South African Journal of Philosophy

ISSN: 0258-0136 (Print) 2073-4867 (Online) Journal homepage: http://www.tandfonline.com/loi/rsph20

Learning, cognition and ideology

Don Ross

To cite this article: Don Ross (2003) Learning, cognition and ideology, South African Journal of Philosophy, 22:2, 139-156, DOI: 10.4314/sajpem.v22i2.31366

To link to this article: http://dx.doi.org/10.4314/sajpem.v22i2.31366

Learning, cognition and ideology

Don Ross

School of Economics
University of Cape Town
Private bag,
Rondebosch 7700
Cape Town

E-mail: <dross@commerce.uct.ac.za>

Abstract

Invited to give the 2000 Rick Turner Memorial Lecture, I pondered the following question: What explains the fact that the sincere thought of a brilliant and heroic person such as Turner can appear preposterous to me, if bad faith or scholarly ignorance on one side or the other are ruled out, as they should be in this case? I address this question by considering what 'ideologies' are from the perspective of cognitive learning theory. I describe the dynamics by which pressures for social coordination cause brains to implement alternative natural softwares for performing inferences in complex domains of association and inference. I conclude by noting that this need not imply normative relativism, since the relative justifications for conclusions produced by different softwares can still be debated. My aim is thus not to contest Turner's ideology or political views, but to partially explain how learning produces differences that transcend factual disagreements and even ethical ones.

What follows is the text of the Rick Turner Memorial Lecture, as given at the University of Natal-Durban on May 11, 2000. I have edited the text only with respect to the elimination of outright errors of reference, of stylistic barbarities, and of one major question left unacceptably unclear in my lecture that was noted by a critical referee for this journal (see note 3). However, I will preface the lecture here with an explanation as to my choice of topic. It was a formidable assignment for me – a person very far, in political sympathies, from the 'left' as it usually characterizes itself – to memorialize a noble man who martyred himself in the organization of labour. (Of course, more importantly, he martyred himself fighting against stupid, vicious and tyrannical racial prejudice, a cause that is certainly not the monopoly of the left.) I was, then, an unusual choice as a Turner lecturer. I accepted, however, because I thought the broad-mindedness of the invitation, if not necessarily the specific vehicle this broad-mindedness lit upon, was admirable. Deep respect for persons, after all, must be compatible with the most profound disagreements over political convictions; and this is a truth that always stands in need of reinforcement. Still, a lecture cannot just be a symbol; it must be *about* something. The crux of my challenge, as I saw it, lay in the fact that it's much harder to be genuinely, as opposed to merely politely, respectful of someone with whom one has few significant beliefs in common. But a memorial lecture demands deep respect, and not just as a matter of social etiquette. I decided that a way to

try to do this was to leave the *content* of my ideological differences with Turner alone, and take the very existence of such differences as a phenomenon to be explained. My topic is therefore not political economy itself, but a meta-political and meta-economic topic that has not been much explored: the cognitive-psychological foundations of ideological difference in general.

As a non- South African, I had been unacquainted with the work of Richard Turner prior to my accepting the invitation to deliver this lecture in his memory. After agreeing, I therefore looked into his book, *The Eye of the Needle*, and at one of Andrew Nash's well-known analytic memoirs of him (Nash 1999). I am going to begin here with some reportage on the nature of my own attempts to engage with Turner's concerns and thought, since my topic – the relationship between path-dependent individual learning processes and ideology – began in meta-reflection on my efforts at framing his view of the world in terms definable within mine. As Nash emphasizes, Turner was of course a dialectical (as opposed to a traditional Soviet-style) Marxist, inspired by the Frankfurt School and the Marxist existentialism of Sartre, rather than by Lenin and his political heirs. He therefore faced the challenge posed for all Western Marxists by their Leninist critics, namely, that of reconciling a programme for political action that must, if it is to be effective, find a basis for optimism within an intellectual matrix that depends, in its very formulation, on a hermeneutics of deep suspicion. All the evidence I have looked at suggests that Turner resolved this tension, at least for himself, a good deal more satisfactorily than Althusser, Jamieson and other famous names that have wrestled with it, whose thought turns caustic and obscure just where Turner's becomes confident and practical. The result seems to me to be more or less completely mad. I use this shocking term technically and descriptively, for reasons shortly to be explained; I of course mean no gratuitous insult here, and the report of the judgement states in any case a fact about me, not a fact about Turner. The interest of the judgement lies in the fact that it survives acquaintance with Turner's nobility of character and purpose; and this situation invites difficult reflections on the nature of political-academic debate. Nash and others clearly had immense respect for Turner; and, more significantly and impressively, so did the real political actors among whom his influence moved. (This applies both to his comrades among union activists and to his lethal enemies in the apartheid state, who did not regard merely academic opposition, no matter how strident, as threatening enough to motivate murder.) Now, if the most apt word I can find to characterize the thought of such an obviously gifted and inspiring person is 'mad', then this must prompt serious questions about how theorists and political actors understand one another. A Marxist, of whichever persuasion, might say that I, as a capitalist, have been duped by history; certain critics on my side, if they followed the lead of Robert Conquest or Francois Furet, would say instead that *Turner* was duped by history. I do not find these responses at all satisfactory. They purport to explain by endlessly postponing explanation, since they explicitly deny that both sides can be serious. The problem I want to talk about today begins precisely in the moment of considering the possibility that both sides *are* indeed serious. And then that problem is self-evidently very important; it is really the problem of whether, and how, politics and ethical debate about politics can be something other than war.

