Homo Oeconomicus 26(2): 1–16 • (2009)

www.accedoverlag.de

Why Private Events Are Associative: Automatic Chaining and Associationism

Robert Epstein
University of California, San Diego, USA
(eMail: repstein@post.harvard.edu)

Abstract That every response is also a stimulus has important implications for how we characterize the private experiences of both people and non-human animals. Acting as stimuli, responses, whether covert or overt, change the probability of subsequent responses. Hence, all behavior, covert and overt, is necessarily associative in some sense, and thinking may be characterized as 'covert autochaining.' According to this view, animals capable of responding to temporally remote stimuli and to characteristics of their own bodies necessarily engage in some form of associative thinking. This characterization of thinking necessarily presumes that private behavioral events adhere to at least some processes that occur in – and have been extensively studied in – overt behavior. To assume otherwise, as do Daniel Dennett, Robert Nozick, and others, is to be unnecessarily pessimistic both about the robustness of evolutionary processes and about our ability to explain complex human phenomena in rigorous empirical terms.

Human thinking has been characterized as associative for at least two millennia. According to this perspective, one thought does not simply follow another closely in time, but rather the first somehow helps to determine the content of the second. Successive thoughts are somehow connected to each other – sometimes tightly, sometimes more loosely – and are not just a series of unrelated ideas bubbling up from the depths, vying one against the other for attention. Although this general idea is widely held, philosophers and psychologists have long disagreed on the reasons why thoughts are ordered and on the mechanisms by which they are connected.

ISBN 978-3-89265-071-3 ISSN 0943-0180

Aristotle thought that remembering involved a sequence of related ideas – ideas that were similar or contrasting, or that had been paired frequently. John Locke is remembered for the phrase 'association of ideas,' even though he in fact wrote little about association. His *Essay Concerning Human Understanding* failed to mention association until its fourth edition in 1700 and even then was concerned only with how the association of ideas might account for certain disorders of thinking. It was Locke's predecessor, Thomas Hobbes, who actually came closer to a modern characterization of thinking in his 1651 *Leviathan*. Hobbes described thinking as an orderly process, with goal-directed thought ('regulated by some desire') being the most orderly. But even the most chaotic 'Trayne of Thoughts' – in which 'the thoughts are said to wander, and seem impertinent to each other, as in a Dream' – was orderly, said Hobbes, if you examine it closely enough:

> And yet in this wild ranging of the mind, a man may oft-times perceive the way of it, and the dependence of one thought upon another. For in a Discourse of our present civill warre, what could seem more impertinent, than to ask (as one did) what was the value of a Roman Penny. Yet the Cohaerence to me was manifest enough. For the Thought of the warre, introduced the Thought of the delivering up the King to his Enemies; The Thought of that, brought in the Thought of the delivering up of Christ; and that again the Thought of the 30 pence, which was the price of that treason: and thence easily followed that malicious question; and all this in a moment of time; for Thought is quick (Hobbes 1651/1950: 17).

In our era we celebrate James Joyce for his unique ability to capture or at least simulate this 'quick' process in fluid prose. And psychiatrists and clinical psychologists have captured dramatic examples of associative prose in the 'word salad' of schizophrenic patients (e.g., Arieti 1974). In a word salad, trivial properties of recent words – often phonetic properties rather than meanings – may shift the speaker willy-nilly from one topic to another, as in 'I want to take a walk around the block, tock, tick tock, I wish I had a clock.'

So why, then, is human thinking associative? Why don't we instead think in a jumble of unrelated ideas that gurgle up from the teeming neuronal depths, each idea breaking into awareness by virtue of winning a kind of neural Darwinian competition? Such 'pandemonium' models of cognition – taking that very descriptive term from Selfridge's (1959) theory of learning – are currently enjoying some popularity. And even Skinner contended that while operant responses 'come under the discriminative control' of specific stimuli, they are not directly 'elicited' by previous stimuli but rather are 'emitted' with only a certain probability of occurrence when in their presence. Why then (to maintain the associative thread with classical thinkers) isn't thought just a kind of Humean process, a landscape upon which one event follows another in time, but without probable causal connection between them?

