ETHL Edinburgh
Summer 2026

The Divine Comedy

Foundational Issues in Linguistic Change and Reconstruction


ncipit Comedia Marci Hale, Americani natione, non moribus

Prologue

In this series of lectures we will explore some foundational notions regarding the methodology of historical and comparative linguistics, with a particular view to the critical assessment of how previous and contemporary work fares under a fairly explicit laying out of our research assumptions. The first topic will review my assumptions about doing 'theoretically informed' historical linguistic research, and explore some ways in which recent scholarly work in this field by others is or is not congruent with these assumptions. It will focus on phonological issues. The second topic will address some contemporary approaches to the syntax and prosody of clitics, with general lessons about diachronic syntax in general, and Proto-Indo-European syntax in particular. That will wrap things up. To what extent we will be able to get through all the material in this document is not at all clear to me at present. We may need to revise our roadmap en route.

1. Prima pars, que dicitur Inferno

L'enfer, c'est les autres. (Sartre, *Huis Clos*)

1.1. Preliminaries: What Is Historical Linguistics?

1. Let us start very broadly, with historical linguistics as a whole. For me, there are two fundamental goals of our work in this area.

2. Fundamental Goals of Historical Linguistics:
    - To help us in our quest to discover the nature of the human mind.
    - To increase our capacity for the interpretion (and appreciation) of the historical products of that mind. To that end, it is useful, valuable, and enriching to develop an understanding of the thoughts, lives and creative products of humans from the past.

3. There can be no question that Historical Linguistics as a discipline must play an absolutely central role in both of these pursuits. Learning how the human mind works is one of the greatest benefits science can give us. In general, though, at present we are so bad at this task that no one would or should use science to regulate their relationship to their spouse, or to set up a system for the rearing of children. Scientific understanding is far from being even vaguely helpful in such domains.

4. However, one of the very few components of the human mind that we *have* made progress on is that of the computational system which gives rise to linguistic output (the grammar), i.e., *language*. I will try to lay out here the unique role that *historical* linguistics plays in developing a proper understanding of this important component of mind.

5. The linguistic analysis of textual records left to us as cultural and artistic artifacts is *critical* to developing an understanding of the content of those records. Language is a computation device that constructs compositional meaning out of stored morpho-syntactic and phonological objects. Therefore, if you do not know what the phonology and syntax of a sentence found in an ancient text is, you do not know what it *really* means. For the ancient languages I work on most, languages like Vedic Sanskrit or Attic Greek, I am not confident that there is a *single sentence* that we truly understand.

6. I will have occasion to make reference to the 'human mind' in these lectures far more frequently than you are probably used to seeing in historical linguistics research. Why do I find this necessary? The need to focus on the *mind* in attempting to explain linguistic phenomena can be made clear by a very simple example. In my dialect, vowels are nasalized before following nasals,[1] thus a form such as /sænd/ 'sand' has the phonetic output form [sæ̃ːnd] in my speech.

7. But can the nasal in 'sand' *cause* the nasalization in the vowel? Let's look more closely:

Figure 1: Causality: Mind vs. World

8. There is a simple, well understood fact about the world that becomes relevant in our efforts to answer this question: the future cannot *cause* the past. In the figure above, *t* represents the 'arrow of time' (i.e., the directionality of time); the dashed line from [n] to [æ̃ː] represents the potential 'causal' role of the [n] on the physical implementation of the [æː]. It is quite easy to see that the 'causality' arrow (from [n] to [æ̃ː]) in this picture *violates* the requirement that the future not cause the past.

9. We can also see this if we imagine that someone stuck a sock in my mouth to get me to shut up (an act for which you will doubtless feel great sympathy by the end of this series of lectures) just as I was reaching the end-point of the pronunciation of the [æ̃ː] in 'sand', and that, therefore, I was not able to finish saying the word—specifically, I never actually utter, in this scenario, the [n] in this string.

[1] To a much greater degree than is physically necessary—thus we posit a rule to get the observed nasalization. There is also vowel lengthening in my dialect in this context.

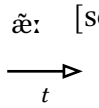

Figure 2: Causality: Mind, World, Sock

10. The [n] in question would therefore never have existed in the physical world (and thus couldn't *cause* anything to happen in that world), but the [æ̃:] would obviously still be nasalized.