Philosophers of science have devoted much attention, for four decades now, to the so-called problem of 'incommensurability'. My remarks today will fall loosely within the ambit of that debate, but I will not draw much upon its main themes. This is be-

cause the philosophical arguments go on at a very high level of abstraction that ignores a distinction between two (or more) quite different levels at which belief systems may be incommensurable. As Ian Hacking (1983) has pointed out, if we try to understand the texts of medieval alchemists, writing about causal – or, *maybe* they're causal; one can't quite tell – processes of transmutation and property transformation in which `barnacle-geese' arise both ex nihilo and from other substances *at the same time*, and so on, then here we *really* have incommensurability. It is just impossible, beyond a certain point, for a twentieth century reader to imaginatively reconstruct and simulate the world view behind these texts. My difficulties with Turner and Nash are not at all like that. Of course I understand their words – whereas I really do not understand the word `barnacle-goose' - but my understanding goes far beyond that point. I understand where and why Turner departs from the assumptions of both capitalists and Soviet Marxists. I understand why this opens certain difficult problems for him, and I understand the logic of his attempts to solve them. In fact, there was a time, about twenty years ago and before I had settled into my present well-worn habits of thought, when in a very amateurish way I could and did *participate* in the project of the dialectical Marxists. There was no element of commitment in this participation, since I was very much in the role of a student then, trying on styles of thought to find out if I could be comfortable in them; but the results show a real enough alertness to the logic of the territory. I could not do that now, because Turner's fundamental categories of analysis seem to me entirely arbitrary and unrelated to reality. In that respect – and *only* in that respect – they are like the concept of a `barnacle-goose' to me. This is the sense in which I apply the adjective `mad', very seriously and advisedly. It seems to me *exactly* the sense in which the term is apt, and it is a much more informative term in the context of political argument than the less clearly focused `incommensurable'.

A full treatment of my topic would require a good deal more analysis in this purely philosophical vein, but today I want to cut closer to the chase, and so I will move on here. I will henceforth drop the modest autobiographical stance, and pretend that I'm speaking as a representative capitalist, and a representative philosopher of neo-classical economics to boot. This is for ease of exposition; nothing central to my topic is going to hinge on the pretense, so if you are a capitalist or a neo-classicist but don't like what I am saying here, you can simply translate my remarks back into the autobiographical mode. So, then:

A point made with great care and subtlety by Nash, and a very important one, is that the capitalist and the Soviet-style Marxist do *not* need to view one another as mad. They disagree with one another about the logic of value-ascription, with the neo-classicist being some sort of subjectivist and the Soviet Marxist being an objectivist. This difference is tremendously important to their rival theories of *government*, but it can be ring-fenced in disputes about political economy if the capitalist grants, as she is apt to, that most people's preference orderings are such that the maximization of their utility quite closely coincides with whatever improves their level and degree of security in material comfort. The capitalist and the Soviet Marxist can then be interpreted as arguing over the conditions under which material productivity growth will be highest. The Soviet Marxist believes that, beyond a certain point, the capitalist's incentive structures get in the way of this by generating inefficient concentrations of wealth; the capitalist retorts that the Soviet Marxist severely underestimates the dynamic significance of technological change in preventing liquidity pools from clotting. Capitalists thus

think that Marx and Lenin are very seriously *mistaken*, but we do not think they are *mad* (at least, as political economists). If all of this gets moralized, as it of course does, that is because each side is apt to suspect the other of self-interested motivations, and this suspicion is magnified by the extent to which any particular sets of opponents admire one another's intelligence. But this factor contributing to moralization is a psychological accident, not something intrinsic to the theoretical debate. Stalin, at least in private, wasn't given to (sincere) moralistic demonization of so-called 'class enemies' even when he tried to justify their mass-murder. His justification rested on the grounds that, objectively, they couldn't help it, not on the grounds that, morally, they *could*; he didn't have to appeal to a massive coincidence to the effect that all kulaks happened to be *voluntarily* morally rotten.[1] Where disagreements, however ferocious, involve perceived facts rather than deep questions of valuation, tactical accommodations and shifts are possible that otherwise would not be; hence the familiar point that Comintern policy zigzags with respect to fascism in the 1930s should not be read as exclusively and *simplistically* cynical. Nash partly explains the ultimate political failure of Turner's project in South Africa in terms of this substantial set of premises shared by the capitalist and the Soviet Marxist. An Alec Erwin, for example, can become a champion of orthodox fiscal policy merely by coming to the conclusion that South Africa's level of capitalist accumulation is not yet high enough to make socialism possible, and this hardly requires any sort of intellectual revolution or gestalt shift. I have no idea whether Erwin's actual views involve any such reasoning. But Nash cites Erwin's current policies as an example of the "loss of the Marxist moment" in South Africa; and it is implied that such policies can be rationalized by Soviet Marxists in a way that is not available to Western Marxists. On Nash's analysis, South African Western Marxists of Turner's sort lost their historical moment when Soviet Marxists politically captured the organized left's agenda; and this in turn made possible the current practical rapprochement between them and the so-called 'neoliberals'. I find this analysis very plausible.

The relevance of all this to my present topic is that capitalists and Soviet Marxists can use largely similar, indeed isomorphic, categories of fundamental analysis in political economy. (Again, this must be qualified by remembering that disagreements at the level of political *morality* can nevertheless be so sharp as to make war almost inevitable; but Nash's point is that this turned out not to be the case in South Africa.) This is no longer the case when the capitalist turns his attention to the views of Western Marxists. Here, concepts that have their origins in the purely philosophical aspects of Marx's thought – concepts he in turn owed mainly to Hegel – dominate the paradigm. That thought of this sort should depart radically from the organizing principles of the liberal tradition is hardly surprising, since, as has been pointed out so often as to be a cliché, the dialectic depends on assumptions of logical holism which are explicitly denied in the first moment of the broad metaphysical conception of which both British

1 The reader of the now-released transcripts of Central Committee sessions at which early phases of Stalin's terror against the party were set in motion is likely to be struck by Stalin's repeated distinctions between judgements of moral character and questions of political "fact". See, for example, Getty and Naumov (1999), p. 321, where Stalin explicitly remarks about this. His comments occur in the context of the mounting attack on Bukharin. When Bukharin tries to appeal to his well-known strength of revolutionary loyalty, Stalin chides him for misunderstanding the entire conversation. "I'm not saying anything personal about you," says Stalin; and the surrounding context provides no grounds for reading this ironically.

empiricism and Cartesian rationalism are species. 'History', 'consciousness', 'freedom', 'thought' itself can all, in a tradition partly descended from Hegel, denote aspects of reality which to thinkers such as Hume and Smith – and, for that matter, Descartes – are simply impossible objects of reference. Familiarly, this is where the divergence between Western and Soviet Marxists is also to be found; for the primitive variety of scientific realism that was the official metaphysic of the Bolsheviks could hardly be further from the dialectical holism of an Althusser or a Gramsci, even though both are derived from emphasizing different sides of Marx's own schizophrenic metaphysics.

