VII The State of SLA

23 SLA Theory: Construction and Assessment

KEVIN R. GREGG

1 Introduction

Anyone who has read the preceding chapters will agree that SLA is a terribly complex process, that understanding the process requires the contributions of numerous fields, from linguistic theory to anthropology to brain science, and that the process is not yet very well understood. In this chapter, we step back a bit from the trees, as it were, of the previous chapters, to look at the forest; to situate SLA within scientific inquiry in general, and within the field of cognitive science in particular.

I speak of SLA as a science here both without apology, despite the arguably quite meager empirical results so far obtained, and without invidious presuppositions about the relative intellectual merit of different objects and methods of study. The world is full of phenomena, only some of which lend themselves to scientific study. It may very well be that only a relatively small part of human nature falls within the area amenable to scientific study.[1] But language acquisition certainly seems to lie in that possibly constricted area, and to that extent I see nothing misleading or pretentious in talking about SLA as a scientific enterprise.[2]

But of course science is anything but monolithic. Scientists can differ not only in the objects of their research, but also in their epistemological stances toward those objects and their methodological stances toward the research, as well as in what they see as the important problems to be solved. In what follows we will look at some of the variation, possible and actual, in SLA theorizing.

2 Attitudes Toward Theories and Theoretical Entities

2.1 Three epistemological stances

To start with, SLA researchers, like other scientific researchers, can differ among themselves in the commitments they make as to what can in principle be known about the phenomena of interest and what the epistemological status is or can be of theories and the entities they posit. Very roughly, one can distinguish three positions: realism, empiricism, and relativism.

Realism is essentially the claim "that the characteristic product of successful scientific research is knowledge of largely theory-independent phenomena and that such knowledge is possible (indeed actual) even in those cases in which the relevant phenomena are not, in any non-question-begging sense, observable" (Boyd, 1989, p. 6). Empiricists would reject the term "knowledge," at least insofar as it is based on non-observable phenomena. "Knowledge" presupposes "truth," and empiricists claim that the most one can attribute to a theory is "empirical adequacy"; we are warranted in believing only what we can observe, although of course we are free to make use in our theories of constructs that go beyond the observable. An empiricist, in other words, can take an *instrumentalist* attitude toward theoretical constructs, using them to make predictions, for instance, but will withhold from them the status of real entities. A relativist denies the theory-independence of phenomena, and further denies, *contra* realists and empiricists alike, that knowledge or empirical adequacy is either actually or potentially of universal validity. Rather, theories are only true relative to some specific personal point of view, cultural or temporal context, Kuhnian paradigm, etc.

I don't propose to spend much time on relativist views of theory, SLA or otherwise, as there seems to be very little reason to take them seriously. For one thing, there are, to a first approximation, no scientists who take a relativist position. This is hardly surprising: it is inherently self-contradictory to conduct empirical research in order to reach conclusions that could be reached without all that bother, and which could not persuade, or even be comprehensible to, anyone outside the researcher's culture/paradigm/mindset. As Long put it, "it is not clear . . . why relativists would bother to do research at all" (Long, 1993, p. 230).

Of course, scientists themselves might be mistaken; they may be blind to the fact that they are not discovering facts but constructing them, as Latour and Woolgar (1986; Latour, 1987) argue (Latour and Woolgar themselves are evidently immune to this blindness; *they* have not constructed facts about how scientific research is done, but discovered them, on the basis of objective observation). Thus, more important than the fact that scientists, including SLA researchers, do not conduct research within a relativist framework is the fact that no one has provided any convincing reason to think that relativism is a defensible, or even a coherent, epistemology. As Brueckner says, "It is difficult to formulate an even remotely plausible view that deserves the title *conceptual relativism*" (1998, p. 295). (For detailed discussion of the problems of relativism, see Laudan, 1990, 1996.)[3] So, since no one has given us any reason to deny the claim, supported by realists and empiricists alike, that "there are some hypotheses and some logically and nomologically possible states of affairs such that we're absolutely warranted in believing the hypothesis if we find ourselves in the indicated state" (Kukla, 1998, p. 112), and since virtually all SLA research takes that claim for granted, we may safely give relativism short shrift, and concentrate on the differences between realist and empiricist takes on SLA.

2.2 Theory and observation

Perhaps the fundamental question dividing realists and empiricists is the theory/observation distinction. For the empiricist, remember, we are warranted in believing only observational statements (although we may, of course, agnostically employ theoretical ones as well). For this claim to go through, however, there must be a non-arbitrary way to distinguish observational statements from theoretical ones. But it has long been argued that any such distinction is inevitably arbitrary, that observation is, as they say, *theory-laden* (Hacking, 1983; Hanson, 1958; Kuhn, 1970; Maxwell, 1962), and these arguments do not seem to have been satisfactorily refuted (Sober, 1994a; see Kukla, 1998, for extensive discussion). Now, if empiricists cannot convincingly maintain the theory–observation distinction, then they have no principled ground on which to withhold belief in the existence of theoretical entities. This becomes particularly germane, perhaps, in the case of sciences like linguistics, which posit entities, such as Universal Grammar, that are on anyone's account unobservable in principle.[4]

If we accept the idea of the theory-ladenness of observation, we can run with it in a couple of very different directions. One direction leads to the relativist claim that theory-neutral observation is impossible: if even the most innocent-seeming observation is tainted with theoretical presupposition, so the reasoning runs, then two observers, starting with different presuppositions, will not be in a position to agree about what inferences from that observation are legitimate. If, as the relativist maintains, there is no theory-neutral observation, we cannot expect there to be observational statements whose veridicality could be accepted by all rational observers. Indeed, on a radical interpretation of theory-ladenness, observers who don't share the same presuppositions actually observe different things.

Fortunately, theory-ladenness simply does not, *pace* the relativists, entail the impossibility of theory-neutral observation. For one thing, as Kukla points out, there is simply no reason to think that our observations are affected by our beliefs in anything like the degree assumed by this relativist position: "To show the impossibility of theory-neutral observation, one would have to establish that *all* cognitive differences have an effect on perception – and this goes beyond what the New Look research program has established on even the most sanguine reading" (1998, p. 115; cf. J. A. Fodor, 1983).[5] And in fact there are countless observational statements about whose veridicality no rational observer disagrees.

For another thing, there is usually no reason to think that a given observation is "laden" with the particular theory being tested by that observation (Hacking, 1983; Nagel, 1997). Nor do we need to demand that *all* observations be neutral relative to *all* possible theories; it is enough (but it is essential) that the observations at issue be neutral with respect to the two or more theories that are being tested (Sober, 1999). Observation of cell mitosis, or of the surface of Mars, is "laden" with various theories from optics that explain how microscopes and telescopes work, for instance. But those theories don't affect the observations, although they may, of course, affect the interpretation of the observations, which is a different question: whether the lines we see are canals or not doesn't alter the fact that we – all of us – see lines. We may refuse to accept the results of a grammaticality judgment test because we think there's no reliable causal connection between the subjects' knowledge and their responses, or because we reject the grammar-theoretical categories being tested, or because we disagree with the judgments; but we will not disagree as to what was in fact observed, which was a set of marks on paper.

This leads to a final, important point about observation, theory-laden or other: as Bogen and Woodward (1988) argue, if "we use 'observe' to mean 'perceive' or 'detect by means of processes which can be usefully viewed as extensions of perception,' then scientific theories typically do not predict and explain facts about what we observe" (p. 305). What we actually observe in a grammaticality judgment test, for example, is the subject making marks on paper or punching keys on a computer keyboard; we do not observe grammaticality judgments. We infer (with a very high degree of confidence, of course) from the observed acts to the judgments, and (with a good deal lower degree of confidence) from the inferred judgments to the hypothesized grammatical knowledge. On the other hand, we want to predict (and explain) not the observable markings and punchings but the judgments, which we can't observe. It is these unobservable *phenomena*, not the observable *data*, that are the objects of inquiry; thus "it is a mistake to think of claims about phenomena as theory-laden observational claims" (p. 315; cf. Woodward, 1989; see Gregg, 1993, for SLA-related discussion).

In short, to say that observations are theory-laden is not by any means to say that objective comparison and assessment of theories are impossible. But there are other inferences one can make from the claim that there is no criterion or algorithm for distinguishing theoretical statements from observational statements. As we saw above, denying the theory–observation distinction opens up the possibility of rationally accepting the existence of "purely" theoretical – that is to say, unobservable – entities, which is precisely what the realist does.[6] But since (unlike relativists) realists do not believe that "anything goes," they must give us some sort of criteria for deciding when a given unobservable construct warrants our belief. Since the criterion of observability is obviously out, it follows that the realist is committed to appealing to, or at least allowing the appeal to, non-empirical virtues such as simplicity, explanatory power, or inference to the best explanation as criteria for preferring one theory over another. This raises a range of problems, as we will see below.

2.3 *Realism and empiricism in SLA*

SLA researchers are not given to publishing their epistemological allegiances or arguing about issues in the philosophy of science. (But see, e.g., Beretta, 1991;

Beretta and Crookes, 1993; Gregg, 1989, 1993; Long, 1990a, 1993; Tarone, 1994.) Even those SLA textbooks that devote some space to more general questions of theory (as opposed to simply presenting and comparing various theories), such as Larsen-Freeman and Long (1991) or McLaughlin (1987), are not that explicit.

One could perhaps characterize the majority of SLA researchers as holding to what Kukla (following Leplin, 1997) calls "minimal epistemic realism," the belief that "it's logically and nomologically possible to attain to a state that warrants belief in a theory" (Kukla, 1998, p. 11). Note that this position does not imply any strong realist commitment. It is indeed a minimalist position, little more than an articulation of the common ground shared by realists and empiricists (notably, the rejection of relativism and its works), leaving open most of the questions that divide those two camps.

Foremost among those questions is the role of non-observables. In SLA, as in all areas of psychology, opinion differs as to how far we should be willing to attribute causal powers to distinct, but of course unobservable, elements of the mind. Empiricist psychologists begrudge every such attribution, wishing to appeal wherever possible to the environment; hence, for instance, their peculiar insistence on operationalization.[7] Realists are perfectly at ease with a proliferation of mental elements, and willing to justify them on theoretical rather than operational grounds. We will see how this opposition plays itself out in SLA below, when we look at SLA property theories.

3 The Domain of an SLA Theory

SLA theorists can differ not only in their epistemological commitments, but in their view of the domain of SLA theory: what is an SLA theory a theory of?

On the face of it, this would seem to be a fairly simple and uncontroversial question: SLA theory is the theory of the acquisition of a second language. Since acquisition is at least something like learning, if not in fact the very same thing, it should follow that SLA falls within the scope of cognitive science, as opposed to social science. SLA research is thus first and foremost an *internalist* rather than an *externalist* discipline, to borrow terms from Chomsky (1995). That is to say that we are primarily concerned with learner-internal changes of state, not with the behavior of learner groups (or even of individual learners), and not with abstract "learner languages," or "E-languages," to apply Chomsky's (1986) term to SLA.

Of course, just because this may seem to go without saying (or at least *with* saying: e.g., Gregg, 1989, 1990) doesn't mean that it has been accepted without question. Firth and Wagner (1997), for instance, make the bizarre claim that current SLA research is too cognitive, although they fail to offer any reasons for changing the direction of research in the way they seem to favor (see commentary by Gass, 1998; Kasper, 1997; Long, 1997). No one, however, has presented a coherent argument against the position that second language acquisition involves individual mental states and their changes, so I think that we can accept that position as a working definition of the domain of SLA theory.