11. How is this possible? There is only one place in the world where there is *apparent* causality from the future to the present: in a mind. We do things all the time *in anticipation* of what will happen. Of course, we are often wrong—thus making it clear that it isn't the *actual* future that is causing our present actions, but our current representations of what we believe the future will be like. The phonological rule in my dialect converts the /sænd/ representation into [sæ̃:nd] *in my mind*, before I engage in any actual behavior. At the point at which I plan the physical implementation of that target, my goal is to say an [æ̃:]. This goal is *no longer* dependent on the following nasal: that dependency was *phonological*, not *physiological*, that is, it is in the *mind*, not in the physics of the natural world.[2]

12. The tool we will be using to make progress on these two important issues is *science*, for the simple reason that this is what we do when we try to develop a systematic conception of some domain of the natural world. Only when science fails us (e.g., in the 'how do I regulate my personal relationships?' question) do we fall back on non-scientific reasoning. And if you've 'enjoyed' some personal relationships, you know how *that* sometimes works out! Science has not yet failed us in the domain we will be exploring in these lectures—indeed, I think there is good evidence that it has given us sufficient encouragement that we should be hopeful that it can provide the appropriate framework for our explorations.

13. This 'scientific' stance imposes certain obligations upon us. For example, we must strive (as best we can) to define the terms we use, taking special care not to use the same technical terms in distinct senses. Moreover, if I claim that processes A, B, and C defined over entities X, Y, and Z explains some range of data ℜ, and you show that the data is equally well accounted for by only processes A and B (defined over the same entities), your theory is the better theory (this is *Occam's Razor*, which holds that *entia non praeter necessitatem multiplicanda sunt*, i.e., 'entities are not to be multiplied beyond necessity').

[2]Of course, we are making a subtle distinction here between the physical world *outside* the mind and the, likewise presumably physical, world *within* the mind. The central point is that within the (physical) world inside the mind the representation /sænd/ is atemporal: it is a representation all of whose parts exist at the same point in time. In the (physical) world outside the mind, as we actually say the word, there is temporal sequencing between the relevant articulatory gestures.

14. Fundamental to modern scientific pursuit are the arguments of Karl Popper. His approach requires that we start with the minimal machinery (following Occam) and expand that machinery only incrementally as required. The most restrictive hypothesis is always to be preferred, because it is more readily *falisified*. Falsification is a crucial concept, but it is important to understand that it does not arise from a simple failure to fit to observed data. The data of the natural world is messy, and historical linguistic data is no exception to this general observation. Building uninsightful models which 'capture' the data, but do so at the cost of capturing any *conceivable* data, is not progress.

15. Take a simple example (to which we will return in out syntactic discussion). Imagine that Scholar A proposes that Attic Greek has 'free word order', and Scholar B proposes that it has clauses of the general structure given below (where C represents a position for subordinators and interrogative/relative elements, and with other details left unspecified at present).

16. Two Theories
    - **Theory A:** free word order
    - **Theory B:** Null Topic C Focus Subject Verb Object

17. It is apparent that the theory of Scholar A will be consistent with every piece of data any of us could ever provide regarding the word order of Attic Greek. In every instance, the words in Attic Greek clauses come in some order! There may be numerous seeming exceptions to the observations of Scholar B. Which is the better theory?

18. It is my strong belief, and I hope it will become clear as we move along that this belief is correct, that the theory of Scholar B is *always* better than that of Scholar A, even though the theory of Scholar A is consistent with much more (100%!) of the data than is the theory of Scholar B!

19. Why is this the case? It is the case because if we examine a given Attic Greek clause, provided it is sufficiently rich in its structure, Scholar B will have made specific claims (hopefully testable, or eventually so) about the interpretation of the clause—that XP is a 'topic', that YP is in 'focus', etc. What will Scholar A's analysis tell us about what the sentence means? Nothing! Indeed, since Scholar A posits no *mechanism* for constructing sentences, and demands no match between syntactic form and meaning, her/his analysis cannot help with the interpretation of the clause at all.

20. In short, it is not enough for our theories to align with our data, they must help us with our task, which includes developing a richer understanding of (1) the textual data we are analyzing and (2) the nature of the linguistic capacities of humans. You will see many types of theory as you go along in your academic life that do a fair job of data-matching, but provide no help to us (in my view) on our more global undertaking.

21. If this orientation is correct, the worst property a theory can have is not a failure to account for all of the data—that is virtually inevitable for all but the one true theory, which we can only approximate; the worst property a theory can have is that it accounts for the data, but achieves that accounting in a manner which makes it *irrelevant* for our higher purposes. A theory which fails to match the data but does so in an interesting and productive manner can move us forward; a theory that 'accounts for' the data by telling us anything is possible—and thus, of course, that whatever we observe must be—can tell us nothing new. After all, we presumably already know that what we observe falls within the possible!

1.2. Synchrony, Diachrony, and Typology: Why Historical Linguistics?

22. You may well be convinced, then, that the areas of investigation stated above are worthwhile, and that *science* is the way to pursue the investigation, and that *linguistics* is uniquely suited to help us make progress in these areas. But why not simply pursue *synchronic linguistics*? Won't that tell us all that we need to know?