It is not my purpose here to review familiar intellectual history. Before I move on, however, I should like to add something specific about Turner's place in this picture. The Hegelian moment of most radical departure from the traditions of Locke and Descartes is carried to its very *furthest* extreme in the philosophy of Sartre, to whose thought Turner devoted his doctoral dissertation. The empiricist and the rationalist understand 'freedom' in a purely relative sense, and as a concept to be unpacked by reference to *causal* relations. Sartre's 'freedom' is an altogether different sort of thing, a freedom to create and recreate logical and normative essences. Locke and Descartes had both been profoundly muddled about essences, but Sartre's work represents the far point of a trajectory that tries to *solve* their problems, whereas the Humean tradition has groped along a path of *dissolving* these problems by denying that the concept of an essence is logically coherent in the first place. A very large part of the 'madness' I find in Turner is really Sartre's madness. Something that is mainly implicit in most Western Marxists is explicit in Turner, namely, the view that the development of class consciousness is a process of *voluntary essentializing* of political conditions so that various sorts of political-economic limitations come to be recognized as *in fact* forms of alienation and oppression. Consistent Humeans, by contrast, cannot think in terms of these concepts because to them *nothing* is or could be politically fundamental outside of a specified psychological and/or ethological context. A Humean need not deny that there are selves, or that they are made by contingent processes over the courses of their biographies; but she must deny that they could be fashioned *ex nihilo* and *by mere thought*. It is when a Sartrean advises me to voluntarily define myself – indeed, tells me that I have an ethical responsibility to do this – that my comprehension fails in a way I was not philosophically sophisticated enough twenty years ago to recognize. I am inclined to describe the difference as follows. Then, I didn't know how to think deeply without essences. After twenty years of learning to philosophize without them, I can no longer imagine them – that is, *really* imagine them, in full sympathetic detail – in the same sense that I can't *really* imagine a process of demonic possession or a personal God. Note, again, that this is different from the sense in which I can't imagine a barnacle-goose; the process of imagination doesn't arise in that case because I don't know what the concept is even *supposed to* denote. I have a quite rich notion of conceptual *roles* for essences and gods, even if I can't imagine anything playing those roles, whereas I *don't* have a conceptual-role slot for barnacle-geese.

Describing the process in this way makes it sound like a variety of emancipation – I was once forced to muddle my thought with essences, now I am free of them. However, this begs the question entirely in favour of my own practice. In order to avoid that, I wish to try to get by as far as possible without further phenomenological or moral metaphors. A major problem with the phenomenological description above is

that it suggests that conceptual roles, and the objects that fill them, exist independently of any particular working cognitive economy. However, this is false; no models of cognition that have a remote chance of surviving through our developing scientific regimentation of psychology can posit intentional entities with characteristics independent of processing dynamics.[2] To put this point another way (and a way which needs to be put here, to avoid confusion): whereas Hegelian *logical* holism seems to me to entail severe, and fortunately gratuitous, epistemological complications, some degree of holism in the ascription of semantic content to particular states of biological systems seems to be forced by the facts.[3] The point I am making here is the following one: There is no fact of the matter to force a preference, one way or the other, for talking about processes of 'ideological settling' as 'losing' or 'gaining' something. Did I understand the metaphysics of the Western Marxists better or worse back when I could reason in their idiom, but without quite understanding the philosophical implications of doing so, than I do now? Either answer, it seems to me, would be arbitrary outside of some specified practical context for the question. To make further progress on the epistemic divides amongst adherents of different ideologies, we must turn to the more normatively neutral discourse of cognitive science.

There are of course many aspects of the ways in which organisms learn about which we are still ignorant. We know that some capacities, such as the ability to process natural language, are learned more or less automatically in normal circumstances, and are characterized by relatively sudden onset once a requisite level of developmental maturation is reached. (This does not presume – we should not presume – that this development is all generated 'from the inside'; the developmental system is coupled dynamically with its environment.) This general model likely applies to all *universal* competencies, including the capacity to theorize. The question as to whether any *particular* theories arise universally in this way turns partly on meta-questions about what we should most usefully mean by 'theory'. Thus, for example, some writers (e.g., Churchland 1979, Gopnick and Meltzoff 1997) insist that our folk-psychological apparatus for interpreting people, including ourselves, as intentional agents is a theory, and one that is almost certainly false in at least its fine details. Others prefer to reserve the label of 'theory' for systems of cognitive organization built through cycles of explicit conjectural generation and testing. Purely for purposes of exposition here – since I do, in fact, find it plausible and helpful to treat folk psychology as a theory – I will confine myself to using 'theory' in the second, narrower sense. This is because, although some theories may plausibly be biologically installed in typical minds, there would seem to be no possible point in referring to such natural structures as 'ideologies', and those are my subject at the moment. So in wondering about the learning processes involved in the acquisition of ideologies, we are focusing on the installation of 'optional' software in brains. (By 'optional' I again don't mean 'independent of the environment'; I mean merely 'not strongly culturally relative'.)

Does this use of 'optional' imply that a person could grow to cognitively functional adulthood without acquiring *any* ideology? This seems to me to be another ill-formed question. If we use 'ideology' loosely, to denote any set of broad schemata for regulating and predicting social relations, then the answer is certainly `no'. The most impor-

2 Note that trying to do this is, in fact, precisely a variety of logical-semantic essentialism.

3 A sophisticated treatment of this issue in the present literature is to be gleaned from the three-cornered debate amongst Ruth Millikan, Timothy Kenyon and Daniel Dennett in Ross *et al* (2000).

tant objects in the child's environment, from the perspective of natural selection, are other people, and evolution *never* leaves the possibility of developing schemata for organizing universally critical domains entirely to the discretion of individual learning. But this is as true of baboons and vervets as of humans; and I do not think it is helpful to think of monkeys as guided by ideologies. Science is likely to be of very limited help in sorting out this semantic issue, but we can fortunately bypass it on pragmatic grounds here. Given our present context, we can stipulate a quite narrow understanding of `ideology' as any generalized system of obligations and responsibilities that distinguishes socially legitimate from illegitimate access to material goods and personal services.[4] It is of course important here that I don't use a conception that begs any questions against Marxists. The conception just stipulated is consistent with Lenin's (1902) insistence that to call a system an `ideology' does not commit one to the claim that it distorts reality. Western Marxists have gone both ways on this question, so it seems safest here to use the neutral conception as our default. If we now make clear that a person need not view their ideology *as* an ideology in order to be said to think and act with an ideology, then we can say that all people who are not raised as farrow-children think and act with and through ideological structures. Subject to these careful constraints on the concept of an ideology, I think that both the Marxist and the liberal can endorse the claim.