This does not, I stress, mean that externalist L2 research is inconsequential, uninteresting, misguided, or irrelevant. And it certainly does not mean that learner behavior can be ignored by researchers trying to explain SLA. The point is simply that we must distinguish between evidence for an SLA theory (learner behavior) and the object of that theory (learner mental states).

3.1 Natural kinds

One way of comparing internalist and externalist approaches to SLA theory is to consider whether they can identify natural kinds among the objects of study. Ideally, that is, a theory should be able to pick out a set of objects such that, however it is defined, one can make interesting generalizations that apply to all and only the members of that set. Biology distinguishes, say, between mammals and fish, but not between terrestrial and aquatic animals, because once you've identified something as a mammal you can predict all kinds of things about its physiology regardless of where it lives, whereas identifying something as an aquatic animal tells you little more of biological interest than that it's an animal and that it lives in the water. ("Interesting," of course, means scientifically interesting, not culturally. The distinction in Jewish law between clean and unclean animals is interesting, but not to the biologist. Chemistry recognizes a class of heavy metals, but not a class of precious metals, no matter how much more interesting you and I find the latter.)

The question, then, is whether an externalist or an internalist approach to SLA is more likely to be able to distinguish natural kinds. Non-native utterances, for instance, are often ungrammatical from the point of view of the target language, but then so are some native utterances, and many if not most non-native utterances (if it even makes sense to quantify them) are grammatical. So it's hard to see how the set of non-native utterances could be characterized so as to distinguish it from the set of native utterances, let alone how we could go beyond the definition of non-native utterances to make other claims than that they are utterances made by non-natives. Again, many L2 learners learn their L2 primarily in classrooms, and thus could fall within the domain of a pedagogical theory, for instance, or a sociological theory about power relations in schools; but then other students learn, and other subjects are taught, in classrooms on the one hand, and many L2 learners learn the L2 outside of classrooms on the other. Many L2 learners are immigrants learning the L2 on the job, and as immigrants in low-paying jobs are the victims of oppression and discrimination; but then many aren't immigrants, and there are many natives who suffer oppression and discrimination. In short, it appears that it will be hard to identify the class of L2 learners in a way that could define them as a natural class for an externalist SLA theory.

Of course, we are not by any means guaranteed that an internalist SLA theory will do better, but the possibilities at least seem a bit more promising.

We need to ask whether the set of L2 learners – or rather, the set of L2 learner grammars – regardless of L1 and regardless of target language, constitutes an interestingly distinct natural kind, different from the knowledge states of learners in general, and from the grammars of L1 learners in particular. And we need to ask whether adult language learners constitute an interestingly different kind from bilingual child language learners. These questions are all still open – albeit to varying degrees, perhaps – and the answers may turn out to be "No." Language acquisition, first or second, could turn out to be nothing other than learning, in which case the class of L2 learners would be of no more specific scientific interest than the class of mathematics learners. More plausibly, perhaps, adult SLA could turn out to be the same as child language acquisition, in which case L2 learners could be merely a source of dirtier than normal data. But at least the possibility seems to remain that there are interesting things to say about the mental states of adult L2 learners *qua* adult L2 learners that one cannot say about children or about learning in domains other than language.

3.2 Idealization

If we are working toward an internalist theory, it may be objected, a theory of mental states and changes of state, while at the same time tentatively assuming that the set of adult L2 learners forms a natural kind, are we not ignoring the seemingly gross variation that obtains across learners? Yes, that's exactly what we're doing. *Any* theory, as a matter of course, idealizes over its subject matter. The very idea of a natural kind presupposes certain attributes shared by all the members and by them only; the problem for the theorist is to identify those common attributes that specify the kind. But that means that, for the purposes of specifying the kind, we can and should ignore variation across members. Once we are in a position to identify what distinguishes the class of L2 learners from other learners, we are in a better position to characterize and explain the variation among L2 learners.

We are tentatively allowing for a couple of different possibilities, which need to be confirmed or rejected on empirical grounds:

i The possibility that nativelike competence can be attained. This view is consistent with, although it doesn't necessarily entail, the position that SLA is essentially the same process as child first language acquisition, and just as first language acquisition theory assumes an ideal learner who attains perfect competence, so would an SLA theory. The self-evident disjunction between this idealization and the reality of SLA variation would then need to be explained (or explained away) by secondary, extrinsic causes: quantitative or qualitative deficiencies in input, motivational, or other affective variation, degrees of acculturation, what have you.

ii The possibility that there are one or more (relevant) universal differences between the initial state of adult learners and that of child learners, hence universal (minimal) "deficits" in final L2 competence. Here again we would be postulating a uniform ideal final state, albeit one that differs from the final state attained by the L1 native speaker. (If we don't mind committing the "comparative fallacy" (Bley-Vroman, 1983), we could say that non-native learners "fail to acquire the target language completely," or have "imperfect L2 competence.") Thus, one could posit the effects of input and affect as in (i), while claiming that even if, in the ideal situation, these were all overcome, there would still be differences between the final state of the ideal native and that of the ideal non-native. One could argue, say, that the adult L2 learner has lost one or more specific learning mechanisms used by children (O'Grady, this volume; White, 1989); or that the adult L2 learner is not able to fix parameter values for the L2 (to "reset" parameters, as it is often put) (Eubank and Gregg, 1999; Hawkins and Chan, 1997; Schachter, 1996; Smith and Tsimpli, 1995; Strozer, 1994).

Note that on either view, (i) or (ii), we are idealizing away from the actual variation that one can observe across individual learners; note further that this is exactly what we should want to do. Take parameter-setting, for one example. One could consistently claim that all L2 learners are different from natives, and identical to each other, in one specific respect – inability to reset parameters – while allowing for, indeed predicting, wide variation across learners according to what specific L1–L2 parametric differences obtain. One might predict, for instance, different L2 English competences – and hence different behaviors – with respect to expletives and pronoun use depending on whether the L1 was a pro-drop language or not. At the same time, by positing a certain uniformity across learners – a uniformity that is, moreover, not observable – we have the possibility of making a principled, testable distinction between possible and impossible variation (or predictable and non-predictable) and thus have a potential means of explaining variation, rather than merely describing it. Indeed, failing to idealize in this way virtually guarantees the sort of theoretical sterility found in much of the SLA variationist literature (Gregg, 1990).

4 SLA Property Theories

Given an internalist perspective – given, that is, that we are hoping to explain the internal state of an individual learner with respect to an L2 – it may not be too question-begging to assume that an SLA theory will characterize the L2 *knowledge*, or *competence*, of an idealized learner. Those terms, of course, have been the object of a great deal of contention, but not from within the internalist perspective, where perhaps the only principled rejection of the terms would come from those who prefer to talk of dispositions to behave rather than of knowledge. The arguments against such a view are well known (see, e.g., Chomsky, 1959, 1980a, 1986), and there's no need to rehash them here. In short, an SLA theory needs to explain the knowledge state of the L2 learner vis-à-vis the L2. Indeed, it needs to explain at least two such states: the initial state, immediately preceding first exposure to L2 input, and the final state, after which input ceases to have any significant instructional effect.

A theory of this sort is sometimes referred to as a *property theory* (Cummins, 1983; for SLA, see Gregg, 1993, 1996a, 2001). A property theory answers the question, "In virtue of what does system S have property P?" (Cummins, 1983); it explains the instantiation of a property in a system. Thus, for example, theories of dispositions – acidity, solubility, heritability, etc. – are property theories. Linguistic theory is an excellent example of a property theory, answering the question of how linguistic knowledge is instantiated in a mind. Property theories do not account for sequential processes – these are the domain of a *transition theory*, which answers the question, "How does system S change from one state to the next?" But this is not to say that property theories are not causal. To claim, for example, that such-and-such a sentence is ungrammatical by virtue of the Empty Category Principle is to claim that there is a causal relation between the ECP and the ungrammaticality. To put it somewhat differently, a property theory that appeals to the ECP in this way is claiming that the ECP is real; not just "psychologically real," whatever that peculiar phrase is supposed to mean, but real: "To be real is to have causal efficacy; to be unreal is to be a mere artefact of some causal process" (Sober, 1994b, p. 220). This is what makes property theories explanations, not mere descriptions.

Property theories of SLA can vary on any number of parameters, but we will look at the following:

i *modularity* (section 4.1): is L2 knowledge in any interesting way modular?
ii *innateness* (section 4.2): is L2 knowledge in any interesting way attained or possessed independently of environmental influence?
iii *the nature of L2 representations* (section 4.3): specifically, does L2 knowledge consist in a hierarchically ordered, structured system of representations, or is it distributed across essentially unstructured representations?

4.1 Modularity

A module is a comparatively autonomous subsystem within a larger system, which acts more or less independently of other subsystems, and has structures and functions that are more or less recognizably different from those of other subsystems. Cognitive science recognizes a couple of different senses of modularity. One difference is in the level of analysis: modularity at the anatomical level vs. modularity at the functional level. A claim of *anatomical* modularity for L2 knowledge would be a claim that L2 knowledge is localized in a specific, well-defined area of the brain. Such a claim, though, stands or falls independently of a claim of *cognitive* modularity, the claim that L2 knowledge, however instantiated physiologically and wherever located, is a module within a larger system of knowledge. The mutual independence of these two modularity claims needs to be stressed, as it is often overlooked in the literature. If, for instance, we were to find that all L2 performance – silent reading, conversation, listening, etc. – activated one specific corner of the brain and no other, that would certainly be suggestive evidence for the cognitive modularity of L2 knowledge. And if that L2 corner were different from the L1 corner, it might suggest that L1 knowledge and L2 knowledge were cognitively different. But such a conclusion would not automatically follow, any more than the conclusion that the books on the third floor of the library stacks are categorically different from those on the first. And by the same token, just as books on the same subject may be shelved in two widely separate locations simply according to age or size or date of acquisition, so would the discovery of multiple "L2 areas" in the brain be consistent with L2 as a cognitive module.

The question of anatomical modularity ("localization of function") is of course an interesting one, but for SLA as a cognitive-scientific discipline, it is secondary to the question of cognitive modularity. As Coltheart and Langdon point out (1998, p. 151), "even if there is anatomical modularity, if the module in which one is interested itself has an internal modular structure, each of these submodules may well be instantiated in a different part of the brain." Coltheart and Langdon go on to draw an important conclusion, one that is often misunderstood in the SLA literature (e.g., Jacobs and Schumann, 1992): "That is why the development of an adequately fine-grained abstract theory of the structure of cognitive systems must precede any attempts to map the neural substrate of cognition."[8]

Putting aside anatomical modularity, we can perhaps distinguish between two different (but mutually compatible) understandings of cognitive modularity, what we might call Chomsky-modularity and Fodor-modularity (see Schwartz, 1998, 1999, for discussion). L2 knowledge would be Chomsky-modular if it is part of a hypothesized language module. The language faculty is modular in that, and to the extent that, it comprises structures and conforms to principles not found in other modules: binding principles, say, or c-command, or the Subset Principle. L2 knowledge would be Fodor-modular (J. A. Fodor, 1983; Schwartz, 1986) if it is (to a significant degree) *cognitively impenetrable* and *informationally encapsulated*: that is to say, if the processing of linguistic input is not significantly affected by or accessible to higher cognitive functions (beliefs, say) or by other input systems.