23. My answer to this is not the one most often found in the historical linguistics literature (at this point that probably won't surprise you). It is not unusual for that literature to claim that diachronic and synchronic linguistic phenomena are 'intertwined' or 'impossible to separate'. Parallel to this is the frequent claim that the syntax of element X is 'undergoing change', is 'on its way' from function A to function B, and that only by understanding the diachronic trajectory of X can we provide a meaningful synchronic account of its behavior.

24. These views, which refuse to keep clear and distinct the question of the nature of an entity in the natural world from that of its aetiology, are, in my opinion, misguided attempts to assert the relevance of diachronic linguistics in an academic landscape in which synchronic approaches have become dominant. For any object in the natural world we can ask what processes earlier entities have undergone to give rise to that object. And clearly, for any object, it will exist and have the properties it has *as the result of* its history.

25. But it is a simple error of logic to assert that because the properties of some entity x have come into being over time, we cannot examine the object and determine what its properties are independent of that origin. Saussure already saw this quite clearly. Whatever linguistic entity x is, it does what it does *synchronically* because it (synchronically) has properties (features) α, β, γ, and those features are computed over by the grammar in which x is situated in some manner, producing some 'behavior' for the entity. What features x had in your great-grandfather's brain cannot be used by you to figure out what x should do, nor, obviously, can the features x may have in the speech of your yet unborn great-great-grandchild's brain.

26. Indeed, until we posit a diachrony-free understanding of what *x* is at both Stage 1 and Stage 2 along some diachronic path, we don't know what it is that the historical account of the diachronic development of *x* is trying to explain. Since this synchronic analysis must feed into any diachronic analysis (Saussure's principle of the 'primacy of synchronic linguistics'), the diachronic analysis cannot play any fundamental role in the synchronic analysis.

27. Okay, so, if we aren't making use of these traditional arguments, why *is* historical linguistics critical to our task? Let's examine Fundamental Goal (b) first. A large part of the investigation of the syntax of Vedic Sanskrit or Attic Greek (or any other archaic Indo-European language) is, in fact, synchronic linguistic analysis. Such research is pursued under the general umbrella of 'historical linguistics' simply because it concerns historical texts. To the extent we recognize the validity of typological or structural claims made by synchronic linguists regarding the nature of syntactic structures, we must respect those claims in our analyses of the *synchronic* syntax of Vedic Sanskrit or Attic Greek (or any other human language). This is broadly acknowledged, though there is of course considerable diversity of opinion regarding what typological or structural claims are thought to be valid.

28. But there is a bit more to it than that. Given the lack of native speakers, it isn't possible to figure out some basic aspects of a dead language (meaning of individual lexical items, for example) by simple inquiry. One of the techniques for informing our theories of core aspects of these archaic languages has therefore been 'historical and comparative linguistics'. By examining subsequent (or anterior) stages of the language in which these records are written, given a sufficiently restrictive theory of change, we can develop ideas about what our ancient records mean. In addition, via the Comparative Method, we can use a reconstructed ancestor of the language under investigation as one of the 'anterior stages' for this type of reasoning. There is no 'synchronic linguistics' substitute for these methods.

29. More interesting is the question of the role of diachronic linguistics in the pursuit of Fundamental Goal (a). What can historical linguistics tell us about the human mind that synchronic linguistics cannot, or can't do as easily? To answer this question is going to require a somewhat lengthy digression on the relationship between synchronic linguistics, diachronic linguistics, and linguistic typology. Fortunately, all of the issues we need to go through to get clear on these matters are ones which we will have to talk about at some point anyway, so we might as well start out with them!

30. Imagine that the circle below represents all grammar-computational systems[3] we can imagine constructing. This includes all actual grammars, of course, but also completely absurd things like 'reverse the order of all words in the sentence if the moon is full' or 'delete final consonants if your heartrate is above 100 beats per minute'. We will call these the set of 'statable' grammars:

[3]The discussion summarizes a lengthier argument in Hale & Reiss (2008).

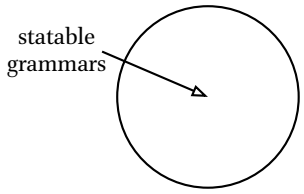

Figure 3: The Set of Statable Grammars

31. Included in this set would be the grammars mentioned above, as well as my grammar and yours, of course. In addition, however, this set includes, e.g., a grammar that is in every way just like mine, but which also has the phonological rule:

Delete vowels in prime-numbered syllables.

32. In such a grammar, we would find a derivation for 'hypercommercialization' something like:

$$həj_{1}pə˞_{2}kə_{3}mə˞_{4}ʃə_{5}laj_{6}ze_{7}ʃən_{8} \rightarrow həjpkmə˞ʃlajzʃən$$

33. Some 'Statable' Grammars:
    - Reverse the order of all words in the sentence if the moon is full.
    - Delete final consonants if your heartrate is above 100 beats per minute.
    - Delete vowels in prime-numbered syllables.