This claim is not, however, as strong as it might sound, because it implies nothing about how to *individuate* ideologies. For example, it does not entail that a person can't use different ideologies in different contexts, nor that a person's social behaviour might not embody an idiosyncratic ideology cobbled together out of pieces of more comprehensive and shared ideologies. I am going to assume here that both of these things are possible – indeed, actual – since any attempt to legislate the possibilities off the table must surely be motivated from *within* some particular ideology. At this point we see that our working conception is really not so narrow after all. This is fortunate for present purposes, because it means that in asking about the learning of ideologies we can in fact just be wondering about the acquisition of developmentally optional social codes in general. And whereas cognitive science has nothing, at present, to tell us about the specific acquisition of elaborate, internally consistent ideological *theories*, we already know quite a lot about the learning of general 'cognitive maps'. I will now leave the domain of politics for a few pages in order to describe this consensus, before returning to consider its implications for the development of ideological gulfs.

The crucial insight that has taken us, in just two decades, from nearly complete ignorance about cognitive learning to the point where scientists have developed and experimented with quite sophisticated and plausible models of it, has been the recognition that 'hard storage' analogies are more of a hindrance than a help to understanding. Such analogies, from Aristotle's mental pictures to Locke's tabulae, were often given highly literal interpretations during the first decades following the invention of the Von Neumann machine. However, the experience in early artificial intelligence research of actually trying to build models based on the Von Neumann architecture re-

4 This obviously leaves room for ideologies organized centrally around conceptions of gender. But note that some quite powerful conceptions of gender might be biologically installed, in which case they will not be elements of ideologies in the sense used here. This is an empirical question for developmental psychologists.

vealed its fundamental inaptness as a biological approach.[5] Von Neumann systems trying to learn outside of highly constrained environments veer between crippling combinatorial explosions in their data loads and embarrassing crashes in performance when designers turn to representational shortcuts to keep processing loads manageable. Two decades of failure in attempting to model human cognitive performance this way provided extended logical reinforcement of the fact that biological realism at the level of mechanisms cannot be sacrificed altogether, as AI researchers initially supposed. The animal learning apparatus cannot resemble a Von Neumann device for evolutionary reasons that complement the logical ones. First, there is no storage medium in brains that can transcribe permanent coded information after the fashion of a silicon chip; all learned information in brains must be stored as dispositions to neuroelectrical and/or chemical *activity* that is merely implicit in patterns of organization when it is not active. (To suggest this point using a crude – too crude – metaphor, all knowledge *that*, at the level of neural coding, must be present as knowledge *how*, while nevertheless often being represented to its 'user', the whole human organism, as knowledge *that*.) Second, evolution cannot design a mechanism that resembles the Von Neumann compiler. That is, natural learning cannot be a process of writing messages in arbitrary code onto a register because such a code would be useless without an interpreter; and, as Daniel Dennett emphasized as early as 1969, we cannot posit intelligent interpreters as *parts* of agents lest we simply push the problem of explaining cognitive capacities deeper inside the system without ever *beginning* to solve it.[6] What is needed to understand biological learning is a mechanism that can simply react to stimuli, both external and internal, by 'tuning' its patterns of activation so that they bear systematic functional relations to patterns in the transduced information. And there can be no intelligent tun*er*; the brain must simply do it automatically. It is only once enough basic scaffolding is in place to allow the learning of a language (with information in discretely storable units of code), that the brain can simulate Von Neumann machines by using external storage sites such as pieces of paper or electronic computers, but the underlying operating system cannot itself use such systems for its procedural learning.

Over the past fifteen years, cognitive scientists have produced several sorts of mechanisms that have this character. (Clark [1987] provides a thorough and accessible account of one such approach, while Caudill and Butler [1990] competently survey the wider field.) None of these mechanisms are, in detail, plausible simulations of actual brains, but, at a sufficiently high level of abstraction there is little serious doubt that they instantiate the basic logical trick across which natural selection stumbled on each of the three independent occasions (in mammals, birds and octopi) at which it has generated intelligent learning systems. What is required are pathways of processing activation that are followed probabilistically, but in which response-probabilities change as functions of the degrees of success (relative, initially, to the organism's pre-set biological goals) of the behavioural patterns in which the neural patterns are causally implicated. We can illustrate the basic idea using a simple sketch of a 'toy' example – though the reader must be cautioned that it is *so* outrageously simplified that its relationship to any real biological example is to be found only at a very elevated level of

5 See the papers in Graubard (1988) for a comprehensive survey of these issues as they struck the participants just at the moment of this paradigm shift in cognitive science.

6 This, by the way, is a shattering and decisive objection to Freudian theory as a model of actual psychological mechanisms.

abstraction. Imagine a recently produced brain learning to distinguish the characteristic pitch and tone patterns sufficient for identification of a crucial source of utility: its mother's voice. For ease of exposition, we will set aside the contribution of its neurochemical apparatus and imagine that it solves the problem exclusively with its neuroelectrical equipment. (It could in fact do this; it's simply a biological fact, grounded in inherited physical processing constraints, not in computational logic, that it doesn't.) Now, suppose that this organism – let us make it a human baby for the sake of familiarity – is indifferently sensitive to all sounds within certain frequency and loudness bands. (This, by the way, is *not* true. Babies are born with the capacity to pick human voices out of the sonic jumble, but must learn to distinguish and reidentify voices of specific individuals. However, it can't harm our example as a pedagogical device if we make the baby's learning task a bit *harder* than it actually is.) Equip the baby only with a completely generalized link between behaviour and attentiveness to noises; that is, imagine that it initially reaches for and turns its face toward physically distinct sounds at random. We can now move inside its brain and describe what happens there as it learns to recognize mother.