Some of the contributions to this volume reflect the various possible stances one can take on L2 modularity. UG/SLA positions, for instance, assume modularity for language as a whole and extend that assumption to L2. But one can with consistency claim that L1 knowledge is modular (in whatever sense) while L2 knowledge isn't (Bley-Vroman, 1990). (For that matter, it's logically possible to claim that L2 knowledge is modular while L1 knowledge isn't, although I can't imagine anyone making such a claim.) A "cognitive nativist" position such as O'Grady's (1996, 1999b, this volume) rejects at least the strong Chomsky-modularity claimed by UG theorists. Although allowing some language-specific principles and mechanisms (such as the Subset Principle), the thrust of cognitive nativism is away from Chomsky-modularity (as reflected (e.g., O'Grady, 1996) in O'Grady's earlier term for his position, "general nativism").

It is, of course, hard to say in principle when a cognitive system is or is not modular "to an interesting degree," but connectionism is certainly anti-modular, at least in practice. This anti-modularity, it is worth noting, is not a logically necessary one. One could have a language module in which linguistic knowledge is acquired by connectionist learning mechanisms, for instance. And, as Ramsey and Stich (1991, p. 308) say, "If the best connectionist models of language acquisition exploit a learning algorithm that is particularly adept at language learning and largely useless in other domains, then again rationalism and connectionism will turn out to be comfortably compatible." But in fact most connectionists would probably reject this possibility (Broeder and Plunkett, 1994) and deny that language has any interesting domain-specific components, such as the principles and parameters of linguistic theory. Just what it is that connectionists think linguistic knowledge *does* consist of, however, is another question, and one that is not at all easy to answer; we'll return to this problem below.

4.2 Innateness

"Innateness" is an infelicitous term, and indeed "many biologists consider the concept of innateness to verge on incoherence" (Sterelny and Griffiths, 1999, p. 6; cf. Ariew, 1996, 1999; Wimsatt, 1986, 1999). Since no one thinks the mind is truly a *tabula rasa*, and no one thinks there are genes for foreign languages, the question is to what extent acquisition of an L2 depends on knowledge that exists independently of environmental input, *and* that applies specifically to the domain of language. Everyone, that is, postulates *some* innate component to language acquisition; at issue is to what extent the innate component is domain-specific, and to what extent the domain-specific component is innate. A comparison may be helpful: knowledge of baseball includes domain-specific knowledge, such as knowledge of what a squeeze play is, of when to throw to first base rather than second, etc. This domain-specific knowledge is learned, not innate. On the other hand, there is an innate component to baseball; bipedal movement, for instance. But clearly, running in baseball is just running; an innate capacity, but not domain-specific. Here again, the point where this "innateness" becomes "interesting" cannot be determined in advance, but useful distinctions can nonetheless be made among SLA theories.

As with modularity, UG/SLA theories stand at one extreme. Although other innate knowledge enters into language acquisition and use – for instance, the "mindreading" capacities (Baron-Cohen, 1995) that enable us to interpret the intentions underlying the utterances of others – the foundation of language knowledge is UG; and UG is innate and domain-specific. Also domain-specific, but not innate, is the peripheral information that varies from one natural language to another – the knowledge that "automobile" is used to refer to automobiles, for instance, or the knowledge of honorifics. But the essence of linguistic knowledge – the principles or mental structures that characterize the language faculty and distinguish it from other mental faculties – is, for UG/SLA theorists, innate. Cognitive nativism would reduce, perhaps drastically, the amount of domain-specific innate knowledge in language; the key properties of the computational faculty, for instance, while innate, are shared, on this view, with mathematical knowledge. Domain-specific non-innate knowledge is of two kinds: the sort of specific learned lexical information for a given language, as in UG theories, but also derived, "module"-specific knowledge. (O'Grady's idea that grammar is a "new machine built out of old parts" (1997, p. 328; quoting Bates and MacWhinney, 1988, p. 147) is reminiscent of Karmiloff-Smith's, 1992, idea of, in effect, learned modularity.)

On the face of it, connectionism would seem to deny the domain-specificity of innate knowledge totally. What is innate, presumably, is merely the general learning capacity that inheres in the system of nodes and the susceptibility of their connections to strengthening and weakening. On this view, linguistic knowledge is almost entirely learned; there are no underlying rules or principles or structures that obtain only for language.

The problem, though, is where do the nodes come from? A connectionist simulation starts with elements of some sort on which to base the growth of a distributed system; say, lexical items and plural forms as in Ellis and Schmidt (1998), or gender markers as in Sokolik and Smith (1992). But what we don't know from the simulation is how those forms were themselves acquired. Does a learner have an inborn concept of plurality, say, or gender? Since gender, at least, is a purely formal (i.e., domain-specific) concept, connectionists would seem to be committed to denying its innateness; yet nothing is said about how the concept of gender (*mutatis mutandis*, plurality, tense, etc.) is learned. But the problem of language acquisition, as Fodor says, "is that of how a child acquires grammatical structure, not how he learns *correlates* of grammatical structure" (J. A. Fodor, 1998e, p. 150). In the absence of specific connectionist proposals about such structures, we seem to have nothing to replace nativist theories such as UG. After all, if they weren't acquired, they must be innate, which is hardly what a connectionist should want to claim. Hence, for instance, Carroll's criticism of Sokolik and Smith, namely that their results merely "show that if the learning device is given *a priori* means to solve a given linguistic learning problem, it does very well indeed. This is just the claim innatists make" (Carroll, 1995, p. 202).

4.3 L2 representations

The crucial distinction between SLA theories here is whether or not they assume that the mental representations of L2 knowledge are structured. Most theories at least tacitly assume some sort of so-called "classical" view of knowledge representation, such that knowledge (e.g., L2 knowledge) is organized in a highly structured system of representations (e.g., a syntax); UG theories, of course, make that structure fairly explicit. Connectionists, on the other hand, generally see L2 knowledge as instantiated in unstructured, distributed representations. (Cf. J. A. Fodor, 1998a, p. 11 fn. 6: "Connectionists are committed, willy-nilly, to *all* mental representations being primitive.")[9] To be more precise, this is the position taken by those connectionists – often referred to as "eliminativist connectionists" – who see connectionism as offering a rival account of mental representation to the "classical" account. Rey refers to this strong form of connectionism as RCON (radical connectionism) to distinguish it from LCON (liberal connectionism), the view of connectionist processes as merely implementing, rather than eliminating, a classical representational system. "The crucial feature that distinguishes RCON from LCON . . . is the claim that there is no causally efficacious constituent structure to the mental representations that play a role in the processing" (Rey, 1997, p. 227). It has often been pointed out (e.g., J. A. Fodor and Pylyshyn, 1988; Rey, 1997; Sterelny, 1990) that a connectionist learning process à la LCON is consistent with a structured, classical representational architecture; but, as Broeder and Plunkett (1994) suggest, most connectionists are not content with that role. Certainly SLA connectionists seem to lean toward RCON (e.g., Ellis and Schmidt, 1998), and in any case LCON is of no interest to us here as a property theory of L2 knowledge, since it doesn't provide an alternative to classical theories like UG/SLA. In what follows, then, I will use "connectionist" and "connectionism" to refer to RCON, or eliminativist connectionism.

What's at stake in the choice between structured and unstructured representations? The fundamental problem is the one raised by J. A. Fodor and Pylyshyn (1988): the problem of *systematicity*. Briefly, it is uncontroversial that in any natural language, if that language allows a sentence of the form, say, "John loves Mary," it will also allow "Mary loves John" (similarly, anyone who can think that John loves Mary – in effect, anyone – can think that Mary loves John). This sort of fact is easily enough explained if you allow syntactic categories and rules that control them: noun phrases are structures that can fill certain roles within a larger structure (a sentence), and if X is a noun phrase, then by virtue of its category membership it can play those roles, etc. Put somewhat differently, the undoubted systematicity of language can be explained if it is *nomologically necessary*: systematicity (syntacticity) is a necessary condition on being a natural language.

On a connectionist account, on the other hand, it would seem that this sort of systematicity is purely contingent: it just so happens that all humans have this capacity. In the absence of appropriate input, it should be perfectly possible for there to be a human who can say "John loves Mary" but cannot say "Mary loves John." This seems, to say the least, counterintuitive; as Fodor says, "I think we had better take it for granted, and as part of what is not negotiable, that systematicity and productivity are grounded in the 'architecture' of mental representation and not in the vagaries of experience. If a serious alternative proposal should surface, I guess I'm prepared to reconsider what's negotiable. But the prospect hasn't been losing me sleep" (J. A. Fodor, 1998a, p. 27).

The systematicity debate rages on, and at least some connectionists have recognized it as an important challenge to connectionism (e.g., Clark, 1993),[10] and have taken it seriously enough to try to overcome it (e.g., Smolensky, 1987, 1995; for responses see J. A. Fodor, 1998c; J. A. Fodor and McLaughlin, 1998), but the consensus so far seems to be that this challenge has yet to be met. Aizawa, indeed, goes so far as to say (1997, p. 126), "Independent discoveries by future science might one day vindicate Connectionism against Classicism, but what future science will not change is the fact that Connectionism cannot explain the systematicity of thought. At most, future science can show that Connectionism is true, *despite* its inability to explain the systematicity of thought" (but cf. Cummins, 1996). In any case, the systematicity problem seems to have been largely ignored in SLA theorizing, and it is hardly likely that it will be resolved there.

5 SLA Transition Theories

Where the SLA property theory will explain the nature of the cognitive states of an L2 learner, the SLA transition theory will explain the causal processes that effect changes in those states such that L2 knowledge becomes instantiated in the learner's cognitive system. As with the property theory, the transition theory should be an idealized one, abstracting away from a specific L1 and L2 or from a specific group of learners. To borrow terms from Sterelny and Griffiths (1999), a general SLA transition theory should aim for a *robust process explanation* rather than an *actual sequence explanation*:

> *Actual sequence explanations* seek to explain the nuances of the causal history of the world we find ourselves in. They explain the contrasts between our actual history and the histories of the nearby possible worlds. For such purposes, the more fine-grained the explanation, the better . . . *Robust process explanations* reveal the *insensitivity* of a particular outcome to some feature of its actual history. Thus an explanation of World War I that appeals to the political divisions of Europe is a robust process explanation, seeking to show that some World War I-like event was very probable. The detailed unraveling of diplomatic and military maneuverings is an actual sequence explanation, showing how we got our actual World War I. (p. 84)

As Sterelny and Griffiths go on to point out, the two types of explanation are not rivals, and each has its own legitimacy. However, in so far as we are trying to formulate a theory of L2 acquisition as such, and not simply an account of how certain learners acquired a certain L2, we need a robust process explanation. As with the case of the property theory and variation in final states, once we have something like a robust process explanation we should be in a better position to offer actual sequence explanations, to account for the various specific deviations from the ideal process that are actually attested.