34. I cite the final, concrete and more directly linguistic example because, assuming that we all agree that no human language could exist which had such a rule in it, I need to point out that the only way build a scientific theory that will get us from the set of *statable* human grammars to the set of *possible* human grammars is to *block* the infinite number of grammars which can be easily stated, but do not appear to be, or to have been, or ever to be going to be, attested.

35. Obviously, we cannot do this by positing an overt constraint to block each impossible one; constraint 1: do not build grammars which reverse the order of all words in a sentence if the moon is full; constraint 2: do not build a grammar that deletes final consonants if your heartrate is above 100 beats per minute; constraint 3: do not build a grammar that deletes vowels in prime-numbered syllables, etc., since we will spend the rest of our lives writing such negative statements.

36. A much shorter, if not much simpler, task is to state what *can* be used to build a grammar, and allow nothing else. So, for example, if we give the grammar *no representational apparatus* that allows it to store information about the phases of the moon or one's current heartrate, the first two grammars above will be blocked (as will an infinite number of additional grammars, such as 'delete final vowels if the moon is new', 'add prothetic t- to every noun if your heartrate is below 80 beats per minute', etc. etc. etc. *ad infinitum*). And if we do not give the grammar the capacity to compute prime numbers, the third grammar (and any others that use prime numbers as a trigger for some process) will be excluded.

37. The critical point here is that it is not necessary for us to make *any statements* at all regarding 'prime numbers': by not explicitly providing the grammar with the computational capacity to generate or evaluate prime numbers, the grammar will be unable to make any use of them in its own computations.

38. The most restrictive theory we can build regarding what is a possible grammar is thus (1) a set of (representational) entities that the grammar must be manipulating, and (2) a set of processes defined over this set of entities. The grammar will not 'see' or 'know about the existence of' anything not in (1), and it won't be able to do anything to the entities in (1) except apply the processes in (2).

39. This is the theory of UG. It does not matter how you approach linguistics, contrary to popular rhetoric, unless you believe that grammars *can be* of the type seen in (33), you must believe in UG. Where do the constraints on what kinds of representation the grammar can make use of, construct, and compute over reside? It is hard to imagine any place these could be except in the mind. The properties of the human mind are determined by the properties of the human genome, just as the properties of the dog mind are determined by the properties of the dog genome, and those of the duck mind by the duck genome.

40. UG is the set of genetic constraints on possible representational entities and processes regarding grammatical computation in the human mind. Note that these constraints are specific to the grammar: human minds *can* compute primes of the sizes needed to have the phonological rule referred to above. Human minds could insert a /t-/ before every noun when the human felt cold or happy. Human languages cannot. The language system operates making relatively modest demands on higher-order cognition precisely because it cannot perform all of the computations the mind can generally. It has a *limited* set of computational capabilities, otherwise the set of natural languages we would find in the world would vary much more than they do.

41. Okay, so the necessary positing of a constrained theory of grammatical computation in humans — a UG — gives us a picture that looks like this:

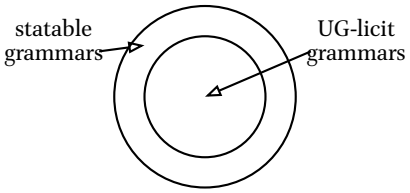

Figure 4: Statable and UG-licit Grammars

42. Of course, only data from a (presumably small) subset of the UG-licit grammars has ever been observed for scientific purposes (future grammars, e.g, as well as the vast majority of grammars of the past, have not), so we can construct a more elaborate and informative figure:

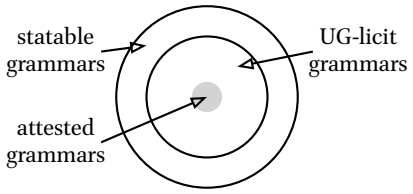

Figure 5: Statable, UG-licit, and Attested Grammars

43. For the record, it is a *tiny* subset of this small grey circle that all typological linguistics is based upon, and from which all evidence used to posit properties of UG has come.

44. Finally, we get to the point! I would like to argue that there is yet another subset of the statable grammars, as in the figure below:

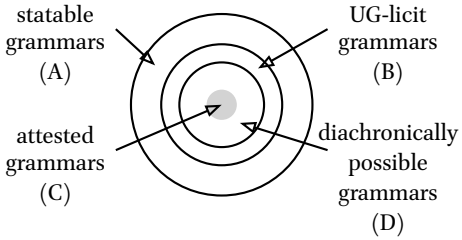

Figure 6: Statable, UG-licit, Diachronically Possible, and Attested Grammars

45. To get at this new category, and to reveal thereby the significance of diachronic linguistics for uncovering the nature of the grammatical systems found in human minds, we will consider a well-attested change, and a widely-claimed linguistic universal. The change I am interested in is **e* > *a* (as in the history of Indo-Iranian, and elsewhere). The universal I am interested in is this: a language may not have more phonetic height distinctions in the *back vowel* space than in the *front vowel* space.