For the last one hundred and fifty years or so researchers and scholars have visited a number of explanatory frames, not staying long on any one choice: thoughts as silent speech or as fractional motor responses (peripheralism); thoughts as events confined to the central nervous system (no time for long motor loops!); thoughts as covert derivatives of verbal or nonverbal responses and serving as signs and symbols (meaningful!); thoughts as symptoms of deeper cognitive processes and principles of organization (Gestalt) or controlled transformations leading to equilibration (Piaget); thoughts directed by the organism's motivational states (Hull); thoughts as implicit trial-and-error responses (Dewey: the Law of Effect internalized to avoid costly bad tries); thought as essentially heuristic (Newell, Shaw, and Simon). But in the competition of ideas about associative thinking, such perspectives have proved to be of only limited value, and no real winner has emerged.

A return to some fundamental principals of behavior may prove helpful. One way of answering the question, Why is human thinking associative?, is to suggest – in a reassertion of an old but still healthy experimental tradition – that thinking is associative because *behavior* is associative. In reframing the issue in this way, we may not only be able to shed light on human thinking but also to venture a reasoned opinion about whether the thinking of animals is similarly associative.

One simple observation establishes an immutable link between the study of overt behavior and the study of many aspects of cognition. The observation has been made before, but because its consequences could not be thoroughly elaborated – the science of behavior was not as robust as it is now – was passed over and (almost) forgotten. That observation is, simply, that *every response is also a stimulus.*

Definitions

Precise definitions of 'stimulus' and 'response' are difficult to formulate and easily challenged, and one might even argue, as I and others have, that more powerful and naturalistic formulations of the activity of organisms in the natural environment are possible (Epstein 1982, 1996). However, the terms are adequate for the purposes of the present discussion, so I will offer brief definitions. A stimulus is an event in the environment which, mainly by virtue of its proximity to the organism, has the potential to alter the organism's behavior. *Note that the organism's own body is part of its environment.* A response is an event in the behavior of an organism. Responses that we name – for example, 'tooth brushing,' 'key pecking,' and 'leg flexing' – are actually classes of nonidentical events which have functional integrity – that is to say, which have roughly the same effect (cf. Catania 1979; Skinner 1938). Responses such as the lever press are relatively brief and simple; others, such as the recitation of the Pledge of Allegiance, are lengthy and complex.

Overt and Covert Stimuli

In the laboratory, stimuli are often presented by the experimenter or by apparatus without regard for what the organism is doing. A tone is sounded, or the silhouette of a predator bird is presented, or a red light is illuminated, no matter what the behavior of the subject. In classical conditioning, two response-independent stimuli are paired. Alternatively, stimuli are sometimes presented contingent upon and immediately following a response. A chimpanzee presses a sequence of buttons and is rewarded with food. A leap over the barrier in a shuttle box terminates shock. A lever press sounds a tone.

In the natural environment, we encounter stimuli in each of these ways. Both response-independent and response-dependent events are common. Stoplights change and phones ring no matter what we do; they are response-independent events. On the other hand, striking a piano key *produces* a sound; operating a light switch *changes* the room illumination; typing *places words* on paper or screen.

Overt stimuli – that is, stimuli observable by others – can be generated another way. Our own bodies often, if not always, produce fairly obvious stimuli. When you wave hello, you are not only responding, you are also producing a fairly obvious visual stimulus for others, as well as for yourself. An autistic child may wiggle her fingers in front of her eyes for hours, an example of behavior usually called 'self-stimulatory.' When one speaks, one presumably also hears.

The bodies of non-human animals also generate obvious stimuli, detectable both by the animals themselves and by others. A bird can presumably detect the flap of its wing visually and otherwise, and the orangutan presumably sees and feels the limb it moves.

Finally, both we and our non-human neighbors are continuously exposed to covert stimuli – events inside our bodies that are not accessible to other observers. Such events are, by definition, 'private.' Irritated nerve fibers may be experienced as aches and pains. While awake, we experience an almost continuous sequence of verbal and perceptual stimuli, and at night, dreams provide vivid stimuli of this sort. Hallucinations are perceptual stimuli so vivid that they overwhelm the signals provided by the sense organs. Note that what I have called 'perceptual stimuli' might just as well be called 'perceptual responses'; they are one and the same thing in this case. A perceptual or verbal response emitted covertly is necessarily a stimulus.

