1 What is "Language"?

> there are situations in dealing with natural data – polemical situations, in particular – in which a statement of principles saves waste and leads to the realization that certain factual problems were ill-formulated; or that some factual problems are related to other factual problems as special, or as more general, cases . . . One soon learns that certain formal principles, appealed to indirectly perhaps, but with an implication of evident validity, are less cogent than they are claimed to be. And yet the practices and concrete interpretations which they are designed to buttress do stand up as reasonable and, on occasion, as verifiable.
>
> *H. Hoenigswald (1973: ix–x)*

It is commonplace to observe that "languages are always changing." Indeed, it seems clear that any conception of the nature of language as an object of scientific study which does not account for this universal property of human languages is, regardless of the apparent richness of its empirical support, in a very real sense fundamentally flawed. On these grounds alone some historical linguists regard with deep suspicion any theoretical approach to the study of language which appears to neglect this diachronic aspect of the nature of language. For their part, theoretical linguists quite correctly observe that the fact of change is not necessarily relevant to the synchronic functioning of language. Past change events do not form part of the "competence" of adult native speakers, for example; our knowledge of such events results not from the acquisition process, but from academic study.[1] This apparent "universal property" of human languages therefore seems to differ fundamentally from other claimed universals (e.g., those of Universal Grammar) in that it is not, at least an any obvious sense, the result of contraints on human cognitive systems. To the extent I believe, with most theoretical linguists, that the object of study of linguistics is precisely those aspects of human cognitive systems which constitute our linguistic competence, it would appear that there may be no place for historical linguistics in such an enterprise.

[1] There is an explicit recognition of the lack of knowledge of language history on the part of native speakers in well-established concepts from historical linguistics such as that of "folk etymology." There appears to be no serious dissent from this view.

It is a fundamental belief of the author of this book that the apparent conflict just sketched between the views of historical linguists regarding the inevitability and universality of language change as a fundamental property of linguistic systems and the concentration of modern theoretical linguists on cognitive systems is in fact illusory. The fact of language change is not only wholly consistent with modern theoretical conceptions of the object of study of linguistics, it is, in my view, highly dependent upon them. I envision a new and critical place for the study of historical linguistics in the modern linguistic enterprise as a result of the resolution of this conflict, which I will attempt to sketch in this book.

Historical linguistics, as a discipline, has existed in a form readily recognizable to the modern practitioner for at least 150 years. The wisdom accumulated during this extended period provides the modern-day historical linguist with a tremendous body of information upon which to draw in pursuing his or her research. Perhaps the greatest benefit of this collective knowledge resides in the fact that most practicing historical linguists have a well-developed and reliable sense of what types of change are attested in human languages and what types are not. On the other hand, this sense is largely a pragmatic one. While it was built up from data gleaned by the detailed investigation of many languages (though, of course, in the end only a small subset of human languages), it was not articulated (nor even usually conceptualized) as *derivable*, in principle, from some broader theoretical construct, but rather as a developed and informed intuition arising from the scholar's individual experience and the collective experience of the field as a whole.

It is, I think, hardly surprising that historical linguists favored this "intuition informed by experience" approach to doing linguistics. Theories of the nature of language were, then as now, usually linked to whatever particular conception of human psychology was in vogue at the time, and these theories displayed wide variation during the long history of the discipline of diachronic linguistics. It was fairly clear to historical linguists, at least since the Neogrammarian era of the late nineteenth century, that linking their investigations too closely to psychological theory impeded, rather than enhanced, the investigation of issues which were uniquely their own concern (such as "sound change"). Indeed, the Neogrammarian doctrine represented the first major "theoretical" breakthrough in linguistics. It was not explicitly tied to any particular psychological model[2] – it was a theory about language as such, the insights of which would need to be incorporated into a coherent social or psychological theory (rather than being dependent upon one). It brought a conceptual independence to linguistics as a domain of inquiry which has been steadily increasing since the Neogrammarian era, giving rise directly to the *synchronic* theorizing of Neogrammarians like Saussure and, indirectly, to much of subsequent, including contemporary, linguistic theory.

For all its advantages over previous work, however, Neogrammarian doctrine represented an attempt to directly implement diachronic generalizations about language without a coherent synchronic theory of the nature of language as an object of study. As Saussure asserted early in the twentieth century, such an

[2] Though of course it presupposed one which has never been made fully explicit.

attempt is doomed from the ourset: diachronic generalizations must hold over what are, in fact, a series of synchronic stages. It is not possible in any meaningful sense to know what "changed" between Stage I and Stage II of some "language" without knowing what Stage I and Stage II were, as synchronic systems. It is therefore not possible to construct a restrictive diachronic theory of possible relationships between any two arbitrary sequential stages without a coherent, equally general, synchronic theory. I will attempt to show below that historical linguists have never fully incorporated this fundamental insight of Saussure's into their methodological discussions, much to the detriment of the discipline.