On any account, the result of SLA is a set of representations of the L2, however different they may be from the native speaker's representations. Since those new representations are representations of the L2, and vary rather neatly according to what L2 input is provided – you need input of English to get representations of English – it seems a safe bet to assume that input is the major causal factor in SLA. (I am using "input" here in the atheoretical sense in which it's generally used in the SLA literature, viz. to refer to the *utterances* of speakers other than the learner, heard (or read) by that learner. The actual characterization of the input to the learning mechanism depends on the property theory being assumed. See Carroll, 1999, forthcoming, for detailed discussion.) This assumption is all the safer given that we are, tentatively at least, restricting ourselves to adult SLA, and hence can eliminate maturational processes as causal powers in the forming of an L2 representational system.

Finally, since our transition theory is an internalist one, we will need to posit some sort of mental mechanism – a learning mechanism – that can act on the input to create the representations. There will no doubt be other internal causal factors – motivation, for instance – but these will necessarily be secondary, for the simple reason that they cannot themselves process linguistic input. Motivation can directly affect the amount and frequency of input, for instance – by getting the learner to go to class and pay attention, say – but motivation alone cannot tell a noun from a verb, let alone parse a sentence or set a parameter value.

In a word, an SLA theory minimally must account for the role of input and must provide for a learning mechanism to create L2 representations based on that input.

5.1 Learning mechanisms

Learning has generally been taken to be an inductive process of trial and error. Hence the often-used term, "hypothesis-testing": on the basis of environmental stimuli, the learner (consciously or unconsciously) makes tentative hypotheses, which are then confirmed or disconfirmed by further stimuli. Certainly some forms of language learning are inductive on anyone's account. A child hears a few examples of verbs in the past and present forms, and finally induces a rule of past-tense formation, say. Of course the term "rule" is highly tendentious; a connectionist, no doubt, would prefer to say that the child develops an extremely strong tendency to associate what a linguist would describe as the past tense form with new verbs, in the absence of disconfirming evidence. The effect is the same, however stated; the learner inductively acquires the past tense marker and can use it expertly once a certain number of exemplars have been presented.

Of course, some learning could not be inductive, again on anyone's account: you can't induce the existence of D, E, and F on the basis of hearing A, B, and C, for instance, but rather must have the entire alphabet presented to you. This is sometimes referred to as brute force enumeration: simply all the exemplars of a given set are presented to the learner. In the same way, it might be possible to present the learner – at least an adult learner – not with a small sampling of the set of past-tense verbs to be learned, but rather with an explicit rule for producing those verbs. But it should be clear that the possibilities in language learning for brute force enumeration will be fairly limited, and I will say no more about it.

In any case, there seems to be a serious insufficiency with inductive learning as an explanation of the language acquisition process. Induction is notoriously *fallible*; the next raven we run into may be white, the next verb irregular. This is a problem, for first language acquisition theorists at least, because first language acquisition is standardly taken to be *in*fallible.[11] This simple but immense fact has, as we all know, led theorists to posit some sort of deductive learning mechanism, pre-eminently parameter-setting triggered by appropriate input. What sorts of input are appropriate, of course, is the big question (see, e.g., J. D. Fodor, 1998; Gibson and Wexler, 1994); but in any case it is assumed that the input feeds into a *parser*, which processes the information if it can, and revises the current grammar if it can't, in order to be able to handle the problematic input. Thus, for a theory that assumes some sort of rule-like, systematic representational system in its property-theory component, the transition theory will largely consist of a theory of grammatical parsing of input.

It will, of course, have occurred to the reader that "infallibility" is not the most apt term for characterizing SLA. And indeed, one might want to attribute the pretty much general failure, or seeming failure, to acquire nativelike L2 competence to the parser's inability to learn from its failures. The adult learner, it could be argued, has a representational system of the same general type as the native – not, *pace* Bley-Vroman (1990), a fundamentally different one – but a parser that can no longer make adjustments in the developing grammar to correct for parsing failures.

5.2 *The role of input*

An SLA transition theory will, of course, vary according to the property theory with which it is linked. Thus, depending on whether the property theory is a "classical" one of some sort – a UG/SLA theory, say, or a cognitive nativist theory – or a connectionist one, we will have different views of the role of input in acquisition.

5.2.1 Frequency

One question, simple enough on the face of it, is the relation between frequency of input and acquisition: how often does input of X need to be provided in order for X to be acquired? Actually, the question is badly put, since we don't receive input of X. What we want to know, rather, is this: in order to create representation R, how often does input that (in some sense that needs to be made clear by the property theory) "contains" R need to be presented to the learner?

On a connectionist view, it would seem that for any and every R, a good deal of relevant input would be necessary. A connectionist learning system learns by adjusting the weights of connections between nodes, and those adjustments, although not monotonic, are comparatively slow and gradual. For connectionists, as Ellis and Schmidt point out, one advantage of connectionist models is that they are "data-driven with prototypical representations emerging as a natural outcome of the learning process rather than being prespecified and innately given by the modellers as in more nativist cognitive accounts" (1998, p. 317). Put a bit differently, "What distinguishes between [connectionist and classical systems] is that, although both can learn, the former can't be programmed but have to be trained. As it turns out, that's a mixed blessing" (J. A. Fodor, 1998d, p. 85). Training takes time, and since the input is the trainer, that means a good deal of input is needed. Language learning seems to be a gradual process, so the gradual, input-based nature of connectionist models is often touted as a plus.

Of course, as Fodor says, this blessing is mixed. Learning a language certainly takes time, but that does not mean that learning any and every specific element of a language takes time. As Sterelny notes, "Lots of human learning is quick; there is a lot of one-shot learning from perception and language. Connectionist learning looks a good model for skill learning, but not for information gathering" (Sterelny, 1990, p. 193; cf. Schmidt, 1994). It is thus perhaps not surprising that, despite talk of representations, connectionists tend to speak in terms of skills.

Theories that, unlike connectionist theories, allow for the existence of rules nonetheless recognize the need for repeated input of R in some cases at least, but there is an important difference. In order to acquire, say, knowledge of plural formation in English, all theories agree that the learner needs to be presented with exemplars of regular nouns in the singular and the plural (putting aside for the moment the possibility of being presented with an explicit rule for plural formation; see "negative evidence" in section 5.2.2). The question is what happens next. A classical theory would see the input acting as the basis for inducing a rule, which would become the *de*ductive basis for determining the plural of nouns not yet presented in input. A connectionist theory, on the other hand, would see the relevant input as merely increasing the strength of association between input nouns and plural *-s*; an asymptotic increase, perhaps, but still only a statistical association, not a rule-based one.[12]

Where connectionist theories may be embarrassed by one-shot learning, UG/SLA theories have the opposite problem. Although UG theories can accept the need for perhaps fairly large doses of input in the formation of specific rules like English pluralization, the core of language learning presumably lies in parameter-setting. And where rule-formation is an inductive process, parameter-setting is supposed to be deductive. Hence the idea of triggering. Input for parameter-setting is not intended as evidence for a hypothesis, but rather as a stimulus that will reliably provoke the learning mechanism to fix one element of the grammar. Triggering is deductive not in the sense that the learner actually engages in anything like deductive reasoning, but rather in the sense that the chain from input to grammar-formation is infallible in the way that the chain from premises to conclusion is infallible in a deductive syllogism. Triggering is deductive in the way that imprinting is: the newly hatched duckling acts *as if* reasoning, "If it moves, it's Mom; that thing just moved; *ergo*, that thing is Mom."

This is all well and good if you're a duck, or if you're an ethologist studying ducks; there's a well-demonstrated relation between cause and effect that should be highly satisfactory to the both of you. The language acquisition theorist is not in such an enviable position; very little is even thought to be known about what specific stimuli in the input could act as the trigger for the setting of a specific parameter in a specific language. And in SLA, discussion of triggering, and of parsing in general, is close to non-existent. In any case, if there is triggering in SLA, one would expect – at least, in the absence of an extenuating explanation – fairly clear-cut results, in the form of a very steep learning curve following the triggering act. Indeed, one exemplar of whatever it is that is necessary to set a given parameter should suffice for that parameter to be correctly and permanently set.[13] We do not seem to have evidence showing such sudden effects in SLA, and indeed there is evidence (e.g., Kanno, 1999; O'Grady, 1999a) that L2 parameter setting may take years, even under seemingly ideal conditions.

5.2.2 Negative evidence and modified input

In first language acquisition, the child succeeds in acquiring native competence without benefit of negative evidence – explicit correction, or explicit metalinguistic information, such as about how to make the past tense. But of course it is widely believed that adults can benefit from negative evidence; so widely, indeed, that there is a multimillion-dollar publishing industry based on this belief. Still, that doesn't mean the belief is incorrect. If I tell you that in Japanese the past tense form of *asobu* is *asonda*, you may very well learn that fact, and even go on to infer that the past tense form of *yobu* is *yonda*. I rather doubt that anyone in SLA has ever believed that *no* negative evidence is *ever* usable, or that negative evidence can never accelerate the *speed* of acquisition. Nor do I imagine that anyone is claiming that negative evidence is always usable; no one, I trust, is arguing that learners will benefit from having ECP violations called to their attention. The real questions for an SLA theory are, is negative evidence ever *necessary* in SLA, and if so, when? One fairly concrete suggestion that has been made (White, 1987, 1989) is that when the L1 and L2 are in a superset/subset relation with respect to a given parameter, the learner will be unable to reset the parameter to the more restrictive subset value, in the absence of negative evidence (for some critical comments on the treatment of the Subset Principle in SLA, see Gregg, 1996a, 2001; White, 1989). In general, as always, the question is still open.

One's position on this question, and the fervor with which one defends it, will depend to some extent on the kind of property theory one supports. A UG/SLA theorist, for instance, should be perfectly comfortable with negative evidence being useful sometimes, so long as the evidence does not implicate UG. UG, after all, is posited as a solution to the problem of the poverty of the stimulus; but if the teacher or the textbook tells you all you need to know about forming the past tense of Japanese verbs, then the stimulus is not impoverished, and there's no puzzle about why you now know about Japanese past-tense formation. The point – or the claim, at least – about language competence is that it vastly transcends the kinds of knowledge that could conceivably be acquired through provision of this sort of evidence; not that 100 percent of language competence consists of such knowledge, only that it includes such knowledge to an important extent. Thus, while UG/SLA theorists can live with a role for negative evidence, that role must needs be a minor one at best.

Ironically, perhaps, a connectionist would seem to need to be more strongly committed to the non-efficacy of negative evidence than a UG/SLA theorist. Language acquisition, like all other learning, is for the connectionist a strengthening of associations, say between verb stems and past tense forms. The strengthening is accomplished by repeated input of the relevant forms, not by explicit metalinguistic reference to the forms. It's not clear (to me, anyway) how input of a sentence like, "The past tense is *-ed*," even repeated a hundred times a day, can be used by a "neural network" to strengthen the connection between verbs and their past tense forms. Thus in this sense it would seem that the connectionist SLA theorist must rely on positive evidence – everyday input – to an even greater extent than the UG/SLA theorist.[14]

A child L1 acquirer also seems to do just fine without any special modifications of the input; despite years of heroic effort, researchers have failed to show the necessity of "motherese," expansions, repetition, recasts, or other forms of input modification. But adult L2 learners *don't* do just fine in general, so perhaps modified input is necessary, or at least useful (Long, 1996). That is a theoretical possibility, of course, although the jury is still not in. But even if it turns out to be the case that input modification is essential for the acquisition of nativelike competence, this would not be much more of a contribution to the framing of a transition theory than would the discovery that motivation is essential, and for the same reason. Modifying the input is basically a way of making the input cleaner, more easily handled by the learning mechanism, whatever that is; in the extreme case, input modification makes the input input. By the same token, sitting attentively in the language classroom each day, rather than hanging out in the quad, makes the input in that classroom input, rather than noises off. But neither a theory of motivation nor a theory of input modification will tell us how the learning mechanism operates on the input, howbeit modified, to produce a bit of grammar.