Non-human animals must also experience private events; there can be no possible basis for assuming otherwise. An animal that reacts to an externally-applied aversive stimulus must also react to irritated nerves that others cannot observe. Animals that can react to temporally remote stimuli, as evidenced, for example, by success in delayed matching-to-sample procedures (Cohen and Roberts 1996; Epstein and Skinner 1981; McSweeney, Weatherly, and Swindell 1996; cf. Barth, Fein, and Waterhouse 1995; Goodie and Fantino 1995), by long-delayed taste aversion learning (Burritt and Provenza 1991; Riley and Mastropaolo 1989; cf. Bitler and Riley 1992; Commons, Mazur, Nevin, and Rachlin 1987), or by deferred imitation (Epstein 1984), must have perceptual experiences of some sort in the absence of the temporally remote stimuli. The nature of these experiences, both for humans and non-humans, is not at issue here. The protracted debate between Pylyshyn (1973, 1981) and Kosslyn (1980, 1994) over the nature of 'mental images' showed how difficult it can be – perhaps even how fruitless it is – to attempt to determine the nature of private events in terms other than physiological. The point here is simply that covert responding cannot possibly be unique to humans, no matter what the precise nature of such responding.

Associationism and Autochaining

To move closer to answering why thinking is associative, we can supplement the observation that *every response is also a stimulus* with a simple behavioral process. The process is called *automatic chaining,* sometimes shortened to *autochaining* (Epstein 1985a, 1990, 1996). In a conventional chaining procedure, the experimenter establishes a specific sequence of responses by shaping a response, then bringing that response under the control of some stimulus, and then withholding that stimulus until a second response occurs. Thus, having taught a rat to press a lever for food when a tone is sounded, the experimenter now waits until the rat steps on a platform before turning on the tone, in effect using the tone to reinforce stepping on the platform. A reliable sequence – a response chain – results: the rat steps on the platform (and thus produces the tone) and then presses the lever (producing food). A third response can be added in similar fashion, and ultimately, long sequences can be engineered (Catania 1979). Note that each response in the chain produces an experimentally arranged stimulus.

The explicit backward chaining procedure of the laboratory – familiar to virtually everyone who has taken an elementary course in learning theory – is nevertheless rare or perhaps nonexistent in the natural environment . Automatic chaining, on the other hand, is commonplace and fundamental. *Our own behavior frequently, and perhaps continuously, changes our environment in some way, and exposure to a changed environment changes the probability* *of subsequent behavior.* Thus, self-generated (or 'automatically' generated) chains of behavior are the rule in the natural environment. For example, you open your refrigerator door and discover a chocolate cake, which you proceed to eat. Opening the door has produced a new stimulus, which has changed the probability of subsequent behavior. A monkey rattles the door of its cage to produce the rattling sound; the shaking loosens the latch, the door opens, and the monkey escapes. Note the opportunity for learning in such cases: the first time the monkey opens the door, it does so 'accidentally'; the second time will likely be 'deliberate.' Even a turn of the head can have a dramatic effect on subsequent behavior, because turning the head changes the visual field.

Automatic chaining is readily apparent in the creative, problem-solving performances I've studied with pigeons and people over the last few decades. For example, human subjects who have apparently given up solving Maier's (1931) classic pendulum problem often create a small pendulum inadvertently. They have tied an object to the end of a long string in order to try to solve the problem by extending their reach, but reaching doesn't work. At some point, they put the object down; invariably, the object and string swing back and forth somewhat, after which subjects solve the problem almost immediately (Epstein 1996). (When his subjects had trouble, Maier sometimes brushed against one string and set it in motion. Most subjects generate hints of this sort on their own.)

Epstein (1985a, 1987) described cases of problem solving in pigeons involving the interconnection of three and four separate repertoires of behavior, respectively. Although the performances were rapid and relatively efficient, the orientation of the pigeon's head was an excellent predictor of what it would do next. Generativity theory (Epstein 1985b, 1990, 1996, 1999) asserts that ongoing behavior in the natural environment is generated continuously as multiple repertoires of behavior are subjected to simple, concurrent transformation functions. Instantiated in a computer model, the theory has proved reasonably successful in predicting novel human performances moment-to-moment in time. Automatic chaining is one of four transformation functions utilized in the model.

The important implications of autochaining for the process of thinking are easily overlooked if we are familiar only with the laboratory technique of explicit chaining, where the presentation of reinforcers and other stimuli are deliberately conditioned upon overt responding. Presumably, the far more ubiquitous phenomenon of automatic chaining operates in covert behavior much the same way it does with overt behavior. I don't see how this could not be so, since verbal and perceptual responses are also stimuli, both covert and overt. In other words – returning to our original observation – *thoughts are necessarily associatively connected because every covert response is also a* *covert stimulus*. The conditions for autochaining are there at the covert level, and therefore it is not just a metaphor to refer to a 'chain of thought.'