46. In general, the universal is 'motivated' by phonetic considerations: given the shape of the vocal tract, the back vowel space is physically smaller than the front vowel space (i.e., you can open your mouth wider in the front than in the back). It is said to be articulatorily and perceptually inefficient to try to maintain more vowel distinctions in a tight space than in a broad one. This is the reason why the constraint is stated in phonetic, rather than phonological terms.[4]

[4]In standard modern phonological terms, /a/ is [+back], making the widely-attested five vowel system {i, e, a, o, u} a violation of the claimed universal, with three back vowels {a, o, u}, but only two front {e, i}. Since the constraint is stated over phonetic, rather than phonological, space, we can leave this issue t one side for the moment.

47. Now, imagine next that we have a language with a classic five vowel system:[5]

Front	Central	Back
i		u
e		o
	a	

48. If the well-attested change we are talking about here, that of **e* to *a*, were to befall such a system, and only that change, we would be left with:

Front	Central	Back
i		u
		o
	a	

49. Note that this system has more height distinctions in the phonetic back vowel space than in the front vowel space, thus violating the putative universal under discussion.

50. If the universal is valid, then the conclusion is obvious: the change of **e* > *a* cannot take place unless we simultaneously get rid of the **o*.[6] But how do we *prevent* the coming into being of this illicit vowel inventory? That is, given that we know that an **e* **can** become an *a*, how do we prevent its doing so in this inventory context?

51. There are at least two ways one could imagine blocking this change. The most obvious one is to say that, while it is trivial to *state* what such a system would be like, a phonological system with the vowel inventory {i, a, o, u} is *computationally intractable* for the human mind (i.e., ruled out by UG) — much like those that reference prime numbers or heart rate are. Linguistic systems with such a vowel system are thus inside set (A) in Figure 6, but outside set (B) in that figure.

52. But recall that, as I mentioned above, we don't want to exclude unattested systems by mere *stipulation*. We want to deduce their lack of attestation from the fact that (1) they involve representations that the linguistic system cannot make use of, or (2) they involve computations over licit representations, but those computations are themselves beyond the capacity of the linguistic system.

[5] I am ignoring several issues here for purposes of explication—if these issues bother you, it is relatively easy to construct your own example parallel to this one. I am intending to make no actual commitment to the position that the claimed universal about height distinctions is actually true—on this I have no specific information. I am *assuming* its validity for purposes of discussion.

[6] We could *first* get rid of the **o*, but in that case the **e* > *a* is operating on an {i, u, e, a} system rather than the envisioned {i, u, e, o, a} one.

53. Regarding (1), we have to understand that having an *o* in a phonological system is always going to be *representationally* possible — UG gives us the capacity to represent an *o*. And the UG representation system could not do any blocking, in any event, based on *phonetic* space: in my view (sometimes called 'substance free phonology') the feature system simply doesn't know anything about the *phonetic space* the things it represents occupy. So, as widely helt, is is phonetic space that is relevant, the blocking cannot be due to the representational capacity of UG.

54. But regarding the other option, (2), we need to understand that an inventory doesn't require any *computation* at all: it simply *is* what one has stored.[7] But if UG *allows* such an inventory (i.e., such systems are in set (B) in Figure 6), and yet they remain completely unattested and, by hypothesis, unattestable (i.e., they can't be in set (C) in that figure), where are we going to block them?

55. Corresponding to the more restricted articulatory space for back vowels is the simple acoustic correlate of that fact: a smaller physical perturbation of the resonating space for vowels in the mouth means a smaller acoustic difference. Thus the *structural* distance between an [e] and an [a] is precisely the same as that between [o] and [a] (MID vs. LO), but the acoustic difference is greater in the former case than in the latter.

56. This brings us to the critical question: What does it mean to say that **e* > *a*? It means that a distinction previously observed (that between *e* and *a*) was not perceived to be linguistically relevant by some acquirer (more on the general theory of change, and my assumptions about historical linguistics, in the next section). That is, the acoustic properties which distinguish MID from LO in the FRONT space were not sufficient to trigger, in the learner, a linguistic contrast. However, the acoustic properties which distinguish MID from LO in the BACK space are necessarily *even less robust* (because they are smaller, in raw acoustic terms) than those that encode the contrast in the FRONT space. So, imagine the following restriction on change events:

57. A Restriction on Change Events:

    - A learner may not pay attention to a contrast in a context in which it is *weakly* instantiated, and at the same time ignore that contrast in a context in which it is *more robustly* instantiated.