6 Evaluating SLA Theories

It's not really clear that we yet have anything worth calling a theory of SLA, in which case it may seem premature to discuss evaluating them. Still, there are at least proposals on the table, if not perhaps very detailed proposals, and we can at least consider what sorts of problems these proposals, or future theories, may face.

6.1 Red herrings

First, though, we need to dismiss a couple of non-problems that get raised all too often in the literature.

6.1.1 Plausibility

Proposals are often evaluated – prematurely, I would say – in terms of their plausibility. Connectionists, for instance, are fond of pointing to the putative similarity of their somewhat tendentiously named "neural network" models to the way the brain really works.[15] More often, plausibility arguments take the form of an attack on a rival proposal for being *im*plausible on one ground or another. These attacks usually are simply examples of what Dawkins (1986) calls the Argument from Personal Incredulity; rather than providing empirical or theoretical evidence contradicting the proposal, one simply appeals to one's sense of what is and is not likely. In SLA, proposals based on theories of Universal Grammar have been especially subject to such attacks, especially perhaps from adherents of what Stoljar and Gold (1998, p. 111) call the "Biological Neuroscience Thesis," the thesis that mental science is *biological* neuroscience, "where 'biological neuroscience' is intended to include only those sciences traditionally regarded as part of neurobiology, roughly: neuroanatomy, neurophysiology, and neurochemistry." Thus Jacobs and Schumann (1992), for instance, along with numerous others, dismiss the constructs of generative linguistic theory as no better than metaphors. Similarly, Ellis tells us that "Innate specification of synaptic connectivity in the cortex is unlikely. On these grounds, linguistic representational nativism seems untenable" (Ellis, 1999, p. 25).

Arguments such as these simply have no force. Implausibility is one of the hallmarks of the natural world, from gravity and quarks to echolocation and metamorphosis, and the mind is one of the most implausible things around. Not, mind you, that we should adopt Tertullian's motto (*Credo quia absurdum est*, I believe because it is absurd), but the appeal to unlikeliness is simply no argument at all. Whether, say, UG exists or not is an open question, of course. But UG is posited not because it's plausible, but because it can explain certain phenomena – phenomena, moreover, about the existence of which there is little dispute. And in science, one normally rejects an explanation only when one has a superior explanation to replace it. At the moment, no adherent of the Biological Neuroscience Thesis has anything like an explanation to rival those offered by linguistic theories that posit something like UG.[16]

Now, scientists do, of course, reject certain theories out of hand, on grounds that may seem like the Argument from Personal Incredulity. No scientist accepts so-called "creation science," for instance, or time travel, or ESP. But these "theories" are not rejected because of their implausibility – a concept for which there is no useful standard against which to measure theories – but because they lack any empirical confirming evidence, while simultaneously contradicting well-confirmed theories that explain a great many phenomena. To accept these truly implausible theories would entail abandoning those well-confirmed theories and the explanations they provide, in exchange for nothing, a price no rational person should be willing to pay.

This is definitely *not* the case with UG, or more generally with cognitive theories that posit mental organization at a level higher than, and not directly reducible to, the neurological. Accepting a UG theory does, of course, require abandoning certain other possible theories of the mind. But unlike creationism or ESP, UG theory does not contradict any well-confirmed theory of the mind, and hence does not require us to abandon well-supported explanations of mental phenomena. The unhappy fact is that we don't *have* much in the way of well-supported explanations of mental phenomena; in fact, linguistics is about the most advanced of the cognitive sciences. At this point at least, McLaughlin and Warfield claim, "there is nothing known about the human brain that gives any reason whatsoever to doubt that it contains a classical cognitive architecture" (1994, p. 381; cf. Smolensky, 1999). This does not mean for a minute that UG theories are correct, of course; it merely means they are not to be rejected on such a flimsy ground as implausibility.[17]

6.1.2 Simplicity

Akin to plausibility is the red herring of simplicity. As we saw above, realists are willing to appeal to non-empirical factors, including simplicity, to adjudicate between rival theories, while empiricists are not. Putting aside the unsettled question whether such an appeal can ever be justified, it certainly is the case that it cannot *always* be justified. Occam's Razor, for instance (entities are not to be multiplied beyond necessity), can only be usefully invoked if it's clear whether a given theoretical construct is truly otiose, where one can compare a theory with the construct and the same theory without it. Thus, in perhaps the first explicit invocation of Occam's Razor in SLA, Gregg (1984) argued that the Affective Filter of Krashen (1981) was otiose in just this way. Such easy targets are rare, however, and it is normally quite difficult to decide, even intuitively, which of two theories is the simpler; all the more difficult when it's not even clear whether one has two theories to compare. As Chomsky said in relation to linguistic theory and first language acquisition, where far greater progress has been made than in SLA, "The issue of relative 'simplicity,' even if this notion can be given some content relevant to choice among theories, can hardly be sensibly raised in connection with theories so meager in confirming evidence and explanatory force as those that have been proposed to account for learning and behavior" (1980b, pp. 288–9).

Appeals to simplicity in SLA tend to be made, ironically enough, against realist positions such as UG/SLA theories. As with plausibility, the appeals are usually general metatheoretical claims, rather than specific comparisons between two theories, say one with binding principles and one without. Thus Ellis, for instance, notes that connectionists are fond of appealing to Morgan's Canon, the principle introduced by the evolutionary biologist Lloyd Morgan, which holds that "in no case may we interpret an action as the outcome of a higher psychical faculty if it can be interpreted as the outcome of one which stands lower in the psychological scale" (Ellis, 1999, p. 28). The canon here seems to be being used as a form of Occam's Razor (although that may not be the appropriate interpretation; cf. Sober, 1998), but in fact the argument doesn't go through. As Morgan himself noted, "the canon by no means excludes the interpretation of a particular activity in terms of the higher processes, if we already have independent evidence of the occurrence of these higher processes in the animal under observation" (Morgan, 1903, p. 59; cited in Sober, 1998, p. 240, fn. 1). But we have such independent evidence, and in abundance, for relevant higher processes in humans; the systematicity argument is based on just such evidence. Morgan's Canon will keep us from attributing beliefs to bacteria and syntax to snakes, but it's of no use in assessing claims about the nature of language and language acquisition.

6.2 *Explanatory problems*

We are still left with plenty of real problems for an SLA theory to overcome, a few of which I'll discuss briefly.

6.2.1 Replacing UG

If we exclude UG/SLA theories for the moment, probably the most fundamental problem facing SLA property theories is that they don't exist. For better or worse, SLA theories of L2 knowledge are theories of UG, using the term loosely to include a number of competing variants, the differences among which we can ignore here. This is regrettable, for although interesting and valuable SLA research can be and is being carried out without an underlying well-articulated property theory, ultimately the question of that underlying theory needs to be addressed. The UG/SLA people – or some subset of them, at least – may turn out to be correct, but we don't know that yet, and it's always beneficial to have competition. At the moment there simply is no well-articulated rival theory of L2 competence against which to measure UG theories.

Now, it might seem that connectionism offers just such a rival, but appearances can be misleading. What one sees, by and large, are connectionist simulations of language acquisition, whose results are (perhaps overoptimistically) interpreted by connectionists as obviating the need for "classical" entities like syntactic rules or principles. But even on the rosiest interpretation of connectionist work on SLA – even, that is, if we were to concede that the simulations are truly successful in "acquiring" the knowledge in question, *and* even if we were to make the much greater concession that the simulations mirror human language learning processes – we still have no explanation of what it is that the learner has acquired. As Sterelny, anything but a foe of connectionism, says, "[T]here is no argument to connectionism as a global theory of the mind from its demonstrated success in dealing with some major portion of it" (Sterelny, 1990, p. 192). Classical theories, including UG theories of language, can explain – whether correctly or incorrectly is another question – such robust phenomena as the systematicity and productivity of language, because classical theories can appeal to rules and principles with causal powers. Connectionist theories are at a disadvantage when trying to explain systematicity and productivity precisely because they reject the concept of non-artifactual rules, without replacing them with anything that can do the job. As McLaughlin and Warfield argue, "connectionists have yet to articulate an alternative to the classical conception of thought, and we think the prospects for its offering an adequate alternative are dim" (1994, p. 374; cf. Gold and Stoljar, 1999; Jackendoff, 1999). And *mutatis mutandis* for language.

6.2.2 "Access to UG"

Among those property theories based on some concept of UG, there is the question of whether or not UG plays an identical role in adult L2 acquisition and in child language acquisition. This question has often been characterized as the question of "access to UG," an unfortunate metaphor that confuses the issue instead of illuminating it.[18] Basically, the question is whether an adult L2 grammar is constrained in exactly the same way as an L1 grammar is constrained by the various principles and parameters of UG. If it is, we would expect, regardless of whatever other "imperfections" or "gaps" there might be in the L2 knowledge representations, to find nothing that violates UG; there should be no "wild" or "rogue" grammars.

The evidence generally seems to indicate an absence of rogue L2 grammars (but see, e.g., Klein, 1995; Thomas, 1991; for detailed discussion, see White, forthcoming, ch. 2). The question, though, is why one would have expected otherwise. That is, under what conception of UG could the "access" question arise in the first place? With earlier characterizations of the access debate, at least – White's (1989) UG-is-dead/UG-is-alive, Gregg's (1996a) theists vs. deists – the assumption seems to have been that UG is a machine to make grammars. Full access would mean that the machine is still in perfect working order; partial access would mean that the machine was in some way impaired and that the final product consequently lacked some parts; and zero access would entail having to build the L2 grammar with different tools. In any case, UG was implicitly being conceived of as separate from any particular grammar.

Now, this is not an incoherent stance to take, but it is inconsistent with most current understandings of UG, where a given grammar (say, the English grammar I carry around in my head) just *is* the set of UG principles, instantiated in a specific way. Without the principles there's no instantiation, which is to say that I am accessing UG every time I open my mouth, and that UG can't die until I do. Of course, on some views of UG there might be UG principles and parameters – and there certainly will be parameter values – that are not relevant to the L1 (subjacency was once one such candidate; cf. Bley-Vroman, Felix, and Ioup, 1988). One could then argue that the learner has no access to precisely those elements of the L2, while having access to those elements relevant to the L1. This would be a "partial access" theory, I suppose; but it seems hard to distinguish it from "zero access" proposals like Bley-Vroman's Fundamental Difference Hypothesis (1990), at least as far as their claims about the nature of L2 representations go.

In fact, it would seem that the various proposals about access to UG need to be formulated as transition theories if they are to be distinguished one from the other and compared. UG, after all, is a set of constraints, and having full access to UG in effect means being fully constrained by UG (see White, this volume). Thus Epstein et al., for instance, define "access" as follows: "We mean by 'X is accessible' only that 'X constrains the learner's hypothesis space'" (1996, fn. 5). But as several of their commentators point out (Gregg, 1996b; Sprouse, 1996; White, 1996; *inter alia*), this definition leaves open all sorts of unwelcome possibilities. Given, for instance, that every parameter setting there is lies within the hypothesis space defined by UG, "access" is fully consistent with the L2 learner (or the L1 learner, for that matter) setting every single parameter to the wrong value. It would be cold comfort to both learner and theorist if learners merely avoided rogue grammars while failing to process input at all successfully. More to the point, however deviant their L2 grammars may be, learners don't in fact do anything so irrational as ignore input, and that fact cannot be attributed simply to "full access to UG" if that merely means full obedience to grammatical constraints.