The so-called 'laws' of association proposed by Hume, Hartley, and others, are not directly relevant and in fact may be of little value to this analysis. Rather, the 'stream of consciousness' is the rule because perceptual and verbal responses are also stimuli. When we envision the Eiffel Tower – that is, when our nervous system lights up in somewhat the same fashion that it would light up were the Eiffel Tower in front of us – a stimulus has been generated to which we will tend to react to some extent. If association is actually covert autochaining, thinking is subject to all of the complexities that autochaining entails, not the least of which are the vagaries of stimulus control itself (Williams 1984).

On a quiet afternoon, with few external stimuli to distract us, covert autochaining is fluid, continuous, and virtually unbounded. Hynagogic phenomena and nocturnal dreams are bizarre, presumably because the autochaining process is unrestrained by external stimuli (Mavromatis 1987). When, on the other hand, stimulus associations are deliberately taught, as they are in investigations of 'stimulus equivalence' or 'mediated transfer,' covert autochaining may produce a variety of new associations that derive from the ones that were originally trained (Horne and Lowe 1996; Peters 1935; Sidman, Rauzin, Lazar, Cunningham, Tailby, and Carrigan 1982; Sidman and Tailby 1982). What is more, this account of associative thinking as autochaining should apply to all organisms, human and nonhuman alike, that engage in covert behavior. Because every covert response is also a private stimulus, organisms that are capable of self-observation and that are also capable of reacting to a temporally remote stimulus must engage in some rudimentary form of thinking, and that thinking process, like all behavior, must necessarily be associative.

Continuity of Overt and Covert Behavior

It would be surprising if what scientists have learned about behavior in the experimental laboratory suddenly became irrelevant as behavior recedes to a less observable covert level. It seems unlikely that even associative thoughts so central that they are devoid of a peripheral accompaniments suddenly abide by a totally different set of governing laws. Only if this sudden transformation were the case could it be argued that autochaining is irrelevant to the process of thinking. And even if the brain somehow occasionally intruded upon the ongoing chain with spurious 'deep thoughts' that differed in kind from known behavior types, it would be unlikely to completely discard behavioral techniques useful at the overt level just because they were engaged in covertly – especially if these covert behaviors could be exploited to enhance the organism's adaptiveness, perhaps by providing new ways for private self-stimulation, or perhaps for vicarious trial-and-error learning.

Explicit behavior chaining originated in a laboratory tradition heavily interested in the concept of reinforcement. The deliberate experimental chains of overt behavior were linked together with reinforcers, giving behavior its direction. Autochaining, too, carries the implication that the organism is tracking reinforcer value during associative thought. (This is especially obvious when one is engaged in 'pleasant' fantasizing.) And although in this present account of autochaining and associative thought we do not need to decide what the role of reinforcement might be for all forms of thinking, there is no good reason to discard all that is known about operant behavior simply because the stimuli and responses are privately experienced.

Daniel Dennett, for example, has repeatedly argued that something like the law of effect must control thought processes, even as he insists that operant conditioning (and presumably, autochaining) is too simple a process to implement what he believes thinking brains can achieve (Dennett 1974, 1978, 1991, 1995). Possibly in an effort to sequester and constrain the unruly Skinnerian meme (a term originated by the evolutionary biologist Robert Dawkins [1976]) – memes being Dennett's favorite vehicle for replicating useful thoughts – Dennett caricatures operant process by describing a Skinnerian 'creature' as one that is limited to blindly trying different responses, occasionally hitting on and having these responses being reinforced by the consequences of a lucky try. Advanced brains, he argues, must be at least Popperian creatures, capable of trying out responses internally, evaluating their likely consequences before executing them overtly:

> Which animals are Popperian creatures, and which are merely Skinnerian? Pigeons were Skinner's favorite experimental animals, and he and his followers developed the technology of operant conditioning to a very sophisticated level, getting pigeons to exhibit quite bizarre and sophisticated learned behaviors. Notoriously, the Skinnerians never succeeded in proving that pigeons were *not* Popperian creatures, and research on a host of different species, from octopuses to fish to mammals, strongly suggests that if there are any purely Skinnerian creatures, capable only of blind trial-and-error learning, they are to be found among the simple invertebrates. (Dennett 1995: 376)

Note (with the sentence starting 'notoriously') that it is up to operant psychologists to prove that pigeons are *not* Popperian. Why does the burden of proof shift? Is it required that Dennett's proposed model of a Popperian creature be accepted by default until disproved, while a model of considerable robustness and voluminous experimental verification, and which may very well account for the behaviors exhibited by a Popperian creature, is labeled as 'too simple' to do the trick? It is not particularly difficult to see how a 'Skinnerian creature' engaging in covert behavior can experience vicarious reinforcement or punishment as various responses are 'tried' covertly. After all, is this not one direct implication of covert autochaining as associative thought?

