behaviour: Emotions as durative dynamic state for action selection. *International Journal of Synthetic Emotions*, *1*(1). (accepted for publication)

Chalmers, M., & McGonigle, B. O. (1984). Are children any more logical than monkeys on the five term series problem? *Journal of Experimental Child Psychology*, *37*, 355–377.

Čače, I., & Bryson, J. J. (2007). Agent based modelling of communication costs: Why information can be free. In C. Lyon, C. L. Nehaniv, & A. Cangelosi (Eds.), *Emergence*

*and evolution of linguistic communication* (pp. 305–322). London: Springer.

Darwin, C. (1859). *On the origin of species by means of natural selection*. London: John Murray.

Dawkins, R. (1976). *The selfish gene*. Oxford University Press.

Dean, T., & Boddy, M. (1988, August). An analysis of time-dependent planning. In *Proceedings of the seventh national conference on artificial intelligence (aaai-88)* (pp. 49–54). Saint Paul, Minnesota, USA: AAAI Press/MIT Press.

Deary, I. J., Whiteman, M. C., Starr, J. M., Whalley, L. J., & Fox, H. C. (2004). The impact of childhood intelligence on later life: Following up the Scottish mental surveys of 1932 and 1947. *Journal of Personality and Social Psychology*, *86*(1), 130–47.

Dunbar, R. I. M. (1992). Time: A hidden constraint on the behavioural ecology of baboons. *Behavioral Ecology and Sociobiology*, *31*(1), 35–49.

Ellenbogen, J. M., Hu, P. T., Payne, J. D., Titone, D., & Walker, M. P. (2007). Human relational memory requires time and sleep. *Proceedings of the National Academy of Sciences*, *104*(18), 7723.

von Fersen, L., Wynne, C. D. L., Delius, J. D., & Staddon, J. E. R. (1991). Transitive inference formation in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, *17*(3), 334–341.

Frank, M. J., Rudy, J. W., & O'Reilly, R. C. (2003). Transitivity, flexibility, conjunctive representations, and the hippocampus: II. a computational analysis. *Hippocampus*, *13*(3), 341–354.

Grosenick, L., Clement, T., & Fernald, R. (2007). Fish can infer social rank by observation alone. *Nature*, *445*, 429–432.

Gruber-Baldini, A. L., Schaie, K. W., & Willis, S. L. (1995). Similarity in married couples: A longitudinal study of metal abilities and rigidity-flexibility. *Journal of personality and social psychology*, *69*(1), 191–203.

Harris, M. R., & McGonigle, B. O. (1994). A model of transitive choice. *The Quarterly Journal of Experimental Psychology*, *47B*(3), 319–348.

Heckers, S., Zalesak, M., Weiss, A. P., Ditman, T., & Titone, D. (2004). Hippocampal activation during transitive inference in humans. *Hippocampus*, *14*(2), 153–162.

Henson, R. N. A. (1998). Short-term memory for serial order: The Start-End Model. *Cognitive Psychology*, *36*, 73–137.

Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review*, *109*(4), 679–709.

Huk, A. C., & Shadlen, M. N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *Journal of Neuroscience*, *25*(45), 10420–10436.

Kobayashi, T., Nishijo, H., Fukuda, M., Bures, J., & Ono, T. (1997). Task-dependent representations in rat hippocampal place neurons. *Journal of Neurophysiology*, *78*(2), 597–613.

McGonigle, B. O., & Chalmers, M. (1977, 23 June). Are monkeys logical? *Nature*, *267*, 694–696.

McGonigle, B. O., & Chalmers, M. (1992). Monkeys are rational! *The Quarterly Journal of Experimental Psychology*, *45B*(3), 189–228.

Norman, D. A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behavior. In R. Davidson, G. Schwartz, & D. Shapiro (Eds.), *Consciousness and self regulation: Advances in research and theory* (Vol. 4, pp. 1–18). New York: Plenum.

Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, *5*(5), 376–404.

Piaget, J. (1954). *The construction of reality in the child*. New York: Basic Books.

Rapp, P. R., Kansky, M. T., & Eichenbaum, H. (1996, October). Learning and memory for hierarchical relationships in the monkey: Effects of aging. *Behavioral Neuroscience*, *110*(5), 887–897.

Schaik, C. P. van, & Noordwijk, M. A. van. (1988). Scramble and contest in feeding competition among female long-tailed macaques (*macaca fascicularis*). *Behaviour*, *105*(1-2), 77–98.

Schaie, K. W., Willis, S. L., & Caskie, G. I. L. (2004). The Seattle longitudinal study: Relationship between personality and cognition. *Aging Neuropsychology and Cognition(Neuropsychology Development and Cognition Section B)*, *11*(2-3), 304–324.

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*(1), 97–136.

Van Elzakker, M., O'Reilly, R. C., & Rudy, J. W. (2003). Transitivity, flexibility, conjunctive representations, and the hippocampus. I. an empirical analysis. *Hippocampus*, *13*(3), 334–340.

Waelti, P., Dickinson, A., & Schultz, W. (2001, July 5). Dopamine responses comply with basic assumptions of formal learning theory. *Nature*, *412*, 43–48.

Wolfe, J. M., Klempen, N., & Dahlen, K. (2000). Postattentive vision. *The Journal of Experimental Psychology: Human Perception and Performance*, *26*(2), 293–716.

Wood, M. A., Leong, J. C. S., & Bryson, J. J. (2004). ACT-R is *almost* a model of primate task learning: Experiments in modelling transitive inference. In *The* $26^{th}$ *annual meeting of the cognitive science society (CogSci 2004)* (pp. 1470–1475). Chicago: Lawrence Erlbaum Associates.

Zentall, T. R., & Sherburne, L. M. (1994). Transfer of value from S+ to S- in a simultaneous discrimination. *Journal of Experimental Psychology: Animal Behavior Processes*, *20*(2), 176–183.