I by no means wish to disparage the "no rogue grammars" argument. If, as seems to be the case, L2 learners do not produce rogue grammars, that is a highly significant fact, calling for an explanation. And indeed, the significance of this fact is often devalued by appeals (e.g., Epstein et al., 1996) to the absence of truly weird grammars among L2 learners: grammars violating structure-dependence, for instance. More to the point would be cases where UG bans a grammar, but common sense and the input don't. For instance, Binding Theory allows certain variations in the scope of anaphors: the Japanese equivalent of *John thinks that Bill should introduce himself to Mary* is ambiguous as to whether John or Bill is to be introduced to Mary. To my knowledge, while there are languages like Japanese that permit reference to either the matrix subject or the embedded subject, and languages like English that permit only reference to the embedded subject, there is no language that permits only reference to the matrix subject.[19] If adult learners still have "access to UG," then English natives learning Japanese should never create such a rogue anaphor system, even if every single instance of anaphora in the input happened to make unambiguous reference to the matrix subject. On the other hand, if learners persist in restricting Japanese anaphora to the embedded subject, they are remaining comfortably within the hypothesis space of UG; they simply aren't getting the appropriate UG-constrained message from the input.

The access question, in short, needs to be illuminated by the transition theory if it is to be settled. We need more than an enumeration of the elements of UG which are no longer "accessible" to an adult learner, if there are such elements. To explain those deficits, we also need an account of how input should affect learning if those deficits were not there.

6.2.3 Variation across final states

Sooner or later, any SLA theory must deal with the fact that the final state, however characterized by the property theory, varies across learners, and differs from the final state achieved by a native speaker. Different theories will have different answers to these two problems, and may have different difficulties in making their answers stick.

To start with, why don't L2 learners acquire the L2 to the same degree as natives? The obvious answer might seem to be age: there definitely do seem to be robust negative correlations between age of onset of acquisition and final proficiency (Long, 1990b). And there is strongly suggestive evidence, at the least, for a critical period or periods for language acquisition, although there still is a good deal of disagreement among researchers on this question (see the papers in Birdsong, 1999; Hyltenstam and Abrahamsson, this volume). But even on the most favorable interpretation of the evidence, age cannot be the whole story, or we would expect absolutely no L1 influence on L2 development, which is clearly not the case. But again, to claim that Maria's English competence surpasses Keiko's because, say, both English and Spanish are SVO while Japanese is SOV is to beg an important question: why should these cross-linguistic difficulties be insuperable? Granting that the word-order difference might benefit Maria, why should Keiko fail to catch up? (Assuming she does fail, of course, and putting aside the question of Mariko, who can already run rings around Maria.) Unfortunately, it's hard to answer these questions yet, in part because most of the research related to age effects has not been conducted within a well-articulated property theory, but rather has contented itself with an unsatisfactory concept of "proficiency," which, while it can be "operationalized" with elegant accuracy (TOEFL over 600, say), lacks theoretically interesting content.

7 Conclusion

There are, of course, numerous other problems confronting the SLA theorist, but it is in fact an encouraging sign that we can specify them as clearly as we now can. The last two decades of SLA research have seen not only a huge increase in the database, but also a much higher degree of conceptual precision and theoretical sophistication. And this is not only in the property theory, where generative grammatical theories continue to change and develop, but also in the transition theory. One should not be misled by the common empiricist origins of SLA behaviorism and SLA connectionism into overlooking how much better articulated and detailed the latter is, which makes it much easier to locate the problems. And connectionists have gone well beyond the dogmatic handwaving of the behaviorists to actually offer simulations of acquisition (it would be nice to see some UG/SLA computer models). It is hardly surprising, though, that theoretical and methodological problems still abound; SLA is a newly emerging scientific field, and problems come with the territory.

ACKNOWLEDGMENTS

I am happy to acknowledge the very helpful comments of William O'Grady and Lydia White, although I imagine they'll neither of them be that satisfied with how I've made use of them. Thanks also to Mike Long and to two anonymous reviewers.

NOTES

1 "Someone committed to naturalistic inquiry can consistently believe that we learn more of human interest about how people think and feel and act by studying history or reading novels than from all of naturalistic inquiry. Outside of narrow domains, naturalistic inquiry has proven shallow or hopeless" (Chomsky, 1995, p. 28).

2 One needs to distinguish between the scientific study of SLA on the one hand and the academic field of "applied linguistics" on the other. The latter, when it isn't simply the respectable field of foreign language education cloaked in a meretricious nomenclature, does not seem to have an object of study, a research program, or a goal. Indeed, there are frequent earnest discussions as to what in fact applied linguistics is or should be (*Issues in Applied Linguistics*, 1 (2); Kaplan, 1980). The amorphousness of the field of applied linguistics – the seeming lack of any of the theoretical and methodological constraints that one expects in an empirical discipline – has allowed many of its adepts to wander in what Shelley called the "intense inane," issuing pronouncements that range from vacuous to incoherent to downright delusional. Pasteur had it right a hundred years ago: "[T]here does not exist a category of science to which one can give the name applied science. There are science and the applications of science, bound together as the fruit to the tree which bears it" (cited by Leiden, 1999, p. 1215). It perhaps goes without saying that this chapter – indeed, this handbook – is not about applied linguistics.

3 Relativistic research, as Long suggests, is a contradiction in terms. This does not mean, however, that there are no SLA researchers who have espoused relativism in some form or other. Such espousals as have appeared in the literature, though (e.g., Block, 1996; Lantolf, 1996; Schumann, 1983; van Lier, 1994), are simply risible at best. See

Gregg et al. (1997), Gregg (2000), and Long (1998) for discussion.

4 The existence itself of UG, of course, cannot be rejected on the grounds of non-observability, nor does any sensible empiricist make such a rejection. The empiricist argument is not that what we cannot observe does not exist, but merely that we are not warranted in believing in the existence of what we cannot observe; a huge difference.

5 The so-called New Look perceptual psychology of the 1950s and 1960s seemed to show a strong influence of belief on perception. (J. A. Fodor's modularity thesis, 1983, is in part an extended refutation of, and indeed was a major factor in the rejection of, New Look psychology.) Kuhn (who, after all, was trained as a scientist) drew on these results as major empirical support for his conclusions about the theory-ladenness of observation. The irony of relying on theory-guided empirical research to justify a position which, if correct, would undermine any reason to accept the research was apparently lost on Kuhn.

6 Note that realists and relativists both oppose empiricists on the issue of the theory–observation distinction, although of course they draw radically different conclusions from this common opposition. As Kukla puts it, "realists and relativists agree that theoretical and observational hypotheses, if they can be distinguished at all, are in the same epistemic boat. They just differ as to the nature of the boat" (1998, p. 112). For realists, the illegitimacy of the distinction allows for the reality of (some) unobservable entities; for relativists, it leads to the subsumption of all observational results under the merely theoretical.

7 The idea that one must operationalize one's definitions is a relic of pre-war positivism that has survived only in psychology, to the bemusement of philosophers of science (e.g., Greenwood, 1991; Hempel 1966; Hull, 1974, 1988; Klee, 1997).

8 Eubank and Gregg (1995, p. 54) make this very point with reference to language acquisition: "Although we think the increased interest in neurolinguistics shown by SLA researchers is a promising sign of increasing sophistication in our field, the fact remains that little progress can be expected in acquisition theory if researchers fail to take linguistics seriously." In response, Schumann (1995, p. 61) insists that "A neurobiological perspective on language is responsible to language only and not to any particular linguistic characterization of language." An exactly parallel argument, of course, can be made by the astrologist: "An astrological perspective on the heavenly bodies is responsible to the heavenly bodies only and not to any particular astronomical characterization of heavenly bodies."

9 Cummins (1991, p. 114), however, argues that "adopting a connectionist architecture does not force one to abandon the 'classical' idea that cognition is to be understood as the computational manipulation of semantically structured representations." And on the other hand, Ramsey (1997) questions whether one need consider connection weights to be representations at all, structured or not: "there isn't anything about connectionism that demands we think the weights function as representations of stored information" (p. 49). Suffice it to say

that in general, connectionists themselves do consider their representations to reside in the varying connection weights, and that they do take these representations to differ from "classical" representations precisely in being unstructured.

10 "Clark [1993, p. 225] says that we should 'bracket' the problem of systematicity. 'Bracket' is a technical term in philosophy which means *try not to think about*" (J. A. Fodor, 1998a, p. 99).

11 Well, not infallible, of course, or why does language change over time? Still, the idealization to an unerring language acquisition device seems eminently reasonable, given the essentially uniform final states achieved by all unimpaired members of a roughly identifiable speech community.

12 Ellis and Schmidt (1998) tested their model of the acquisition of plural morphology on a nonce noun, and it did quite well in producing the regular plural of their artificial language. (Actually, it wasn't a true nonce word, but rather one that had only been presented in the singular.) Significantly, Ellis and Schmidt did not bother to test their human subjects on the same nonce word, merely suggesting that their test of the model was analogous to doing so. But of course on a rule-based account of plural-learning, one would predict that human subjects who had acquired the rule would not merely do quite well, but would score at or very near 100 percent, and would not benefit from further input.

13 This is, in fact, a problem for a triggering account of language acquisition. In the case of imprinting, there is a definite, albeit statistically minute, chance that the hatchling, say, will first see something other than its mother, and will form an irreparable bond with that "incorrect" stimulus object. What if the first relevant input for a given parameter happens by misadventure to be inappropriate for the target language? If parameter-setting were as deterministic as imprinting, the unfortunate child would presumably be stuck with a deviant grammar.

14 One might be tempted to treat the weakening of a connection due to the absence of strengthening stimuli as a form of indirect negative evidence; this temptation should be resisted. The idea of indirect negative evidence, as proposed by Chomsky (1984), is that "if certain structures or rules fail to be exemplified in relatively simple expressions, *where they would be expected to be found* [emphasis added], then a (possibly marked) option is selected excluding them in the grammar" (p. 9); in effect, a sort of unconscious deduction by *modus tollens*. But connection-weakening has nothing to do with rules, expected or otherwise; the failure of the learner to make a connection between singular nouns and [ba], based on the total absence of input of [ba] immediately after a singular noun in the input, may well serve to keep that learner from acquiring [ba] as the English plural marker; it won't lead the learner to acquire /z/.

15 The similarity of neural network models to neural networks lies more in the eye of the beholder, or rather the modeler, than in reality, connectionist protestations notwithstanding; for discussion see, e.g., J. A. Fodor (1998e); Rey (1997).

16 "To advocate the biological neuroscience thesis . . . is to claim

that eventually we will have explanations of mental phenomena that are couched in the concepts of neurobiology. This view is extremely interesting, but one would need considerable evidence to accept it" (Stoljar and Gold, 1998, p. 111; see Gold and Stoljar, 1999, for detailed discussion).