From the surface-vibrations of the eardrum to the output of neuroelectrical patterns always caused by and only by discrete noises, the baby's fixed audio processing module will do a certain amount of work that cannot be over-ridden by general learning. The resulting neuroelectrical signals will cause some of the infant's synapses to 'spike', that is, to fire pulses down their connecting cords, called 'dendrites'. Each dendrite will lead to another synapse. Some connections will be inhibitory, making it less probable that the receiving synapse will emit a pulse. Others will be excitatory, making it more probable that the receiving synapse will emit a pulse. Depending on the detailed connections at the input point from the audio module, the first result of the sound will be an electrical 'storm' distributed over portions of the baby's brain. Perhaps the storm will have no systematic effect this time, its consequences swallowed amidst other concurrent storms triggered by other stimuli. But perhaps some of the connections reach synapses in the motor-control system, causing the baby to turn its head. Now, suppose this time the original signal was something uninteresting (though we have to work hard to think of something uninteresting to a human infant; they are voracious pursuers of information); perhaps the fridge has come on. The baby has turned its head as a *result* of the fridge's activity, but it would be premature to say that the baby has *attended to* the activity, since by hypothesis its movements are still random; no learning has yet established an association between any results of head-turning and noises. This time the head movement produces no distinct rewards (or discomforts), so the probability of this causal pattern of activity being repeated will remain as it was. Now change the stimulus. This time mother has cooed. Now the resulting head-movement brings a warm cuddle and/or a meal. Later, when the child's eyes have focused, it will be able to receive the reward evolution has programmed it to value most, namely a smile. At that point, still ahead on the learning curve, patterns of synaptic dispositions to excite one another will establish a direct associative link between mother's voice and mother's smile, so that seeking out the first will be a good way of getting access to the second. At that point, our subject will be well along the road to having learned a *concept* of 'mother'. Today, however, the system is still well short of such lofty heights of abstraction. The cuddle is a pleasant sensation. To the *brain* – and there is nobody else home here yet – this just *means* that signals

'back-propagate' through the system, strengthening the synaptic connections implicated in the random causal series that began with the sound and ended with the hug. Mother will unwittingly assist the learning process by intensifying the volume and frequency of coos when the baby 'looks at her'. (Note the scare-quotes here. *Mother* is guilty of some harmless anticipatory anthropomorphization at this point, since head-movements will eventually indicate that baby is looking at her, but don't yet, since baby doesn't now look *at* anything, and not just because it can't focus; there are as yet no *things* in its inchoate notional world.) Mother's response sends a flurry of possible paths for strengthening through the infant's brain, as a result of which its learning will be far more rapid than would be the case if its environment were passive. The probability of mother's voice causing the head to turn in its direction will rise until it approaches 1, though not in a linear way. If, on the fourth occasion mother speaks, the result is a blast of stale dog-breath, the connective dispositions will be inhibited. This is a set-back along the learning path, but a highly useful one, since the child must simultaneously learn *not* to do certain things, such as turn toward its (let us suppose) jealous and spiteful sibling. Once foeveael concentration is at hand, even doggie's barking will trigger stares and reaching, because the innate satisfactions of looking at animate beings will dominate aversion to bad breath. What crucially pulls the baby upward along the path to a coherent representation of the world amidst this potentially confusing flux of excitations and inhibitions are a host of redundancies *in the world* that the learning brain can exploit. As baby is learning to turn its head towards mother's voice it is also learning to reach toward that source, and this too brings happy effects. Synaptic patterns implicated in reaching will excite patterns implicated in turning. To the child's delighted observers, the resulting improvements in coordination of behaviour appears as the dawning of agency, as the emergence of a *person* – and the observers aren't wrong about this, though the parents will consistently project intentions ahead of the facts, thereby helping to pull their little learning system along into both self-construction and socialization. The gradually emerging result in normal cases really *is* an agent, a brain whose dense network of synaptic dispositions simulates the logical coherence of a deliberately designed software package, but has a level of adjustment flexibility to new stimuli – both external *and* internal – that no system of hard storage, such as a physical Von Neumann machine, could achieve in real time. Or, at least, *for a period of time* it has such flexibility. For as the child learns, its brain is acquiring habits, and many of these will be hard, or even impossible, to break.

Before we return to pick up this thread, which will lead us directly to the relationship between learning mechanisms and ideologies, we must pause for *some* (incomplete) cleaning up of the very hasty science implicit in the story just told. First, and most important, the mechanisms that actually support learning are a complex *medley* of interacting processes, not the monolithic assembly-line suggested by taking the above account too literally. Biological systems are built by blind trial-and-error winding through historically contingent adaptive landscapes, not by engineers. Second, the story elides over some points of considerable ignorance. For example, though we know (through a mix of behavioural and neuropsychological evidence) that brains use some kind of back-propagation, there are several possible mechanisms that might be involved in it, and very little is known about them or about how they might interact with one another. Third, as already noted, our story has assigned too much responsibility for ultimate behavioural patterns to neural learning, and too little to evolutionary

biases; connection-weights and patterns of connectivity do *not* begin as random. Fourth, we must be careful in extrapolating generalizations about learning in the prelinguistic context to situations in which culturally deposited information can be accessed. *Brains* do not need to learn many facts and principles that *people* nevertheless know. We will need to take this precaution particularly seriously when we turn our attention back to the formation of ideologies. All this being said, however, the sketch of infant learning of a simple association provides a basic model of all neural learning, especially as regards two fundamental principles of explanation in the domain: (1) We must not posit hard storage within the brain; all learning must be dispositional and all learned capacities, facts and biases must be coded dynamically; (2) Though some brain-processes certainly, and crucially, monitor others, the *basic* learning dynamics cannot be directed by applications of intelligent interpretation on pain of regress and emptiness of explanatory logic. We may summarize the implications of these two principles for cognitive learning in general as follows. Each stage of *propositional* learning, which essentially consists in finding access routes to entries in a 'cultural encyclopedia' mostly stored outside the brain itself, is built by and on a platform of non-propositional dynamic dispositions that organize subsequent data. Because of their inherited biological biases, brains do not begin as *general-purpose* computers. However, the degree of plasticity found in a particular brain must steadily shrink as it, in effect, loads increasingly special-purpose software that equips it for tasks – mainly social in character – presented by its environment. De-installation of such software will prove difficult for at least two reasons: First, its very non-propositional character renders it relatively invulnerable to rational argument. Second, it is constitutive of the individual *character*, the coherence of which is, as a result of biological pressures, among a typical agent's most basic and enduring self-maintenance goals. This second pressure is in need of elaboration, since it creates and maintains the strongest sort of link between neural learning and the formation of aspects of culture, including ideologies.