Is Dennett evincing a bias that disallows using principles developed from overt behavior to explain behavior at the covert level? If so, he is not alone in this regard. Robert Nozick has stated the bias more directly. After characterizing operant behavior as limited to provide only a local optimization of behavior, lacking the more advantageous adaptation of more 'thoughtful' global optimization techniques, Nozick points out that just because complex behavior can be generated by synthesizing simper behavioral principles does not mean that's the way nature does it. It might just evolve a new function. In fact, the implication is that nature almost certainly will:

> [T]he following principle seems appropriate. If (a) we can see the limitations of certain processes, such as operant conditioning; (b) we can see what other processes would be selected for if they arose, such as cognitive processes; and (c) apart from these considerations, we naturally think anyway that such other processes are taking place in us, then we should be very suspicious of theories that deny the existence of these processes. Such theories carry a heavy burden of proof (Nozick 1981: 706).

And just what might these other processes be?

> Many of the things cognitive psychologists speak of, information processing, *trying out hypotheses in imagination* [italics added], and so on, would seem to suit this purpose of transcending some limits of operant conditioning (p. 705).

There is something strange – if not just dead wrong – about an argument that supposes evolution will necessarily rectify a supposed limitation of function by creating an incompletely specified mechanism in which intuition urges us to believe. Is this not precisely the kind of wishful thinking that science has been developed to restrain?

The issue here is not necessarily whether the process of autochaining is *sufficient* to explain the complexities of, say, a Popperian creature's anticipatory behavior, but rather whether behavioral processes studied at the overt level retain their identity when covert. Even if evolution deemed it necessary to include new and better functions at the covert level, is it parsimonious to believe that adaptive behavioral processes would be summarily dismissed when they come to be covertly available, relying exclusively on a totally new process to take over at the boundary of the skin?

Perhaps in Dennett's case, the bias is reserved for Skinnerians only (the title of one of Dennett's (1978) early papers was 'Skinner Skinned') and is not a general disqualification of overt behavior's continuity into covert manifestations. In fact, one of his 'just so' stories offered to help explain how humans evolved to the level of conscious thought features the internalization of speech, receding to covert autostimulatory behavior. He even goes so far as to suggest that as the overt speech chain receded to the private level, it helped to precipitate – via the Baldwin effect (an evolution-accelerating process that allows 'good tricks' learned at the phenotypic level to reach the genome) – the modification of brain pathways to better secure this autostimulation:

> In particular, we can speculate that the greater virtues of *sotto voce* talking to oneself would be recognized, leading later to entirely silent talking to oneself. The silent process would maintain the loop of self-stimulation [because *responses are also stimuli*] but jettison the peripheral vocalization and audition portions of the process, which weren't contributing much (Dennett 1991: 197).

Is behavior that is reduced to the covert, divesting itself of the longer and slower motor component so that it becomes genuinely central, still properly described by the laws of performance characteristic of overt behavior? Or does it suddenly become discontinuous with those laws, requiring a new formulation? Overt operant behavior can be described mathematically with one or another version of the 'matching function,' derived from an even more basic response-reinforcer input-output function (Herrnstein 1970). Are autochains of covert fantasies subsumed under the same rule, and what serves as reinforcers in that case?

Experimental verification of covert events must necessarily await the development of more sophisticated investigative techniques. Or, as with the limiting cases of high-energy physics and relativistic kinematics, which are assumed to be continuous with observable phenomena at more tractable ranges of heat and velocity, we may ultimately have to rely on indirect evidence and theoretical coherence, with its associated mathematical descriptions, to explain what must be happening at the extremes.

Donahoe and Palmer (1994), for example, simulate reinforcement (including internal reinforcement) processes with neural networks that closely mimic neurophysiological processes and brain structures rather than inventing novel internal cognitive processes:

> We must avoid the temptation to think of covert behavior as a *kind* of behavior, with properties essentially different from overt behavior. Rather, all behavior lies on a continuum of observability (Donahoe and Palmer 1994: 275).