Thus, although historical linguistics provided the discipline of linguistics with its first explicit "theory," this theory was highly stipulative. It did not follow in any direct way from a combination of (1) a conception of the nature of the object of study and (2) a conception of the principles governing the "change" of that object over time. The first is, of course, an explicit theory of synchronic linguistic structure (the definition of "language" as an object of study is a necessary prerequisite to such a theory). The second is the goal of the discipline that historical linguistics is striving to be in the modern era. The necessity of the development of such diachronic principles has long been noted by historical linguists, but progress, in spite of the rich compilation of empirical data which should bear directly on their construction, has been frustratingly slow, particularly in areas other than phonology.

The shortcomings in conceptualization and method which have given rise to this slow development lie in two areas: the failure to incorporate a coherent conception of the nature of language into the working methodology of diachronic linguists and the lack of clarity (in part arising from this failure) surrounding the notion of "change." In what follows I will attempt to reveal where I believe the shortcomings in each of these areas lie and suggest ways in which they might be reduced. I will begin with two fundamental questions: what is "language," as the object of study of linguistics? and what is "language change"?

1.1 Synchronic "Language" vs. Diachronic "Language"

There are, in fact, two distinct issues which arise in attempting to answer the first question in the context of historical linguistics – what is "language" as a synchronic object in the world (independent of its history), and what is "language" in its diachronic aspect (i.e., what does it mean to say that a language "changes" over time)?[3] The first I will simply call "the synchronic question," and it arises quite

[3] I will not at this juncture address arguments, such as those in McMahon (2000), that no principled distinction can be drawn between synchrony and diachrony. The reasons for this will probably become apparent in what follows, but the matter will be explicitly discussed in the concluding chapter.

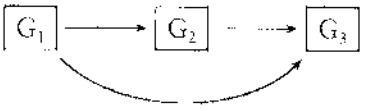

Figure 1.1: Synchrony and diachrony

clearly in the work of Ferdinand de Saussure, perhaps most unambiguously in his discussion of what has come to be known as the "Primacy of Synchronic Linguistics":[4]

> For example – and to begin with the most obvious fact – they [i.e., the "synchronic" and "diachronic" approaches to linguistic analysis] are not of equal importance. Here it is evident that the synchronic viewpoint predominates, for it is the true and only reality to the community of speakers. The same is true of the linguist: if he takes the diachronic perspective, he no longer observes language but rather a series of events that modify it. (Saussure 1959: 90)

We can represent Saussure's model of the relationship between synchronic linguistic states and language change as in figure 1.1. In this figure, each box represents a synchronic state of the language. Diachronic linguistics concerns itself with the differences between these successive states (e.g., between G_1 and G_2, or G_2 and G_3, or even, by extension – as discussed in some detail below –, between G_1 and G_3).

What is in the boxes in figure 1.1? That is, what is the nature of the "synchronic states" G_1, G_2, and G_3? This is the *synchronic* version of the "What is 'language'?" question. As is not uncommon in scientific disciplines which concern themselves with domains of human activity, the answer to this question has proven difficult. In particular, much of the linguistic literature shows clear mixing, in one and the same discussion, between an empirically useful conception of "language" (as the object of study of linguistics) and more everyday uses of this term. Not surprisingly, such ambiguities do little to further the cause of scientific clarity. In this book, I will attempt to keep clearly distinct the widespread, pretheoretical concept of "language," with its predominantly sociopolitical foundations, and the theoretical concept of the "grammar" as the object of study of linguistics as a discipline.[5]

[4] "Par exemple – et pour commencer par le fait le plus apparent – ils n'ont pas égale importance. Sur ce point, il est évident que l'aspect synchronique prime l'autre, puisque pour la masse parlante il est la vraie et la seule réalité. Il en est de même pour le linguiste: s'il se place dans la perspective diachronique, ce n'est plus la langue qu'il aperçoit, mais une série d'événements qui la modifient" (Saussure 1984, 128).

[5] Given the lack of success of my fellow linguists in trying to maintain terminological clarity in this domain, I can only imagine that I, too, at times will fail. I beg the reader's pardon in advance for any such lapses.