58. With this principle of learning (or 'change', which results from learning), we see why someone with the classic five-vowel system cannot lose the *e : a* contrast (which involves lots of physical space, and thus a comparatively robust acoustic signal encoding the MID : LO contrast) but maintain the *o : a* contrast (which involves a tight physical space, and thus a comparatively weak acoustic signal encoding that same contrast). Note that UG is not implicated: if we could *directly implant* linguistic systems in the brain, without having to go through the human *perceptual system*, we could implant one with words having *i, u, o,* and *a* (but none with *e*), and that human would be able to talk just fine!

[7] I leave aside here for reasons of time phonological theories which attempt to compute *surface* inventories, such as OT. These fail with respect to the issues under discussion for numerous uninteresting reasons (see Hale & Reiss 2008).

59. The diachronic filter—which says simply that there are systems which, given initial conditions and possible paths of change, cannot come into existence—thus keeps us from mistakenly attributing to the computational mind (and thus to our theory of that mind) restrictions which arise from other domains (in this case, the physics of acoustics and the nature of the raw perceptual system(s)). Understanding historical linguistics—in particular, what types of change are trivially attested, what types rare but attested, and what types completely unattested—allows us to build a model of the 'diachronic filter'. Understanding this filter is absolutely central to developing a proper characterization of UG (i.e., the human mind, with respect to linguistic computation), because it prevents us from mis-attributing to UG observed cross-linguistic regularities and absences which find their explanation in the extra-linguistic factors which feed into linguistic diachrony.

1.3. Definitional Matters

60. There are two basic, but technical, issues which we need to get clear about before we do any real work in the area of historical and comparative linguistics. It is best if we get clear about them now, before we get entangled too deeply in data-related issues. The first of these is the concept of LANGUAGE, the second that of CHANGE.

61. It will be useful to have **some** data to think about as we approach this task. In the table below you see a pretty standard 'comparative linguistics' table for a set of descendants of the Proto-Western-Micronesian language (PWMc).[8] This is at this point largely a toy example, against which we can bang our heads for bit. The data is real (largely, or probably exhaustively, taken from the Proto-Micronesian reconstructions in Bender et al. 2003), but there are many languages not in the table, and there are many issues we are not going to discuss. I have replaced the orthographic symbols sometimes used by Bender et al. with standard IPA characters to which those orthographic symbols correspond.[9]

[8]The only reason to use this subfamily of the Micronesian languages is that it leaves to one side some, for our purposes here not relevant, complications offered up by Kosraean and Kiribati.

[9]The language names are abbreviated here: CHK is Chuukese, PUL is Puluwat, WOL is Woleaian, PUA is Pulo Annian, SNS is Sonsorolese, and MRS is Marshallese. I have substitued the standard IPA length indicator (ː) for all double-written segments in Bender et al., and IPA /j/ for Bender et al.'s <y>. I have not done the substitution for Marshallese vowels, which are phonetically underspecified along the back and round dimensions. Try not to be confused by the fact that the vowels written <a>, <e>, <é> and <i> are low, mid, upper-mid, and high vowels with no backness or roundness specifications. They thus do *not* have their IPA values, there being at present no IPA symbols to represent such objects.

		— — —The Chuukic Languages— — —					Non-Chuukic
gloss	**PWMc**	**Chk**	**Pul**	**Wol**	**PuA**	**Sns**	**Mrs**
ashes	*falaŋa	fanaŋ	falaŋ	felaŋe	θanaŋa	faraŋa	—
branch	*rae	ræː	ræː	raː	laː	laː	rˠauɰ
coconut	*niu	nyː	nyː	lyː	nyː	nuː	nʲij
cry	*taŋi	—	hæŋ	taŋi	taŋi	taŋi	tʲaŋ
eyebrow	*fasu	—	faːt	faːty	θaːti	—	jatˠ
fathom	*ŋafa	ŋaːf	ŋaːf	ŋaːfa	ŋaːθa	—	ŋéjŋéj
fire	*afi	æːf	jæːf	jaːfi	jaːθi	jaːfi	—
housefly	*laŋo	nɔːŋ	lɔːŋ	laːŋo	naːŋo	raːŋo	lˠaŋʷ
land	*fanua	fəny	fany	falywe	θanya	—	janʲéj
mosquito	*ñamʷu	nɔːmʷ	nɔːmʷ	laːmʷu	naːmʷu	raːmʷu	nˠamʷ
north	*afaŋi	efeŋ	yefæŋ	yefaŋi	eθaŋi	—	jaŋ
path	*ala	aːn	jaːl	jaːle	jaːna	jaːra	jalʲ
shoulder	*afara	afar	jefar	jefare	jaθala	—	ɰajerˠ
star	*fituː	fyː	fyː	fysy	θiθi	fiti	jitʲiw
swim	*afe	—	jaf	jafe	jaθa	—	ɰaɰeɰ
weave	*faufau	fəːfə	fəːfə	fœːfœː	θauθau	—	jawjéw