Their rigorous selectionist approach to complex behavior (what Dennett would call 'greedy reductionist') is in sharp contrast to attempts to 'explain' the supposed limitations of simpler processes in accounting for complex behavior by assuming that evolution will just automatically create processes of a better kind just because we can imagine, at least roughly, what these might be!

> The fact that conditioning phenomena at the neural level can only be studied with instrumental intrusion poses overwhelming practical problems to our attempt to provide a complete experimental analysis of human behavior, but it does not pose problems in principle. We believe that we are justified in considering covert events – thoughts, images, unexpressed feelings – in our interpretation of complex behavior provided that we do not introduce *ad hoc* principles, principles that are not founded in the experimental analysis of overt, measurable, quantifiable behavior (Donahoe and Palmer 1994: 276–277).

And even if it is true that well-known processes, studied at the level of overt responding – e.g., operant conditioning and autochaining – have limitations that induced evolutionary add-ons (perhaps an innate capacity for language, or at least an innate tendency for children to learn it quickly), that does not mean that operant conditioning does not occur in humans (it does), anymore than the existence of operant behavior in animals means that Pavlovian conditioning fails to occur in animals (and humans) sensitive to the law of effect. Just because the principles of associative autochaining may not explain *all* of thought does not mean that they do not explain *some*, or even a very large part, of the stream of consciousness. Evolution may supplement or modify, but rarely does it completely replace.

A further quote from Nozick, who wishes to introduce a new moral theory, perhaps reveals the real motivation for his and Dennett's bias:

> The principle also would apply to consciousness and self-awareness, provided we identify the functions these perform and the limitations they overcome. And the principle applies to our current concern, free will. If we can describe an evolutionary function for free will so that *however it works* [italics added], we would expect it to be selected for if it arose, because it overcomes limitations of its absence, and if apart from these considerations we naturally think anyway that we do make (some) free choices, then we should be very suspicious of theories denying free will, and should view these theories as carrying a heavy burden of proof (Nozick 1981: 706)

Again the burden of proof has been shifted, this time in regard to consciousness and free will! But Dennett (1991) has already dispatched the issue of free will to his satisfaction and turned his attention to the grandest question in psychology – the problem of consciousness. Dennett wagers on what one might call the 'illusionist' choice, offering a detailed deconstruction of the supposed phenomena of consciousness into its component pieces, which individually are made to dissolve before one's very eyes. His characterization of the Skinnerian and Popperian creatures are just minor asides on his journey to a grand theory of consciousness.

As part of the larger vision, Dennett proposes an account of thinking which clearly contrasts with the one offered here. Consciousness, Dennett says, results from imposing a kind of serial virtual machine, implemented in software, upon the brain's parallel structure:

> There is no single, definitive 'stream of consciousness,' because there is no central Headquarters, no Cartesian Theater where 'it all comes together' for the perusal of a Central Meaner. Instead of such a single stream (however wide), there are multiple channels in which specialist circuits try, in parallel pandemoniums, to do their various things, creating Multiple Drafts as they go (Dennett 1991: 253–254).

Dennett is fighting old ghosts, and achieves his theory of consciousness by modeling the brain from the inside-out, arriving at the seriality of thought only after imposing a mechanism – a virtual machine composed largely of 'memes' – at the top to ride herd on the confusion of parallel events occurring at the lower levels. The present account takes a more direct route, which one might call outside-in: from the autochaining of overt behavior to covert associative thought.

Others have taken on consciousness, as well, offering a full range of grand theories, most of which are distinguishable from Dennett's by assuming consciousness is explainable at the physiological level. Most have little to tell us about the laws of thought (Frick 1994; Penrose, 1989, 1994; Chalmers 1996), although some, like Edelman (1989), have tried to build a comprehensive neuronal model combining perceptual and conceptual skills at all levels, from the simplest to the most complex. Searle (e.g., 1997) backs a causal theory of consciousness, arguing that the obvious correlation of brain states with reported levels of awareness compels us to believe that consciousness is an emergent feature of the brain, as real as any other natural property, and not some grand illusion produced by a zombie brain boot strapping itself to act 'as if' conscious. At the bottom of the heap we find discouraging theories such as those proposed by McGinn (1991) and Fodor (1992), which suggest that consciousness is so intractable as to make it cognitively unfathomable to us at our present level of evolution. (Would Nozick expect us to evolve out of this limited cognitive state just so the problem of consciousness could be solved by future philosophers?)