It is a truism that as the individual human organism matures, its behaviour is decreasingly predictable on the basis of biology alone, and increasingly sensitive to cultural variables. The acquisition of a public language makes for immense differences in the nature of cognitive capacities. Aside from the possibilities for fixed external storage of information already mentioned, language also brings capacities for cognitive self-manipulation unavailable to creatures whose signalling systems are biologically non-arbitrary and meagre in redundancy of storage-space.[7] People can implicitly remember vast complexes of information merely by knowing how to produce sentences that state propositions dense with entailment relations (both logical and pragmatic). Furthermore, they can talk to themselves, aloud or *sotto voce*, thereby triggering new associative learning in their own brains. These two capacities, which in fact imply one another, constitute the basic dynamics for *self-making*, that is, the learning of nested networks of relatively idiosyncratic dispositions that amount, taken in their entirety, to distinctive characters and personalities. Social pressures to produce *stable* selves lead to constraints on internal coherence in behaviour and judgement, in any given person, which are much tighter than those imposed by genetic biases, either individual or species-wide. This is the element of truth in existentialist claims that people create their

7 This point, along with those that follow, is extensively stressed and illustrated by Dennett (1991).

own characters, though it is *very* far from true that they do so autonomously. As Dennett vividly puts the point, each individual is merely the most persistent member of the team of authors that spins, over years, her character.[8] Self-making is subject, particularly, to the constraint that a biological individual is to have only one self. Actors and, to a lesser extent, public officials are permitted to depart from this maxim for certain instrumental purposes; but if these people cannot maintain 'core selves' to which they return when out-of-costume or off duty, they are regarded as ill.

Characters are sufficiently basic as units of social interaction, and in typical adults sufficiently complex, that it might be tempting to think of ideologies as aspects or even proper parts of them. Such a temptation should be resisted. Social pressure encourages characters to be comprehensible by others but, to varying extents in different cultures, unique and distinguishing. Ideologies, by contrast, are treated more like coordination devices, as indicated by the fact that many – perhaps most – people are less tolerant of variations within their ideologies that challenge for canonical status than they are of self-evidently distinct, even 'enemy', ones.[9] However, there are important similarities between characters and ideologies as cognitive-cultural constructs. People rely crucially on their characters to predict their *own* behaviour and to avoid the systematic disadvantages that accrue from acting on inconsistent preferences. We might refer to this as the 'book-keeping' function of character. It seems evident that ideologies play an exactly analogous role at the level of social action that must be coordinated. This can apply even across different ideologies, so long as they are reasonably standard; thus Yuri Andropov expressed relief when Ronald Reagan succeeded Jimmy Carter, and Nelson Mandela once expressed a preference (once there was no political risk in doing so) for the "straightforward" P.W. Botha over the "cunning" F.W. De Klerk. The first instance might involve some red herrings as an example of the point, since by 1980 the Soviets probably felt closer affinities to American leaders who simply fought the Cold War instead of complicating it further with muddled ethical intuitions. Mandela's remark is more clearly to the point here. It seems highly unlikely that Mandela had, in any relevant domain, more *in common* with Botha than with De Klerk, and he is obviously aware that Botha would have presented much more serious obstacles to his political goals. His comment should be read as simply a vivid reference to the difficulty of anticipating a political rival once the interpretive and predictive wedge of ideology has begun to blur.

Ideologies, like characters, cannot be reliable predictive devices unless they are actually thought to motivate thought and action. My character will not help me to predict my behaviour if I am convinced it is a sham, and it is famously unwise to rely on expectations of ideological consistency where skilled politicians are concerned.[10] Neither characters nor ideologies will be able to serve as trustworthy motivators if they are en-

8 It is important to reiterate, as noted earlier, that existentialist stress on autonomy is intended to make a set of *moral* points, not causal ones. One cannot use science to refute an ethical stance, though one *can* use science to show that a particular ethical stance may be silly and impractical.

9 The political history of communism provides particularly vivid, and repeated, illustrations of this claim; but almost any other well-studied political or religious movement will do as well. For an account of this phenomenon's game-theoretic logic, see Hardin (1995).

10 It has sometimes been suggested that the mistake foreign statesmen, including even the pathologically cynical Stalin, made with Hitler, lay in their taking his apparent political skill as evidence that his explicit, nutty, ideology was not serious. Did Churchill fare better in this case because his own grasp of routine politics was, as he often showed, unreliable?

tirely under the agent's own control, or if they motivate in the fashion of premises in arguments, always subject, like any other premises, to revisions in the face of evidence.[11] Characters, as noted above, are not even entirely under their authors' control during early stages of formation when their trajectories are still maximally plastic, and over time the weight of commitment implied by coherence constraints, and then limned by observable judgements and actions, can restrict the scope of control very sharply.[12] Character gains much of its reliability as a motivator from its rich web of implication in all action, including much that is non-deliberative. Ideology is somewhat different in this respect, since its scope is more restricted to deliberative, even intellectual, judgements and activities. Of course, fellow ideologues, keeping pervasive watch on group solidarity, take up much of the slack here. But this cannot be the basic mechanism; an ideology maintained only by external coercion, and doubted privately by all its cynical adherents, will soon go the way of the Soviet one. We thus arrive at the following puzzle. Ideologies, to be effective as political coordination devices, require stable anchors below the level of explicit propositional affirmation, since all explicit propositions except logical truths and simple observation reports are continuously subject to sceptical examination. This will be especially true of ideological propositions, which are typically under devoted siege by adherents of competing ideologies. But ideologies are highly propositional structures. So where is their internal cognitive anchor?

The answer, I suggest, lies in the tendency of ideologies to encode special semantics and even idiosyncratic logic. I return here to my opening remarks on my own inability to *think like* a Western Marxist; the casual phrase is perhaps not so casual. The special meanings attached to words – praxis, consciousness, liberation, dialectic itself – cause the circles of definitions that fix them, and the sentences in which they occur, to float outside of my integrated field of felt experience, and this is *despite* the fact that I spend many hours each week reflecting on the same phenomena as preoccupy the Western Marxists. Of course, I am quite sure that Western Marxists do not feel this way about their tropes and concepts; after all, they live them and think with them. On the other side, as a neo-classical economic theorist I am explicitly conscious of the extent to which I am teaching my students not a set of *facts*, but a complex (and, in this case, highly formalized) *logic*, in which they must try to become *fluent*. Indeed, I routinely *tell them this*, so that they understand the intended learning exercises. The most successful economics students are distinguished mainly by their having made this logic 'second nature' (*there* is a very striking idiom in the present context!), seeing any problem of human or animal motivation in economic terms 'at a glance' – that is, without needing time for extensive deliberation. Now, academic disciplines are not ideologies, though they are often closely related to them. Some sociologists 'speak' (and, hence, can think) Western Marxist without *being* Western Marxists, and though all orthodox economists can 'speak' capitalist, many are not capitalists. This distinction is possible because ideology involves some necessary connection with political *action*, even if the action consists only in expression of sincere normative judgements