"Language," like many terms which have widespread informal, pretheoretical use, can be conceived of in a variety of ways.[6] The discipline of linguistics could, in principle, adopt any of a number of these. "Language" as a sociopolitical notion (reflected in terminology such as "the English language") represents a set of distinct ways of speaking – the speakers of this "language" typically feel bound to one another through their self-designated identity as, for example, "speakers of English." It has been clear for some time that this particular conception of "language" is difficult to systematically implement in linguistic research: there is no empirical test which allows us to determine that, e.g., the English of a Yorkshire village and of Ypsilanti, Michigan, are manifestations of the same entity ("English"), while a given dialect spoken in Denmark and a similar dialect spoken in Norway represent distinct entities ("Danish" and "Norwegian"). The standard attempts to develop tests to establish "languagehood" (such as mutual intelligibility, shared core vocabulary, etc.) have not proven fruitful.[7]

This is not to say that one couldn't establish some criterion or set of criteria for identifying a grammar as "English" – one could pick some particular piece of morphology, a syntactic structure of some type, a lexeme, or some combination of these. There is no limit to the human ability to create arbitrary categories. The question is whether there are any motivated, i.e., nonarbitrary, criteria available – criteria which solve the *scientific* problem of defining "sociopolitical" languages.[8]

Likewise, this is not to deny that linguistic self-identity on the part of speakers may affect their behavior (including their linguistic behavior, e.g., in willingness to adopt innovations), but only that at base the notion is not a "linguistic" one. After all, many factors influence the spread of linguistic innovations – socioeconomic status, dress, ethnicity, age, gender, etc. – without it being necessary to see these various aspects of speaker status and behavior as the proper object of an explicit linguistic theory. We must distinguish between the behavior of individuals and the nature of their linguistic competence. The interaction of issues such as language choice – including here selection between dialects – and general theories of human behavior provides valuable material for the study of human behavior, but it is only when we can precisely characterize the linguistic competence of an individual – in detailed, explicit, and empirical terms – that we can even begin to coherently investigate the much more complex question of the implementation of that competence to perform specific (e.g., social) tasks.

[6] The discussion of Chomsky in *Knowledge of Language* (1986) is still foundational on this matter.

[7] For a clear demonstration of the shortcomings of these techniques in actual practice, see Schütz (1972).

[8] Note that arbitrarily selecting some property *x* as criterial for "English," some unrelated property *y* for "French," and yet a third criterion, *z*, for Marshallese, would never provide us with a scientific answer to the general question of what is "sociopolitical" language. For each new language we come across, or which comes into being afresh, a new criterial property would need to be posited.

We next to turn to the question of how the concept of "language" as an object of scientific inquiry has come to differ from its common pretheoretical predecessor.

1.2 "Language" as a Synchronic Object

What is "language" as a *synchronic* object of scientific study within the science of linguistics, then? As noted above, linguists over the centuries have answered this question in a variety of ways. For example, in the nineteenth century, when the focus of linguistics as a discipline was on a relatively small set of dead languages, it was clearly commonplace to assume that a "language" was a *corpus*, a set of linguistic forms encoded in the surviving records for a given period of a given linguistic tradition. This conception of the object of study of linguistics turns out to be seriously inadequate when we, as linguists, turn our attention to living, spoken languages. First, corpora have a number of properties which living languages lack. Through their fixed nature, for example, they have limited vocabularies, indeed, even a highly restricted set of utterances (those which happen to be attested in the extant record of the language). By contrast, living languages possess numerous processes which allow the productive creation of new words (e.g., via morphological derivation, including compounding, and inflection) and cannot in any meaningful sense be said to consist of a finite and listable set of sentences. It seems unlikely, for example, that, although you are a "speaker" of English, you have ever before been confronted by the specific sentences you have read thus far in this book. Second, living languages generally provide evidence – e.g., detailed phonetic data – which we could never extract from the historical record of "dead" languages.

It is universally acknowledged that the features found in "dead" languages which make them unlike living languages are not due to the fact that these languages were in fact different in kind from their modern sisters, but rather must be attributed to the accidental nature of their attestation (which artificially restricts our access to them). No one believes that the sentences preserved on our Hittite tablets were the only sentences Hittites could say![9] It seems clear that we should not define the object of study of linguistics on the basis of such

[9] While this is, I assume, universally recognized, some of the relatively straightforward implications of this conception of things appear to be missed in some "corpus" linguistics. For example, one often sees fairly dramatic conclusions being drawn on the basis of the statistical distribution of some particular forms (in some particular sentences) in a dead language. Nevertheless, there is no reason to believe that the Hittites said (the Hittite equivalent of) "I will destroy his land" more often than they said "Meet me here tonight," but the former sentence, and thus the morphological objects in that sentence, could easily occur in the corpus *far* more frequently than the (unattested) latter. This is the normal state of affairs when dealing with a dead language.

fundamentally irrelevant properties as accidentally preserved records. That is, we must not confuse the nature of the attestation of a language with what that language was when alive. While we will need special tools – philological tools, in particular – to extract linguistically relevant and significant data from the corpora of dead languages, the so-called Uniformitarianism Principle demands of us that we assume that the languages of antiquity and other earlier historical periods were in fact no different in their fundamental properties than the living languages around us. That is, the extant record of a dead language provides us with a glimpse at the "tip of the iceberg" of what was in fact a vital, living language, no different in its basic design properties than the living languages which surround us today. The limited set of sentences and related morphological limitations are simply accidents of attestation. This matter will be discussed again when we turn to the question of the diachronic nature of "language" below.