To be sure, the present suggestion that associative thought is a consequence of autochaining at the covert level says little about consciousness per se. Whether, as we engage in private, covert responding, we are genuinely conscious, or just acting as if we were, cannot now be decided. Research in psychology will proceed, as it always has, without the solution to this higher mystery, and until someone can explain *why consciousness matters in the first place*, the science of behavior, and of thought, can continue profitably within a stance of careful neutrality. In effect, I am suggesting that we restrict our thinking to the process of thinking itself, while offering an empirically-grounded account of why thinking displays its characteristic connectedness.

There is perhaps one tangential connection between the present account and the grand question of consciousness. In identifying thinking with covert autochaining, I am supposing that non-human animals must experience at least some human-like thinking, which suggests a certain level of awareness on their part. To the extent that an organism can react to temporally remote events in its own behavior, one might even argue that the organism is 'self-aware' (cf. Epstein and Koerner 1986, Parker, Mitchell, and Boccia 1994; Radner and Radner 1989). For Searle, with his theory that consciousness is an emergent feature of the brain, it is a foregone conclusion that some animals, at least, must be conscious. Dennett, perhaps, would probably not even accept that animals act 'as if' conscious.

The idea that every response is also a stimulus has been invoked before in psychology's past, but so has the assertion that a brain must 'try out hypotheses in imagination' – a requirement of Dennett's Popperian creature. The latter, in fact, has been a common characterization of thinking since at least William James's and John Dewey's time. The implications of combining the responses-as-stimuli idea with the laboratory-studied phenomena of autochaining gives the original concept more substance, I believe, with perhaps even more to follow. Indeed, that is the main advantage of theory that is grounded in an empirical analysis of observable phenomena; data collection constrains speculation.

References

Arieti, S. (1974). *Interpretation of Schizophrenia*. Basic Books.

Barth, C., Fein, D., and Waterhouse, L. (1995) Delayed matching-to-sample performance in autistic children. *Developmental Neuropsychology* 11: 53–69.

Bitler, D. A., and Riley, A. L. (1992). Memory for tastes in an operant delayed discrimination. *Bulletin of the Psychonomic Society* 30: 385–388.

Burritt, E. A., and Provenza, F. D. (1991). Ability of lambs to learn with a delay between food ingestion and consequences given meals containing novel and familiar foods. *Applied Animal Behaviour Science* 32: 179–189.

Catania, A. C. (1979). *Learning.* Prentice-Hall.

Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.

Crick, F. (1994). *The Astonishing Hypothesis: The Scientific Search for the Soul*. Simon and Schuster.

Cohen, J. S., and Roberts, R. (1996). The role of trial tracking on rats' successive delayed matching-to-sample modality performance. *Behavioural Processes* 36: 277–287.

Commons, M. L., Mazur, J. E., Nevin, J. A., and Rachlin, H. (Eds). (1987). *The effect of delay and of intervening events on reinforcement value*. Erlbaum.

Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press.

Dennett, D. C. (1974). Why the law of effect will not go away. *Journal of the Theory of Social Behavior* 5:169–187.

Dennett, D. C. (1978). Skinner skinned. In D. Dennett (Ed.), *Brainstorms*. Bradford Books.

Dennett, D. C. (1984). *Elbow room: The Varieties of Free will Worth Wanting*. MIT Press.

Dennett, D. C. (1991). *Consciousness Explained*. Little, Brown.

Dennett, D. C. (1995). *Darwin's Dangerous Idea.* Simon and Schuster.

Donahoe, J. W., and Palmer, D. C. (1994). *Learning and Complex Behavior*. Allyn and Brown.

Edelman, G. (1989). *The Remembered Present: A Biological Theory of Consciousness.* Basic Books.

Epstein, R. (1982). A note on the mythological character of categorization research in psychology. *The Journal of Mind and Behavior 3*: 161–169.

Epstein, R. (1984). Spontaneous and deferred imitation in the pigeon. *Behavioural Processes*9: 347–354.

Epstein, R. (1985a). The spontaneous interconnection of three repertoires. *Psychological Record* 35: 131–141.

Epstein, R. (1985b). Animal cognition as the praxist views it. *Neuroscience and Biobehavioral Reviews*9: 623–630.

Epstein, R. (1987). The spontaneous interconnection of four repertoires of behavior in a pigeon. *Journal of Comparative Psychology* 101: 197–201.