62. Since full data is not being given to you, it may not be possible to justify the reconstructions as given. The full data can be found in Bender et al. 2003. Several 'regular sound changes' should jump out at you from a cursory examination of the data. For example, it seems clear that final postconsonantal vowels were lost in polysyllables in Chuukese, Puluwat and Marshallese. As does the fact that PWMc *l seems to show up as /l/ in Woleaian, Puluwat, and (with secondarily conditioned palatalization, velarization, and rounding) in Marshallese, as /r/ in Sonsoralese, and as /n/ in Chuukese and Pulo Annian. It is from data arranged in the fashion we see above that sound changes into the daughters are typically inferred.

63. This is, as I said, a very straightforward table seen in almost any work on 'comparative linguistics', showing the lexical correspondences between a set of related languages. We won't problematize today the notion of 'relatedness', but let's take a look at some issues involving the table itself, because it turns out that nothing in this 'straightforward' table is straightforward at all.

64. The first question is the following: what do the columns headed by language names represent? If Pul represents the Puluwat language in the socio-political sense, then it represents the speech of some 1200 people. For Mrs over 70,000!

65. But suppose that each and every Puluwat speaker had their own grammar, and that these grammrs could differ in phonology, morphology, syntax, or all three! Should Puluwat get 1200 columns instead of one? What does the single column *mean*?

66. And what about the column labelled PWMc (Proto-Western Micronesian)? Is a proto-language a 'language'? How many speakers does it represent? Did it have dialects, like all living languages do?

67. Turning next to the data *within* the living language columns, Bender et al. use the linguist-developed orthographic representations for the languages in the table, which is intended to represent lexical items in their *phonemic* form. As you know, the phonemic form gets converted by the grammar into a phonetic form. Is the phonemic form the correct form to place in a table like this (i.e., to make use of when reconstructing on the basis of daughter-language lexical items)?

68. And if we deduce our 'sound changes' from such tables, since the data is in phonemic form, this implies that 'sound changes' are essentially phoneme-to-phoneme mappings. Is this correct?

69. To answer these, and related, questions, we need to establish some terminological clarity in as explicit form as we can. The next few sections will attempt to do that, and explore the implications of the posited definitions for tables like the one under examination.

1.3.1. Language

70. We will adopt here the standard modern conception of LANGUAGE as a computational component of the individual human mind (I-language). Like all systems, it acquires its unique identity through (1) the set of entities over which it operates and (2) the set of processes it uses to manipulate those entities.

71. We will assume that this computational system is modular: the distinct modules of the system operate on distinct entities. The phonological system operates upon phonological, including prosodic, entities. The morphosyntax operates on morphosyntactic entities.

72. The most restrictive system is one in which the phonological system can *only* operate on phonological objects, the morphosyntactic system *only* on morphosyntactic objects. There is fairly good evidence that this restrictive system is the most productive one.

73. The entities over which the grammar computes possess properties (indeed, this is the only way entities can be distinguished from one another; two entities are different if they have different properties, the same if all of their properties are the same). We call those properties 'features'. There is no strong reason to believe, and some compelling reasons not to believe, that at this point we have correctly identified the pool of features UG makes available to the computational system. We will in general, however, not assume anything exotic in this domain in these lectures.

74. The operations of the computational system involve the modification of these features (their addition, deletion, value-changing) and their grouping (contatentation) and ordering. I will assume, though I suppose not a lot hinges on this, that no features are present in outputs that could not have also been present in inputs: i.e., the same *representational alphabet* is made use of in underlying and in output representations.

75. The grammar thus represents a mapping from one set of features, bundled and linearized in various ways, to another set of features, modified in the course of computation, and potentially bundled and linearized differently that at the initial state.

76. This, then, is what I assume a 'language' is. Of course in your own work you are not bound by any of *my* assumptions: the important matter is not whether you accept some particular set of assumptions of this type, but that your scientific reasoning be coherently built on, and consistent with, whatever assumptions you do make. This is virtually impossible to keep track of if you leave your assumptions implicit.

77. We can explore one of the issues that arise when we make an alterative assumption when we turn to a consideration of the PWMc column—what kind of entity does the data in that column represent? There have been a variety of answers (many implicit, rather than explicitly stated) to this question in the literature on the Comparative Method, and implied by the work of its practitioners. In an important methodological paper, Lichtenberk (1994) distinguishes between an 'instrumentalist' and a 'realist' view of the nature of proto-languages.