The question of the nature of synchronic language states as the object of scientific investigation has been addressed perhaps most prominently by Noam Chomsky in numerous works, nowhere as clearly as in his well-known *Knowledge of Language* (1986). Chomsky distinguishes between "externalized" or E-Language – language as something observable, out there in the world – and "internalized" or I-Language – language as a ***knowledge state***.

> Let us refer to such technical concepts as instances of "externalized language" (E-language), in the sense that the construct is understood independently of the properties of the mind/brain . . . Let us refer to this "notion of structure" [in the mind of the speaker] as an "internalized language" (I-language). The I-language, then, is some element of the mind of the person who knows the language, acquired by the learner, and used by the speaker-hearer. (Chomsky 1986: 20–2)

As Chomsky has argued at length, the proper object of empirical linguistic investigation must be I-Language.[10] The sociopolitical (or "normative," as Chomsky calls it) concept of E-Language is, unfortunately, fairly close (to the extent it can be made sensible) to our everyday use of the term "language" in phrases such as "the English language" or "the Marshallese language."[11] In my opinion it is quite commonplace in the literature for major debates to arise as a result of linguists' deft ability to switch unconsciously in the course of their discussion between

[10] As Chomsky notes regarding E-Language: "That any coherent account can be given of 'language' in this sense is doubtful, none has been offered or even seriously attempted. Rather, all scientific approaches have simply abandoned these elements of what is called 'language' in common usage" (Chomsky 1986: 15).

[11] As Charles Reiss regularly points out to me, equating the "sociopolitical" or "normative" conception of language fully with Chomsky's concept of "E-language" represents a simplification of Chomsky's characteristically complex thinking on these matters. I do not consider the simplification to do any particular damage in the present discussion, and will therefore persist in this usage. The reader should of course refer directly to the discussion in *Knowledge of Language* for a more sophisticated version.

these two quite distinct senses of "language." In this book, I will have occasion to point out several instances of this unfortunate phenomenon.

It will therefore be useful if we can be as precise as possible about what the notion "I-Language" encompasses and precisely how this concept differs from the more common notion of "E-Language." However, it is of some value to point out that, in general, discussion of linguistic issues is phrased, even in relatively nontheoretical work, in such a manner that it implicitly recognizes the importance of I-Language for linguistics. For example, "language acquision" is typically treated as the coming into being of a particular knowledge state (I-Language) in a particular human being. The notion of E-Language – whether defined as a "set of utterances" in use in a particular "speech community" or in some related manner – plays no central role in the study of this process.[12] The transmission of grammatical competence does not consist of the transfer of sets of utterances from one generation to another, but rather in the coming into being in the mind of the acquirer of a system for generating linguistic representations (which may then be passed off either to a production system, in the case of speech, or to a processing system, in the case of parsing). It is clear that parts of this generative system (the I-Language, which we will call the "grammar") owe their existence to the input data received by the acquirer, but it is equally clear that the relevant aspects of that input data are not the utterances themselves, i.e., the specific set of sentences, but rather the set of linguistic structures which the acquirer has posited in developing an account of the primary linguistic data which formed the basis for his/her grammar construction.

Speech communities, to the extent this notion could be turned into something empirically well-defined, would seem to need to acquire their properties through the concatenation of the properties of the individual I-Languages which are present within the community. The set of "potential utterances" of such a speech community,[13] i.e., the E-Language of that community, is therefore, in any event, at best a derivative notion, deducible from a sufficiently rich understanding of the relevant set of I-Languages, which we will therefore be required to construct anyway. Advocates of E-Language approaches not only lack a clear definition of what an E-Language might be, but also fail to provide any compelling statement of what it is we gain by seemingly arbitrarily constructing these E-language objects from our posited I-languages. What is it we can learn from studying these new constructs that can't be discovered from an examination of their component I-language entities?

To answer the question of what a particular I-Language is with some greater degree of precision, it is important to acknowledge in the first instance that

[12] That is, acquisitionists do not typically collect the set of sentences in use in a particular speech community and then contrast these with the set of sentences used by acquirers. What are compared are linguistic *structures* rather than surface (output) strings.