Epstein, R. (1990). Generativity theory and creativity. In M. A. Runco and R. S. Albert (Eds.), *Theories of Creativity* Sage.

Epstein, R. (1996). *Cognition, Creativity, and Behavior: Selected Essays*. Praeger.

Epstein, R., and Koerner, J. (1986). The self-concept and other daemons. In J. Suls and A. G. Greenwald (Eds.), *Psychological Perspectives on the Self* . Erlbaum.

Epstein, R., and Skinner, B. F. (1981). The spontaneous use of memoranda by pigeons. *Behaviour Analysis Letters* 1: 241–246.

Fodor, J. (1992, July 3). The big idea: Can there be a science of mind? *Times Literary Supplement*, p. 5.

Goodie, A. S., and Fantino, E. (1995). An experientially derived base-rate error in humans. *Psychological Science* 6: 101–106.

Hermstein (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior* 13: 243–266.

Hobbes, T. (1950). *Leviathan.* London: J. M. Dent and Sons. (Originally published by Andrew Crooke in 1651)

Horne, P. J., and Lowe, F. L. On the origins of naming and other symbolic behavior. *Journal of the Experimental Analysis of Behavior* 65: 185–241.

Kosslyn, S. (1980). *Image and Mind.* Harvard University Press.

Kosslyn, S. (1994). *Image and Brain: The Resolution of the Imagery Debate.* MIT Press.

Horne, P. J., and Lowe, C. F. (1996). On the origins of naming and other symbolic behavior. *Journal of the Experimental Analysis of Behavior* 65: 185–241.

Locke, J. (1700). *An Essay Concerning Humane Understanding* (fourth edition) A. and J. Churchill and S. Manship.

Maier, N. R. F. (1931). Reasoning in humans. II. The solution of a problem and its appearance in consciousness. *Journal of Comparative Psychology* 12: 181–194.

Mavromatis, A. (1987). *Hypnagogia: The Unique State of Consciousness Between Wakefulness and Sleep*. Routledge.

McGinn, C. (1993). *The problem of consciousness: Essays Towards a Resolution.* Blackwell.

McSweeney, F. K., Weatherly, J. N., and Swindell, S. (1996). Within-session changes in responding during delayed matching-to-sample and discrimination procedures. *Animal Learning and Behavior* 24: 290–299.

Nozick, R. (1981). *Philosophical Explanations*. Belknap Press.

Parker, S. T., Mitchell, R. W., and Boccia, M. L. (Eds.). (1994)*Self-Awareness in Animals and Humans: Developmental Perspectives.* Cambridge University Press.

Penrose, R. (1989). *The Emperor's New Mind.* Oxford University Press.

Penrose, R. (1994). *Shadows of the Mind: A search for the Missing Science of Consciousness*. Oxford University Press.

Peters, H. N. (1935). Mediate Association. *Journal of Experimental Psychology* 18: 20–48.

Pylyshyn, Z. (1973). What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin* 80: 1–24.

Pylyshyn, Z. (1981). The Imagery Debate: Analogue Media versus Tacit Knowledge. *Psychological Review* 88: 16–45.

Radner, D., and Radner, M. (1989). *Animal Consciousness.* Prometheus.

Riley, A. L., and Mastropaolo, J. P. (1989). Long-delay taste aversion learning: Effects of repeated trials and two-bottle testing conditions. *Bulletin of the Psychonomic Society* 27: 145–148.

Searle, J. R. (1997). *The mystery of Consciousness*. New York Review, Inc.

Selfridge, O. G. (1959). Pandemonium: A paradigm for learning. In D. V. Blake and A. M. Uttley (Eds.), *Proceedings of the Symposium on Mechanisation of Thought Processes*. London: H.M. Stationary Office.

Sidman, M., Rauzin, R., Lazar, R., Cunningham, S., Tailby, W., and Carrigan, P. (1982). A search for symmetry in the conditional discriminations of rhesus monkeys, baboons, and children. *Journal of the Experimental Analysis of Behavior* 37: 23–44.

Sidman, M., and Tailby, W. (1982). Conditional discrimination vs. matching to sample: An expansion of the testing paradigm. *Journal of the Experimental Analysis of Behavior* 37: 5–22.

Skinner, B. F. (1938). *The Behavior of Organisms: An Experimental Analysis.* Appleton-Century-Crofts.

Williams, B. A. (1984). Stimulus control and associative learning. *Journal of the Experimental Analysis of Behavior* 42: 469–483.