17 It is often argued that UG is inconsistent with evolutionary theory (see, e.g., Deacon, 1997; Lieberman, 1984, 1991). If this were true, it would be a strong argument against UG theory, given that evolutionary theory is well confirmed, to say the least. But in fact there is no reason to believe that there is any contradiction; see, for example, J. A. Fodor, 1998f, 1998g; Pinker and Bloom, 1990; for arguments from very different perspectives against the argument from evolutionary implausibility.

18 There is a sizeable literature on the "access" question. See, for instance, the papers in Eubank (1991); Epstein, Flynn, and Martohardjono (1996) and the commentaries thereon. There is also a set of papers given at a colloquium on the issue at the 1998 SLRF (by Bley-Vroman, Carroll, Gregg, Meisel, Schwartz, and White), available on the internet: <www.lll.hawaii.edu/nflrc/NetWorks/NW9>.

19 Actually, if Iatridou (1986) is correct, there is at least one such language (Greek). Still, the point remains that there could be UG-banned but plausible – inductively acquirable – IL grammars, grammars that should not, on the "access" account, be attested.

REFERENCES

Aizawa, K. 1997: Explaining systematicity. *Mind and Language*, 12, 115–36.

Ariew, A. 1996: Innateness and canalization. *Philosophy of Science (Proceedings 63)*, S19–S27.

Ariew, A. 1999: Innateness is canalization: in defense of a developmental account of innateness. In V. G. Hardcastle (ed.), *Where Biology Meets Psychology: Philosophical Essays*. Cambridge, MA: MIT Press, 117–38.

Baron-Cohen, S. 1995: *Mindblindness: An Essay on Autism and Theory of Mind*. Cambridge, MA: MIT Press.

Bates, E. and MacWhinney, B. 1988: What is functionalism? *Papers and Reports on Child Language Development*, 27, 137–52.

Beretta, A. 1991: Theory construction in SLA: complementarity and opposition. *Studies in Second Language Acquisition*, 13, 493–511.

Beretta, A. and Crookes, G. 1993: Cognitive and social determinants of discovery in SLA. *Applied Linguistics*, 14, 250–75.

Birdsong, D. (ed.) 1999: *Second Language Acquisition and the Critical Period Hypothesis*. Mahwah, NJ: Lawrence Erlbaum Associates.

Bley-Vroman, R. 1983: The comparative fallacy in interlanguage studies: the case of systematicity. *Language Learning*, 33, 1–17.

Bley-Vroman, R. 1990: The logical problem of foreign language learning. *Linguistic Analysis*, 20, 3–49.

Bley-Vroman, R., Felix, S., and Ioup, G. 1988: The accessibility of Universal Grammar in adult language learning. *Second Language Research*, 4, 1–32.

Block, D. 1996: Not so fast: some thoughts on theory culling, relativism, accepted findings and the heart and soul of SLA. *Applied Linguistics*, 17, 63–83.

Bogen, J. and Woodward, J. 1988: Saving the phenomena. *Philosophical Review*, 97, 303–52.

Boyd, R. 1989: What realism implies and what it does not. *Dialectica*, 43, 5–29.

Broeder, P. and Plunkett, K. 1994: Connectionism and second language acquisition. In N. C. Ellis (ed.), *Implicit and Explicit Learning of Languages*. New York: Academic Press, 421–53.

Brueckner, A. 1998: Conceptual relativism. *Pacific Philosophical Quarterly*, 79, 295–301.

Carroll, S. E. 1995: The hidden dangers of computer modelling: remarks on Sokolik and Smith's connectionist learning model of French gender. *Second Language Research*, 11, 193–205.

Carroll, S. E. 1999: Putting "input" in its proper place. *Second Language Research*, 15, 337–88.

Carroll, S. E. forthcoming: *Input and Evidence: The Raw Material of Second Language Acquisition*. Amsterdam: John Benjamins.

Chomsky, N. 1959: Review of B. F. Skinner, *Verbal Behavior*. *Language*, 35, 26–58. Reprinted in J. Fodor and J. Katz (eds) 1964: *The Structure of Language: Readings in the Philosophy of Language*. Englewood Cliffs, NJ: Prentice-Hall, 547–8.

Chomsky, N. 1980a: *Rules and Representations*. New York: Columbia University Press.

Chomsky, N. 1980b: Some empirical assumptions in modern philosophy of language. In H. Morick (ed.), *Challenges to Empiricism*. Belmont, CA: Wadsworth, 287–318.

Chomsky, N. 1984: *Lectures on Government and Binding*. 3rd revised edition. Dordrecht: Foris.

Chomsky, N. 1986: *Knowledge of Language*. New York: Praeger.

Chomsky, N. 1995: Language and nature. *Mind*, 104, 1–61.

Clark, A. 1993: *Associative Engines*. Cambridge, MA: MIT Press.

Coltheart, M. and Langdon, R. 1998: Autism, modularity and levels of explanation in cognitive science. *Mind and Language*, 13, 138–52.

Cummins, R. 1983: *The Nature of Psychological Explanation*. Cambridge, MA: MIT Press.

Cummins, R. 1991: The role of representation in connectionist explanations of cognitive capacities. In W. Ramsey, S. P. Stich, and D. E. Rumelhart (eds), *Philosophy and Connectionist Theory*. Hillsdale, NJ: Lawrence Erlbaum Associates, 91–114.

Cummins, R. 1996: Systematicity. *Journal of Philosophy*, 93, 591–614.

Dawkins, R. 1986: *The Blind Watchmaker*. New York: Norton.

Deacon, T. 1997: *The Symbolic Species: The Co-Evolution of Language and the Brain*. New York: Norton.

Ellis, N. C. 1999: Cognitive approaches to SLA. *Annual Review of Applied Linguistics*, 19, 22–42.

Ellis, N. C. and Schmidt, R. 1998: Rules or associations in the acquisition of morphology? The frequency by regularity interaction in human and PDP learning of morphosyntax. *Language and Cognitive Processes*, 13, 307–36.

Epstein, S. D., Flynn, S., and Martohardjono, G. 1996: Second language acquisition: theoretical and experimental issues in contemporary research. *Behavioral and Brain Sciences*, 19, 677–758.

Eubank, L. (ed.) 1991: *Point Counterpoint: Universal Grammar in the Second Language*. Amsterdam: John Benjamins.

Eubank, L. and Gregg, K. R. 1995: "*Et in amygdala ego*"? UG, (S)LA, and

neurolinguistics. *Studies in Second Language Acquisition*, 17, 35–57.
Eubank, L. and Gregg, K. R. 1999: Critical periods and (second) language acquisition: *divide et impera*. In D. Birdsong (ed.), *Second Language Acquisition and the Critical Period Hypothesis*. Mahwah, NJ: Lawrence Erlbaum Associates, 65–99.
Firth, A. and Wagner, J. 1997: On discourse, communication, and (some) fundamental concepts in SLA research. *Modern Language Journal*, 81, 286–300.
Fodor, J. A. 1983: *The Modularity of Mind*. Cambridge, MA: MIT Press.
Fodor, J. A. 1998a: *Concepts: Where Cognitive Science Went Wrong*. Oxford: Clarendon Press.
Fodor, J. A. 1998b: *In Critical Condition: Polemical Essays on Cognitive Science and the Philosophy of Mind*. Cambridge, MA: MIT Press.
Fodor, J. A. 1998c: Connectionism and the problem of systematicity (continued): why Smolensky's solution *still* doesn't work. In J. A. Fodor, *In Critical Condition: Polemical Essays on Cognitive Science and the Philosophy of Mind*. Cambridge, MA: MIT Press, 113–25.
Fodor, J. A. 1998d: Review of Paul Churchland's *The Engine of Reason, the Seat of the Soul*. In J. A. Fodor, *In Critical Condition: Polemical Essays on Cognitive Science and the Philosophy of Mind*. Cambridge, MA: MIT Press, 83–9.
Fodor, J. A. 1998e: Review of Jeff Elman et al., *Rethinking Innateness*. In J. A. Fodor, *In Critical Condition: Polemical Essays on Cognitive Science and the Philosophy of Mind*. Cambridge, MA: MIT Press, 143–51.
Fodor, J. A. 1998f: Review of Richard Dawkins's *Climbing Mount Improbable*. In J. A. Fodor, *In Critical Condition: Polemical Essays on Cognitive Science and the Philosophy of Mind*. Cambridge, MA: MIT Press, 163–9.
Fodor, J. A. 1998g: Review of Steven Pinker's *How the Mind Works* and Henry Plotkin's *Evolution in Mind*. In J. A. Fodor, *In Critical Condition: Polemical Essays on Cognitive Science and the Philosophy of Mind*. Cambridge, MA: MIT Press, 203–14.
Fodor, J. A. and McLaughlin, B. 1998: Connectionism and the problem of systematicity: why Smolensky's solution doesn't work. In J. A. Fodor, *In Critical Condition: Polemical Essays on Cognitive Science and the Philosophy of Mind*. Cambridge, MA: MIT Press, 91–111.
Fodor, J. A. and Pylyshyn, Z. W. 1988: Connectionism and cognitive architecture. *Cognition*, 28, 3–71.
Fodor, J. D. 1998: Unambiguous triggers. *Linguistic Inquiry*, 29, 1–36.
Gass, S. 1998: Apples and oranges: or, why apples are not orange and don't need to be. *Modern Language Journal*, 82, 83–90.
Gibson, E. and Wexler, K. 1994: Triggers. *Linguistic Inquiry*, 25, 407–54.
Gold, I. and Stoljar, D. 1999: A neuron doctrine in the philosophy of neuroscience. *Behavioral and Brain Sciences*, 22, 809–69.
Greenwood, J. D. 1991: *Relations and Representations: An Introduction to the Philosophy of Social Psychological Science*. London: Routledge.
Gregg, K. R. 1984: Krashen's Monitor and Occam's Razor. *Applied Linguistics*, 5, 79–100.
Gregg, K. R. 1989: Second language acquisition theory: the case for a generative perspective. In S. M. Gass and J. Schachter (eds), *Linguistic Perspectives on Second Language Acquisition*. Cambridge: Cambridge University Press, 15–40.
Gregg, K. R. 1990: The Variable Competence Model of second language acquisition, and why it isn't. *Applied Linguistics*, 11, 364–83.

Gregg, K. R. 1993: Taking explanation seriously; or, let a couple of flowers bloom. *Applied Linguistics*, 14, 276–94.

Gregg, K. R. 1996a: The logical and developmental problems of second language acquisition. In W. C. Ritchie and T. K. Bhatia (eds), *Handbook of Second Language Acquisition*. San Diego: Academic Press, 49–81.

Gregg, K. R. 1996b: UG and SLA: the access question, and how to beg it. *Behavioral and Brain Sciences*, 19, 726–7.

Gregg, K. R. 2000: A theory for every occasion: postmodernism and SLA. *Second Language Research*, 16, 383–99.

Gregg, K. R. 2001: Learnability and second language acquisition theory. In P. Robinson (ed.), *Cognition and Second Language Instruction*. Cambridge: Cambridge University Press, 152–82.

Gregg, K. R., Long, M. H., Jordan, G., and Beretta, A. 1997: Rationality and its discontents in SLA. *Applied Linguistics*, 18, 538–58.

Hacking, I. 1983: *Representing and Intervening*. Cambridge: Cambridge University Press.