[13] This conception of the object of study of linguistics was widely held by American structuralists. Bloomfield (1933) states that "The totality of utterances that can be made in a speech-community is the *language* of that speech-community."

the notion does not necessarily refer to the generative systems which give rise to *all* of the linguistic representations of a single individual. Multilingualism and polydialectalism are both well attested (if not nearly universal) properties of individuals. It seems clear that attempting to develop a linguistic analysis of a so-called "Spanish-English" bilingual by treating, e.g., what is usually referred to as "Spanish"-type output and what is usually referred to as "English"-type output as the product of a single linguistic system (with a single morphological, phonological, and syntax module) is not likely to be very productive. In the case of a so-called "Spanish-English" bilingual speaker it is useful to posit *two* (at least) distinct grammars in the brain of a single individual.[14] It also does not seem possible, or even desirable, to impose the sociopolitical distinction between "language" and "dialect" on our I-Language models. Multiple production systems, whether responsible for generating representations which are treated in pretheoretical terms as "different languages" or as "different dialects of the same language," are, by definition, different grammars (i.e., distinct I-Languages). Grammars must be distinct if they generate nonidentical sets of representations.[15]

Each I-Language ("grammar") is thus responsible for generating a particular set of output representations. It is of some significance, since I will return to the point in the discussion of specific change mechanisms later in this book, that "output representation" is distinct from "output." The actual acoustic[16] output of an individual is, of course, the result of "playing" the grammar's output representation through a particular system (the "performance" or "production" system) under some specific set of (cognitive and physical) conditions. Acoustic output, due to the nature of the conditions governing its form, may be highly variable. The output representation generated by the grammar, which consists, if we consider only production for a moment, of a set of features, is assumed to be constant for some particular sentence (to the extent it was generated by a single grammar, almost certainly the norm). Clearly, if a particular set of features is played through a system which is producing output at a rapid rate, for example, the acoustic output will differ considerably from the same set of features played through the same body in a context in which it is producing its output at a slower rate. The variability of output in these contexts is conditioned extralinguistically. As noted above, the contrast between those aspects of the output which are

[14] Note that this assertion is neutral on the question of the precise details regarding, e.g., the storage of "overlapping" information. It is entirely compatible with a model in which duplicated "knowledge" is stored only once, but accessed by both grammars, much as the interface modules – responsible for, e.g., prelinguistic acoustic parsing and articulation – are.

[15] The amount of "overlapping" information, possibly stored once but accessed by both generative systems, may differ considerably in the cases of bidialectal vs. bilingual speakers, of course. The evidence is not clear on this point, nor does it seem particularly important to worry about the question at this time.

[16] I will take the license of using terminology appropriate to spoken language in this discussion, though the same generalizations hold of other natural human linguistic systems, such as signed language.

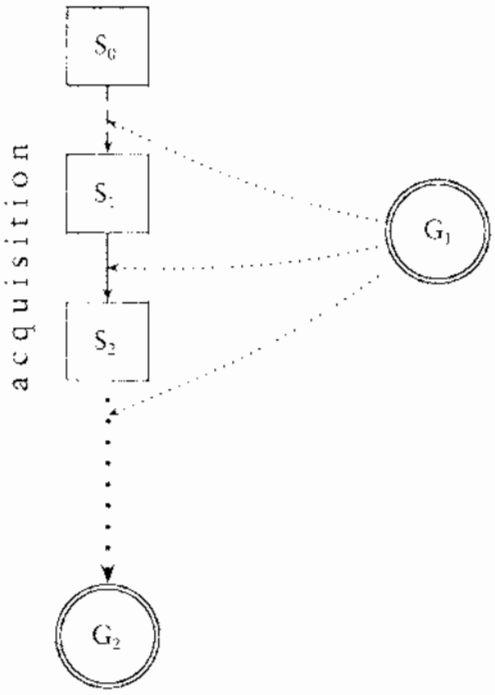

Figure 1.2: The nature of acquisition

conditioned by features of the grammar and those aspects which are conditioned by extragrammatical factors will play a key role in the discussion of certain types of change events to be considered below.

"Language," therefore, for linguistic purposes is equivalent to I-Language, i.e., the "grammar." The grammar itself is a knowledge state. This knowledge state arises in the acquirer through the interaction of the primary linguistic data (PLD) to which the acquirer is exposed during the acquisition period with the initial state of knowledge of language. This innate initial state is usually referred to as "Universal Grammar" (UG). We might represent the acquisition process as in figure 1.2. In this figure, S_0 represents the state of knowledge in an acquirer at birth, before exposure to the primary linguistic data of his/her environment. This knowledge state changes upon exposure to the PLD, such changes being the "stepwise" developing chain of knowledge states ("grammars") represented as S_1, S_2 . . . in figure 1.2.[17]

The most important feature of the acquisition process for our purposes in this book is the "terminal" state, G_2. The existence of this state in figure 1.2 reflects

[17] There are several niceties involved in this figure whose full explication must await the discussion of language acquisition and language change in the concluding chapter. One of these is the difference between the lines connecting S_0 and S_1 and those connecting other stages in the acquisition process. I'll defer discussion of these points at this time.

my belief that the acquirer, in the course of a given acquisition path, actually attains a point at which acquisition of that particular grammar can be said to be complete. Subsequent data in the PLD which appears to diverge from the current knowledge state of the acquirer is taken as evidence for *other* grammars after this point, whereas before this point, it may be taken as evidence that the grammar currently hypothesized (e.g., S_2 in figure 1.2) is in need of revision to some other knowledge state (e.g., S_3 in figure 1.2).