11 See Elster (1999) for an extended defense and study of this fact.

12 One is reminded here of the comic English television character Reginald Perrin, whose forlorn attempts to escape from the constrictive effects of others' expectations of his character as a bland and successful businessman led him to repeatedly fake elaborate suicides on beaches in order that he could then attempt reinvention; the comedy lay in the mechanics by which these attempts always failed.

of certain sorts. But if ideologies are anchored in distinctive logical-semantic fields, and particular disciplinary matrices are anchored in those very same fields, then we should expect strong correlations between being a successful economics student and being a capitalist, on the one hand, and being a social theory student and being a Western Marxist, on the other, even when teachers scrupulously try to avoid deliberate indoctrination. And this is surely just what we do find. The neoclassical economist who is not a capitalist must use one logical-semantic field for thinking in the workplace and another for arguing politics in the pub and snorting over the newspaper; and while this is obviously possible it is not typical.

I have contended that entering into an ideology involves learning to think with a particular logical-semantic field. What the cognitive-scientific learning theory I have sketched shows is that learning a logic is importantly *not* like learning a set of facts. As I have discussed, the latter activity need involve very little *neural* learning; all that is required is that a few generic referential tokens be memorized as keys to bodies of entailments and addresses for hard storage sites outside the brain. But no one is said to have learned a logic if she merely knows where to look up its principles of inference. Similarly, a person who deliberately conducted himself in political arguments by looking up the meanings of words in Bottomore's *Dictionary of Marxist Thought* would be regarded as, at best, a painful sort of student, or, at worst, a victim of intellectual tyranny; he is not an ideological *practitioner*. So: becoming an ideologue is largely a matter of acquiring certain habits of thought. In light of what was said earlier about neural learning, saying this is no longer a matter of invoking a tired and empty figure of speech. Acquiring habits of thought is not something a *person*, as a participant agent in explicitly public and rational processes, can do by himself. He must recruit his *brain* – literally! – to the project.

One of the more remarkable facts about brains is that they can establish alternative associative pathways encoded using the *very same* synapses for use under different triggering circumstances. We have always known, of course, that many people are multilingual, and that excellent musicians can usually play more than one instrument at performance standard. It is only since we have been able to look at their brains using MRI and other recent technologies, however, that we have been able to confirm that separable but similar skills do not require separate storage. Now, this might suggest that *ideological bilingualism* is also possible. Is it? To return to my opening autobiographical setting, was I on the road to ideological bilingualism twenty years ago before I decided to concentrate on just one logic? I think that the answer to this question is 'no', for reasons that have to do *both* with neural learning and with politics.

I will first make the point through the cognitive science route. I will do this by setting up a fanciful image. I hope that I have said just enough about the details of learning in brains, however, to indicate that what is fanciful about my image is its simplistic view of philosophical history, not its computer analogy, which I think is quite literally apt. Let us imagine a brain, about twenty years old, that has learned a natural language and a number of useful and pleasing technical skills, acquired a distinctive character, been socialized in such a way that it is recognizably 'Western' in its cultural habits, but otherwise ideologically 'unset'. Adolescents in this society, let us suppose, are under little or no social pressure to make consistent political commitments, but this changes as they begin to be sized up for participation in formal institutional life. Making her political judgements in an unsystematic and ad hoc way, our agent finds that

she earns only the scorn of *everyone* whose political thought is guided by one or another set of roughly consistent principles. (She won't even make it as a politician this way; for that one must first learn to think with a coherent political logic and then learn how to rationalize actions that prima facie contradict it.) This needn't strike her *deliberatively* as a practical problem, if her environment is typical; perhaps she knows only that 'education', in some general, fuzzy sense, will make her 'seem more intelligent', that is, earn her various rewards that she probably conceptualizes more as conversational and broadly social than strictly mercenary. She takes a few courses in, let us say, philosophy, political science, economics and sociology. She will succeed in mastering all this material to passing levels only if she uses different logical-semantic fields for different assignments. For this, her *brain* will have to acquire a high level of logical flexibility; and, since it is a human brain, it will do that if she works it through enough learning exercises. Broadly, the hardware of her neural learning mechanisms acquires pathways that simulate either of two software packages. One takes any abstract concept it is asked to reflect on and rummages it for all of the external relations it can find. Some of these turn out to be contradictory, and every time the software finds such contradictions it splits the initial concept into two 'modes'; this gives it a new set of relations, the inter-modal ones, to link into its network of expanding complexity. Let us name this software HEGEL. Running it, our student can write essays in social theory and continental philosophy that earn good grades. Unfortunately, it performs disastrously in economics and analytic epistemology. Fortunately, our student's versatile brain has also learned RUSSELL. This software hunts down contradictions not to exploit them but to eliminate them from possibility. When it discovers that the coexistence of two concepts in an ontological frame creates a contradiction, one portion of the ontology must be explained as mistaken, using the conceptual resources of the remaining portion. Where HEGEL generates sweeping but very complex theories of large domains, RUSSELL outputs elegant models of constrained ones. Its products earn solid Bs in economics and analytic philosophy.

If our student never comes to care about the opinions on political subjects to which uses of these software packages lead her, then she can go on happily alternating between them as the circumstances demand. But in this case she is not ideologically bilingual because she is not really *ideological* at all. If she *does* feel commitment to some of her theoretical opinions, she will have a problem. Using the two packages on the same political problem will often – perhaps even usually – produce contrary, or at least entirely disconnected, opinions; and this is incompatible with commitment. If she wants A's in economics, or wants to go on to graduate study, she will probably find that the increase in dexterity with RUSSELL that she needs *automatically* brings with it, as something integral to its effectiveness, an aesthetic and/or moral *repugnance* for contradictions. But this internalization of a deep intellectual *value* will severely interfere with her ability to run HEGEL. Furthermore, the distaste she will feel for a fellow-student's HEGEL-generated essays – replete with dialectical turns and ontological proliferations – will cause a reciprocal annoyance in the emerging HEGEL pro. Why, the latter wonders, does our heroine insist on logic-chopping all his flights of criticism to death before they can get anywhere near synthesis? If our students continue all the way along their respective intellectual trajectories, they are likely *not* headed for professional conflict in the seminar room. They will ultimately not find themselves in the

same seminar rooms, because they will not be able to understand one another well enough to have arguments worth either's time.