Hanson, N. R. 1958: *Patterns of Discovery*. Cambridge: Cambridge University Press.

Hawkins, R. and Chan, C. Y.-H. 1997: The partial availability of Universal Grammar in second language acquisition: the "failed functional features hypothesis." *Second Language Research*, 13, 187–226.

Hempel, C. G. 1966: *Philosophy of Natural Science*. Englewood Cliffs, NJ: Prentice-Hall.

Hull, D. L. 1974: *Philosophy of Biological Science*. Englewood Cliffs, NJ: Prentice-Hall.

Hull, D. L. 1988: *Science as a Process*. Chicago: University of Chicago Press.

Iatridou, S. 1986: An anaphor not bound in its governing category. *Linguistic Inquiry*, 17, 766–72.

Jackendoff, R. 1999: Parallel constraint-based generative theories of language. *Trends in Cognitive Sciences*, 3, 393–400.

Jacobs, B. and Schumann, J. H. 1992: Language acquisition and the neurosciences: towards a more integrative perspective. *Applied Linguistics*, 13, 282–301.

Kanno, K. 1999: Acquisition of verb gapping in Japanese by Mandarin and English speakers. In K. Kanno (ed.), *The Acquisition of Japanese as a Second Language*. Philadelphia: John Benjamins, 159–73.

Kaplan, R. (ed.) 1980: *On the Scope of Applied Linguistics*. Rowley, MA: Newbury House.

Karmiloff-Smith, A. 1992: *Beyond Modularity*. Cambridge, MA: MIT Press.

Kasper, G. 1997: "A" stands for acquisition. *Modern Language Journal*, 81, 307–12.

Klee, R. 1997: *Introduction to the Philosophy of Science: Cutting Nature at its Seams*. Oxford: Oxford University Press.

Klein, E. 1995: Evidence for a "wild" L2 grammar: when PPs rear their empty heads. *Applied Linguistics*, 16, 87–117.

Krashen, S. D. 1981: *Principles and Practice in Second Language Acquisition*. Oxford: Pergamon Press.

Kuhn, T. S. 1970: *The Structure of Scientific Revolutions*. 2nd edition, enlarged. Chicago: University of Chicago Press.

Kukla, A. 1998: *Studies in Scientific Realism*. New York: Oxford University Press.

Lantolf, J. P. 1996: SLA theory building: "Letting all the flowers bloom!" *Language Learning*, 46, 713–49.

Larsen-Freeman, D. and Long, M. H. 1991: *An Introduction to Second Language Acquisition Research*. London: Longman.

Latour, B. 1987: *Science in Action*. Cambridge, MA: Harvard University Press.

Latour, B. and Woolgar, S. 1986: *Laboratory Life: The Construction of Scientific Facts*. 2nd edition. Princeton, NJ: Princeton University Press.

Laudan, L. 1990: *Science and Relativism*. Chicago: University of Chicago Press.

Laudan, L. 1996: *Beyond Positivism and Relativism*. Boulder, CO: Westview Press.

Leiden, J. M. 1999: Gene therapy enters adolescence. *Science*, 285, 1215–16.

Leplin, J. 1997: *A Novel Defense of Realism*. New York: Oxford University Press.

Lieberman, P. 1984: *The Biology and Evolution of Language*. Cambridge, MA: Harvard University Press.

Lieberman, P. 1991: *Uniquely Human: The Evolution of Speech, Thought, and Selfless Behavior*. Cambridge, MA: Harvard University Press.

Long, M. H. 1990a: The least a second language acquisition theory needs to explain. *TESOL Quarterly*, 24, 649–66.

Long, M. H. 1990b: Maturational constraints on language development. *Studies in Second Language Acquisition*, 12, 251–85.

Long, M. H. 1993: Assessment strategies for second language acquisition theories. *Applied Linguistics*, 14, 225–49.

Long, M. H. 1996: The role of the linguistic environment in second language acquisition. In W. C. Ritchie and T. K. Bhatia (eds), *Handbook of Second Language Acquisition*. San Diego: Academic Press, 413–68.

Long, M. H. 1997: Construct validity in SLA research: a response to Firth and Wagner. *Modern Language Journal*, 81, 318–23.

Long, M. H. 1998: SLA: breaking the siege. *University of Hawai'i Working Papers in ESL*, 17 (1), 79–129.

McLaughlin, B. 1987: *Theories of Second-Language Learning*. London: Edward Arnold.

McLaughlin, B. P. and Warfield, T. A. 1994: The allure of connectionism reexamined. *Synthese*, 101, 365–400.

Maxwell, G. 1962: The ontological status of theoretical entities. In H. Feigl and G. Maxwell (eds), *Scientific Explanation, Space and Time*. Minneapolis: University of Minnesota Press.

Morgan, C. L. 1903: *An Introduction to Comparative Psychology*. 2nd edition. London: Walter Scott.

Nagel, T. 1997: *The Last Word*. New York: Oxford University Press.

O'Grady, W. 1996: Language acquisition without Universal Grammar: a general nativist proposal for L2 learning. *Second Language Research*, 12, 374–97.

O'Grady, W. 1997: *Syntactic Development*. Chicago: University of Chicago Press.

O'Grady, W. 1999a: Gapping and coordination in second language acquisition. In K. Kanno (ed.), *The Acquisition of Japanese as a Second Language*. Philadelphia: John Benjamins, 141–58.

O'Grady, W. 1999b: Toward a new nativism. *Studies in Second Language Acquisition*, 21, 621–33.

Pinker, S. and Bloom, P. 1990: Natural language and natural selection. *Behavioral and Brain Sciences*, 13, 707–84.

Ramsey, W. 1997: Do connectionist representations earn their explanatory keep? *Mind and Language*, 12, 34–66.

Ramsey, W. and Stich, S. P. 1991: Connectionism and three levels of nativism. In W. Ramsey, S. P. Stich, and D. E. Rumelhart (eds), *Philosophy and Connectionist Theory*. Hillsdale, NJ: Lawrence Erlbaum Associates, 287–310.

Rey, G. 1997: *Contemporary Philosophy of Mind*. Oxford: Blackwell.

Schachter, J. 1996: Maturation and the issue of Universal Grammar in second language acquisition. In W. C. Ritchie and T. K. Bhatia (eds), *Handbook of*

Second Language Acquisition. San Diego: Academic Press, 159–93.
Schmidt, R. 1994: Implicit learning and the cognitive unconscious: of artificial grammars and SLA. In N. Ellis (ed.), *Implicit and Explicit Learning of Languages*. New York: Academic Press, 165–209.
Schumann, J. H. 1983: Art and science in second language acquisition research. In A. Guiora (ed.), *An Epistemology for the Language Sciences*. *Language Learning* special issue, 33, 49–75.
Schumann, J. H. 1995: *Ad minorem theoriae gloriam*: a response to Eubank and Gregg. *Studies in Second Language Acquisition*, 17, 59–63.
Schwartz, B. D. 1986: The epistemological status of second language acquisition. *Second Language Research*, 2, 121–59.
Schwartz, B. D. 1998: The second language instinct. *Lingua*, 106, 133–60. Reprinted in A. Sorace, C. Heycock, and R. Shillcock (eds) 1999: *Knowledge Representation and Processing*. Amsterdam: Elsevier, 133–60.
Schwartz, B. D. 1999: Let's make up your mind: "special nativist" perspectives on language, modularity of mind, and nonnative language acquisition. *Studies in Second Language Acquisition*, 21, 635–55.
Smith, N. and Tsimpli, I.-M. 1995: *The Mind of a Savant: Language, Learning, and Modularity*. Oxford: Blackwell.
Smolensky, P. 1987: The constituent structure of mental states: a reply to Fodor and Pylyshyn. *Southern Journal of Philosophy*, 26, 137–60.
Smolensky, P. 1995: Connectionism, constituency, and the language of thought. In C. Macdonald and G. Macdonald (eds), *Connectionism*. Cambridge, MA: Blackwell, 164–98.
Smolensky, P. 1999: Grammar-based connectionist approaches to language. *Cognitive Science*, 23, 589–613.
Sober, E. 1994a: Contrastive empiricism. In E. Sober (ed.), *From a Biological Point of View: Essays in Evolutionary Philosophy*. Cambridge: Cambridge University Press, 114–35.
Sober, E. 1994b: Evolution, population thinking, and essentialism. In E. Sober (ed.), *From a Biological Point of View: Essays in Evolutionary Philosophy*. Cambridge: Cambridge University Press, 201–32.
Sober, E. 1998: Morgan's Canon. In D. D. Cummins and C. Allen (eds), *The Evolution of Mind*. Oxford: Oxford University Press, 224–42.
Sober, E. 1999: Testability. Presidential address to the Central Division of the American Philosophical Association in New Orleans, May. *Proceedings and Addresses of the APA*, 73 (2), 47–76.
Sokolik, M. E. and Smith, M. E. 1992: Assignment of gender to French nouns in primary and secondary language: a connectionist model. *Second Language Research*, 8, 39–58.
Sprouse, R. A. 1996: Appreciating the poverty of the stimulus in second language acquisition. *Behavioral and Brain Sciences*, 19, 742–3.
Sterelny, K. 1990: *The Representational Theory of Mind: An Introduction*. Oxford: Blackwell.
Sterelny, K. and Griffiths, P. E. 1999: *Sex and Death: An Introduction to Philosophy of Biology*. Chicago: University of Chicago Press.
Stoljar, D. and Gold, I. 1998: On biological and cognitive neuroscience. *Mind and Language*, 13, 110–31.
Strozer, J. R. 1994: *Language Acquisition after Puberty*. Washington, DC: Georgetown University Press.
Tarone, E. E. 1994: A summary: research approaches in studying second-language acquisition or "If the shoe fits . . ." In E. E. Tarone, S. M. Gass, and A. D. Cohen (eds), *Research Methodology in Second Language*

Acquisition. Hillsdale, NJ: Lawrence Erlbaum Associates, 323–36.

Thomas, M. 1991: Do second language learners have "rogue" grammars of anaphora? In L. Eubank (ed.), *Point Counterpoint: Universal Grammar in the Second Language*. Amsterdam: John Benjamins, 375–88.

van Lier, L. 1994: Forks and hope: pursuing understanding in different ways. *Applied Linguistics*, 15, 328–46.

White, L. 1987: Against comprehensible input. *Applied Linguistics*, 8, 95–110.

White, L. 1989: *Universal Grammar and Second Language Acquisition*. Amsterdam: John Benjamins.

White, L. 1996: UG, the L1, and questions of evidence. *Behavioral and Brain Sciences*, 19, 745–6.

White, L. forthcoming: *Universal Grammar in the Second Language: From Initial to Steady State*. Cambridge: Cambridge University Press.

Wimsatt, W. C. 1986: Developmental constraints, generative entrenchment, and the innate-acquired distinction. In W. Bechtel (ed.), *Integrating Scientific Disciplines*. Dordrecht: Martinus Nijhoff, 185–208.

Wimsatt, W. C. 1999: Generativity, entrenchment, evolution, and innateness: philosophy, evolutionary biology, and conceptual foundations of science. In V. G. Hardcastle (ed.), *Where Biology Meets Psychology: Philosophical Essays*. Cambridge, MA: MIT Press, 139–79.

Woodward, J. 1989: Data and phenomena. *Synthese*, 79, 393–472.