Note that nothing in this model precludes an acquirer from simultaneously engaging in several distinct learning paths towards distinct targets (e.g., towards a variety of "English" and a variety of "Spanish," to put the matter informally). Nor does it prevent the acquirer from beginning a new learning enterprise well after their first grammar has been fixed (as in, e.g., late childhood or even adult second-language acquisition). The key distinction between L2-acquisition and the "stages" in acquisition represented in figure 1.2 can be described as follows. State S_2 in figure 1.2 represents a revision to knowledge state S_1 based on the evidence provided to the acquirer by the PLD. Basically, it reflects the fact that the acquirer has reached the point where they *reject* state S_1 in favor of S_2 as a plausible grammar for generating the data in the PLD. Since the target of acquisition is a knowledge state which will generate such data, the rejection of S_1 in favor of S_2 is a direct result of the acquirer's informed hypothesis that S_2 offers a "better match" to the system generating the data the acquirer is receiving than does S_1. The "stages" in acquisition thus represent *replacements* of hypothesized grammars which were well grounded at an earlier point but which ultimately were deemed less consistent with the evidence.

We can contrast this process with that of the acquisition of a "second language" by an acquirer who has already attained the terminal state, G_2, for their first acquisition target. It is manifestly obvious that a 7-year-old child who is a competent speaker of some variety of "English," suddenly exposed to steady and sufficient input evidence for some variety of "Spanish" such that s/he begins to acquire this variety of "Spanish," does not engage in a *replacement* of their current knowledge state (of some variety of "English") by this new, "Spanish-like," knowledge state. That is, whereas in the intermediate stages of a given acquisition path additions to knowledge represent replacements of now outdated and contraindicated hypotheses, in the case of L2-acquisition no already acquired, steady-state L1 systems in the acquirer get "replaced" by new, L2-consistent knowledge states. Instead, a new learning path, with a new acquisition target, is begun, the initial L1 knowledge remaining intact. Put simply, the recognition of the existence of, e.g., a variety of "Spanish" in one's linguistic environment does not trigger a complete rejection of an acquired grammar of a variety of "English," in spite of the fact that the existence of this variety of "Spanish" in one's environment is inconsistent with the knowledge state a monolingual speaker of a variety of "English" is in.

It is clear that in this model – a model which will play a critical role in our discussion of the nature of language change – there is an important distinction

between easily rejected, tentative intermediate hypotheses (the S_1, S_2, etc. of figure 1.2) and the fixed, established knowledge state (G_2) which represents an acquired grammar. There is implicit recognition of this contrast even in pretheoretical work in linguistics. For example, historical linguists do not accept, as the basis for the study of language *change*, a comparison between my grammar and that of my 4 year-old child. Put another way, the "synchronic states" being compared in Saussure's sketch (figure 1.1 above) of the object of study of diachronic linguistics must be mature G_2 states – the differences between my grammar and that of my child do not represent language change.

Imagine a contrary hypothesis – one which held that *grammatical* (as opposed to lexical) learning continued throughout one's lifetime.[18] That is, imagine a model within which G_2 is never acquired such that throughout one's life one continued to modify one's knowledge state just as one did during the early stages of acquisition, constantly "correcting" one's knowledge on the basis of additional evidence provided by the PLD. We could depict this hypothesis as in figure 1.3.[19]

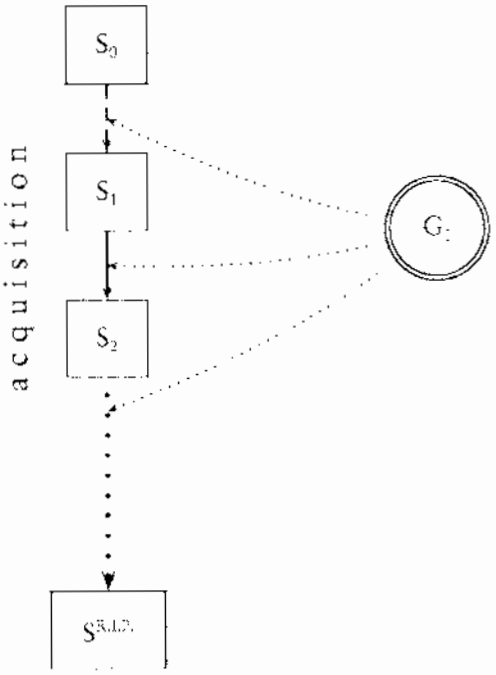

Figure 1.3: Acquisition without end

[18] There is little question that one's language behavior can change during the course of one's mature adulthood. Once one acknowledges the necessity of a contrast between "knowledge" and "behavior" – in particular, once one accepts that the set of factors which influence one's behavior, including one's "language behavior," is much broader than just the "grammar," the relevance of this fact becomes considerably more problematic.