Perhaps it doesn't matter much that our student will have to commit to one logical-semantic field if she aims at professional academic competence. As Turner would have emphasized, where ideology really *matters* is on the street or the shop floor. Could our student achieve *modest* ideological bilingualism – just enough, say, to listen with unbiased but critical attention to Andrew Nash *and* Jeremy Cronin *and* Themba Sono? It might be *possible* – I'm not after a metaphysical necessity here – but it is, I think, highly unlikely. The problem is that our student *must* care about her opinions if they're to be worth bidding for with arguments in the first place. Indeed, remember that her concern to acquire a political personality, even one just good enough for pub conversation, is what led her to learn HEGEL and RUSSELL to begin with. These very different styles of thought are deep enough to colour the whole semantic field in which serious political debate goes on – so much so that when I read Richard Turner talking about 'freedom', I can't bring myself to sincerely believe that his concept could denote any real state of a real human being because I can't *imagine* it.

Note that the account of the cognitive basis of ideology, and of ideological disagreement, I have given here need not be seen as *contradicting* more traditional ways of thinking about moral and political conflict.[13] It might be perfectly true and useful to say that the capitalist and the Western Marxist attend to different non-normative facts because of divergent self-interest, for example. After all, on my account a genuine ideology becomes partly *constitutive* of the self that uses it, so different ideologies will tend to correlate with different 'self'-interests in an unusually deep and literal sense. Furthermore, given that selves *are*, inter alia, devices for achieving social coordination stability against environmental pressures to update beliefs and values too often, it follows from my account that rival ideologues will notice or make salient different, but equally genuine, facts. I thus suggest that a Marxist (Western or Soviet) could agree with most of what I have said here. This is important, because, as I began by saying, my purpose has not been to try to overthrow Turner's normative landscape but to explain why I find it too alien to think in *without* having to suppose that *either* of us is necessarily wicked or thick-headed. This leaves it as open as ever to someone to try to show that one (or, for that matter, both) of us *is*, in fact, systematically overlooking morally significant facts or other considerations. Explaining the causal origins of ideological polarization at one level, the cognitive, does not imply that all other levels of analysis must be reduced away. It implies that cross-ideological political and moral debate is *hard* – I take it that a good account had *better* imply this obvious fact – but it doesn't imply that it's misguided or unimportant activity. I don't plan to give it up, and I'm not asking Turner's political heirs to do so either.

Postscript: It might seem to some that what I have defended in this lecture is a kind of smug political relativism, somewhat along Richard Rorty's line. Have I not, after all, suggested that ideologies are built on neural software, that a person can run one or the other, and the software setting in any given case is a product more of biographical accident than of reasoned consideration? It is true that I have not defended or attacked any ideology here; that hasn't been my purpose. I *do* think it is true that one cannot

13 I am grateful to an anonymous referee for this journal, who pointed out that my original text left this interpretation open.

usefully try to persuade adherents of a rival ideology using only the special resources available in the logical-semantic field of one's own. However, this does not imply that tests external to *both* ideologies can't be found. This is true both as between rival ideologies themselves and the software packages on which they run. For example, one can write proofs in formal languages, of the sort that underlie modern information technology, using RUSSELL, and one can do physics with it, and various other things one can't do with HEGEL. Is there anything, besides finding social-political engineering plausible, that one can do with HEGEL but not with RUSSELL? My ideological opponent might say that with HEGEL he can describe an inspiring vision of political community, whereas RUSSELL gets me only a squabbling bourgeois parliament of partial – in both senses of the word – opinions. But notice that he needs a parochial (to him) ideological referent here, namely, a notion of 'inspiring' I don't really understand. And, in any case, this has brought us to the level of the ideologies themselves. I will defend mine simply by claiming, on inductive grounds, that communities of RUSSELL operators have done a much better job of feeding and housing people – without having to murder a lot of them – than communities of HEGEL operators have done, anywhere or under any circumstances. If my opponent is a *Western*, as opposed to a Soviet Marxist, he might say that she cares about other things, like his special notion of 'freedom' perhaps. At that point, I will stand on the preferences of the electorate and see what happens. But let me emphasize that my point in this postscript has been *only* to disavow relativistic implications of what I have said. It would be disrespectful to the memory of the very good and noble man I am memorializing here to try to overthrow his ideology itself with merely a little afterthought of an argument. I have instead sought to pay him the best tribute I think I can, by devoting a serious effort to understanding why I so profoundly cannot follow his intellectual path despite the shining example of his moral one.

References

Bottomore, T. 1978. *A Dictionary of Marxist Thought*. Cambridge, MA: Harvard University Press.

Caudill, D., and R. Butler 1990. *Naturally Intelligent Systems*. Cambridge, MA: MIT Press / Bradford.

Churchland, P. M. 1979. *Scientific Realism and the Plasticity of Mind*. Cambridge: Cambridge University Press.

Clark, A. 1987. *Microcognition*. Cambridge, MA: MIT Press / Bradford.

Dennett, D. 1969. *Content and Consciousness.* London: Routledge.

Dennett, D. 1991. *Consciousness Explained*. Boston: Little Brown.

Elster, J. 1999. *Alchemies of the Mind*. Cambridge: Cambridge University Press.

Getty, J.A., and O. Naumov 1999. *The Road to Terror*. New Haven: Yale University Press.

Gopnick, A., and A. Meltzoff 1997. *Words, Thoughts and Theories*. Cambridge, MA: MIT Press / Bradford.

Graubard, S., ed. 1988. *The Artificial Intelligence Debate*. Cambridge, MA: MIT Press / Bradford.

Hacking, I. 1983. *Representing and Intervening.* Cambridge: Cambridge University Press.
Hardin, R. 1995. *One For All*. Princeton: Princeton University Press.
Lenin, V. 1902. *What Is To Be Done?* Harmondworth: Penguin.
Nash, A. 1999. 'The Moment of Western Marxism in South Africa.' *Comparative Studies in South Asia*, *Africa and the Middle East* 19: 83-94.
Ross, D., A. Brook and D. Thompson, eds. 2000. *Dennett's Philosophy: A Comprehensive Assessment*. Cambridge, MA: MIT Press / Bradford.