[19] $S^{R.I.P.}$ in this figure represents the speaker's knowledge state at his/her death.

Note that under such a hypothesis there would be no justifiable reason for setting some arbitrary point, S_n, during the acquisition process, after which we could use the evidence provided by this grammar as reflecting a well-defined synchronic state. The grammar of a 30-year-old and the grammar of a 30-month-old would both represent simply a point somewhere along this never-ending chain of development. Under such an assumption, the differences between my grammar and that of my 4-year-old child *would* represent "language change," at least with equal justification as that gained by adopting any other point, clearly an undesirable result.[20]

To summarize, then, it would appear that the most coherent and productive conception of "language" as a synchronic object of study is that of the I-language grammar. Indeed, it is not clear that there is any even vaguely coherent alternative. In my experience, traditional historical linguists find this conclusion unsettling, feeling that it may undermine many of the hard-earned scientific gains of their field. This book argues in some detail that this discomfort is unnecessary.

1.3 "Language" as a Diachronic Object

It is possible that it would turn out that the optimal scientific conception of the nature of language for synchronic linguistics could be distinct from the optimal conception of that object for diachronic purposes. This seems *a priori* unlikely,[21] but in any event an even cursory examination reveals that sociopolitical, E-Language conceptions of language fail to provide a usable foundation for diachronic work, just as they do for synchronic.

Imagine that we attempt to define language, for diachronic purposes, as something akin to, though hopefully more precise than, the pretheoretical notion of "the English language" and the like. We could conceive of language in this sense as, e.g., a set of grammars (i.e., I-languages) which share some criterial property or properties.[22] Note that it would no longer be the case that such "languages"

[20] Or at least not a result which would be accepted by any historical linguist, theoretical or traditional. If you believe that this represents a desirable result, it would be incumbent upon you to give some coherent explanation for how one might do historical linguistics under such an assumption and for why historical linguistics appears relatively successful using what would turn out to be, if you were right, strongly counterfactual assumptions.

[21] Note that it would be rather surprising that under such a conception "diachronic" languages, like the distinct objects known as "synchronic" languages, have a phonology, with features and segments and rules, or processes, or constraints or whatever, and a morphology, and a syntax, etc.

[22] Since this conception will be summarily rejected in what follows, I will not belabor the problems (mentioned earlier) involved with establishing in a valid, motivated, and compelling manner what criteria might be invoked in such a case.

have things like "a phonological system" in any meaningful sense (they "have" only their members, i.e., individual I-languages). If one wanted to create some superordinate entity which did "have" things like a phonological system, we would instead need to build this object by extracting information from the individual I-languages and compiling it in some manner into a "language." One can imagine ways to do this, but none seem likely to be of much use in diachronic linguistics. For example, one could adopt the position[23] that the phonological inventory of the superordinate "language" was the union of the phonological inventories of the relevant I-languages. The inventories of "languages" being used for diachronic purposes would then suddenly be massively unlike those found in any natural human language, containing far more segments and making *phonological* distinctions nowhere attested. Alternatively, one could take the position that the phonological inventory of the superordinate "language" was the intersection of the phonological inventories of the relevant I-languages. Unfortunately, "languages" in this sense would have dramatically smaller inventories (in some cases possibly null) than any attested human language.

It is of course not possible for me to discuss the likely uselessness of every possible operation one could perform on a set of I-languages to try to generate some conception of "language" for diachronic purposes which differs from that of I-language itself. However, it seems clear that any such notion would be at best a derivative one, since it builds on the concept of I-language, already available for scientific exploitation. It is hard to see what "language change" would be under any "superordinate" set definition of "language" in any event: Does the "language" contain tokens of I-languages or types?[24] Does the "language" change if my linguistic descendants develop front round vowels even if some speakers in Northern England already have them? What if the distribution of the front rounded vowels differs in the two grammars? Does this affect the phonological system of the superordinate "language"? Fortunately, no such conception has ever been seriously *employeds* in diachronic linguistics, in spite of the fact that many historical linguists have asserted that some such "superordinate system" conception of "language" is what they are adopting. And this grants me an opportunity to emphasize a point to which we will return time and time again in this book. The rhetoric of scholarly disagreement and dissent – both of which are healthy and necessary elements of the scientific enterprise – is shaped largely by the sociocultural landscape of the field, including the complexities of interpersonal (and interinstitutional, and international) relationships, and a myriad of other out and out mysteries of the ways of the world. It is very often the case that the

[23] It is worth pointing out that since these matters are purely definitional, one is free to adopt whatever definition one likes. However, purely stipulative and unmotivated definitions, designed merely to avoid the technical problems arising from an inaccurately structured initial definition, are generally not very compelling.

[24] Note that if the answer is "tokens," the causes of death in human speakers become mechanisms of language change, a seemingly undesirable consequence.

rhetoric employed by historical linguists -- both more "traditional" in orientation and more "theoretical" in orientation -- in discussing the foundational assumptions behind their work bears little or no relationship to the actual assumptions which must be posited if we are to ground their *methods* logically. If we look beyond the rhetoric we will often find that scholars rabidly opposed to one another's research, nevertheless use methods which make identical assumptions on matters of seeming controversy.

So let us return to the I-language conception of language as a possible foundation for diachronic as well as synchronic investigation. While the synchronic version of the "What is 'language'?" question concerns what is in the boxes in the Saussurean figure 1.1, the diachronic question asks what licenses the connection between the boxes by the arrowed lines. That is, under what conditions would we be justified in treating a given G_1 as an appropriate object to be linked to G_2 by an "arrow of diachrony"? Put another way, the question about the contents of the box G_1 is about synchronic identity (under what conditions can we treat output as evidence for "the same" grammar), and the question about the arrow between G_1 and G_2 is about diachronic identity (under what conditions can we treat the relationship between G_1 and G_2 as one of "descent")? We can work towards the development of a useful answer to this question by a consideration of the nature of historical evidence and the scientific investigation of historical records, i.e., by introducing the field of philology.

1.4 Discussion Questions and Issues

A. Consider the following alternative definition of "language": "Language" is a set[25] of distinct grammars (i.e., grammars with different – although perhaps only "minimally" so – properties), such as, e.g., those which we might, sociopolitically speaking, consider to be present in the minds of "English" speakers, or "French" speakers, or "Marshallese" speakers, or whatever. What would "language change" be, under such an assumption, and what practical or methodological issues about the object of linguistic study might arise in pursuing research under such an assumption?

[25] I'll phrase these definitions in terms of set theory like this, because I think it's a useful way to think about the possibilities I have in mind. It is worth noting that a set doesn't have the properties of its elements – e.g., the set of all squares does not itself have four right angles, or four sides of equal lengths, etc. So a set of grammars will not have the properties of those grammars. This should become clearer as you read the subsequent proposed alternative definitions of "language," which are designed to tease these issues apart.

B. Consider the following alternative definition of "language": "Language" is the *intersection* of all the elements of a set of distinct grammars such as, e.g., those which we might, sociopolitically speaking, consider to be present in the minds of "English" speakers, or "French" speakers, or "Marshallese" speakers, or whatever. Thus, for example, if every one of the grammars of "Marshallese" (sociopolitically defined) had a /k^w/, or postposed definite articles, or whatever, then the "Marshallese language" would have said property, but if a single "Marshallese" grammar lacked that property, the "Marshallese language" would also lack it. What would "language change" be, under such an assumption, and what practical or methodological issues about the object of linguistic study might arise in pursuing research under such an assumption?

C. Consider the following alternative definition of "language": "Language" is the *union* of all the elements of a set of distinct grammars such as, e.g., those which we might, sociopolitically speaking, consider to be present in the minds of "English" speakers, or "French" speakers, or "Marshallese" speakers, or whatever. Thus, for example, only if any one of the grammars of "Marshallese" (sociopolitically defined) had some property such as /k^w/, or postposed definite articles, or whatever, then the "Marshallese language" would also have that property. What would "language change" be, under such an assumption, and what practical or methodological issues about the object of linguistic study might arise in pursuing research under such an assumption?

D. Consider the following slightly varied version of the [B] scenario above. Instead of selecting the grammars sociopolitically, imagine that we take a set of properties P_1, P_2, P_3, . . . P_n and use these linguistic properties to define a set of grammars. Any grammar defined by possession of all of these criterial features will be considered to be part of "language X." Take the defining properties to be the (relevant) properties of the "language," and ask yourself what would "language change" be under such a scenario, and what practical and methodological issues about the object of linguistic study might arise in pursuing research under such an assumption.

E. I assert above (in the discussion of figures 1.2 and 1.3), controversially it seems, that there is a termination point to the acquisition process – i.e., that at some point, relatively early in life, one is "done" with acquisition proper. A contrasting view, briefly dismissed above, holds that one continues to change one's grammar (i.e., "learn") throughout one's life (at least potentially). Discuss the implications of removing the "termination point" in the acquisition process and using the alternative assumption of continuous, life-long grammatical learning for the development of a useful definition of "change." That is, if acquisition simply continues unabated throughout one's life, at what point does it make sense to think about a "change" having taken place, and what justifies selecting the particular point you posit?