> Do we assume that at some time in the past there really was a language that had the properties that we have reconstructed (the realist view), or is such an assumption irrelevant to our concerns (the formalist view)? ... On the realist view, reconstructed proto-languages are viewed not as formal devices but as real entities, as real as the languages around us.

78. Lichtenberk's paper is an extended argument in favor of the 'realist' view, the orientation which is certainly dominant in contemporary reconstruction methodology. The implication of that view which he is most interested in exploring in that paper is the following: if the proto-language represents our best hypothesis about the properties of a 'real language', then shouldn't our reconstructed proto-languages have dialects and show the type of variability and diversity that we find when we examine the languages around us?

79. Traditionally the answer to this question was a simple, if somewhat surprising, 'no.' The justification for that answer was pretty straightforward: given two 'dialects' of some language X, the Comparative Method provides a mechanism for the reconstruction of the ancestor of those two dialects. This ancestor is the proto-language. The two 'dialects' are the daughters of that proto-language.

80. Since this can be done for any dialects of any language (including alleged dialects of any earlier 'reconstructed' stage of a set of related languages), there is no reason to treat any dialectally-divided entity as the proto-language—the proto-language would always be, by definition, the reconstructed ancestor of that dialectally-divided entity.

81. We seem, then, to be left with two pretty compelling arguments: if proto-languages are languages like those around us, then they should have dialects (as the ones around us do). On the other hand, the Comparative Method itself (the only established procedure for the creation of the hypothesis we call a 'proto-language') would need some additional mechanism that essentially said: stop applying at this point, even though there are still observed, recoverable differences between the 'dialects' you have reconstructed for which you *could* reconstruct a 'common ancester'. These would be the putative 'proto-language dialects'. At present, no such mechanism exists.

82. There are two ways this situation could resolve itself, even within the almost universally accepted 'realist' conception of things. We could amend the Comparative Method so that it would generate 'proto-language dialect variation', or we can find some justification for proto-languages seeming to differ so radically from 'observed languages'. Lichtenberk attempts the former, but I do not see that his efforts can be judged a success. See the extensive discussion of my concerns with his arguments found in Hale (2007).

83. Adopting the conception of what a 'language' is advocated above, the issue resolves itself cleanly. Individual grammars do not have dialects. Proto-languages are thus posited grammars — like all 'languages' are, and therefore cannot have 'dialects'.

84. What does all this mean for our table above? The columns headed by 'language' names must represent the data provided by *a grammar* — i.e., the linguistic system of an individual, because this is what 'languages' are. In practice, of course, many, many speakers of English have a lexical item which we could represent phonemically as /kæt/ that means something like 'cat'. It does not matter in such a case which individual grammar we take our form from. When it does matter (i.e., when the grammars differ in the form or features associated with a given morpheme) we are going to need to be careful to draw our data consistently from one of the options, rather than taking some forms from one grammar, other from from a differing one.

85. This has a potentially convenient and desirable corollary, in my view. If the data in the table represents phonemes, for example, then we should recognize that there is only one place in the world where we find such objects: in the minds of humans. The 'English language', rumors to the contrary notwithstanding, does not have a 'phoneme inventory'. Nor does 'the Chuukese language', 'the French language', or any of those other socio-political notions with which you may be familiar. If we want to have phonemes in our table, the columns need to represent systems that are found in human minds, because only these systems possess *phonemes*.

86. It follows, of course, that if you want to use some other definition of 'language' in your historical linguistic work, your table better contain things that exist for the entities created by your definition. I don't myself see how a non-I-language definition could allow the 'languages' in the table to have things like phonemes, morphemes, syntactic features, etc., but then maybe I just lack imagination.

1.3.2. Change

87. If we are going to investigate 'language change', we presumably need 'change' in the relevant sense to be defined over 'languages' in the relevant sense, i.e., we are going to want 'change' to be a relationship between grammars.

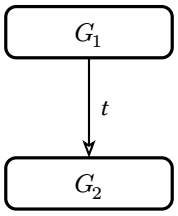

88. Clearly, the grammars need to stand in some kind of relevant relationship for us to say that that relationship is mediated by 'change' events. But not just any kind of relationship: Samoan *tapu* and English *taboo* are historically connected in some sense, but it is not *that kind* of relationship we are trying to capture with the concept of 'change' (as opposed to 'borrowing'). What we would like is to require that G_2 be *descended* from G_1 in some meaningful sense. What might a useful construction of that sense be?

89. At the same time, by the way, it would be very nice if we had some understanding of why it is that change happens at all. Why don't we all just have the same grammar as our linguistic ancestors? What makes change *inevitable*?

90. There is a healthy amount of evidence that one of the major sources of regular diachronic change of the type we are interested is *nonconvergent acquisition*, i.e., the failure of the acquirer to converge on the same knowledge state (grammar) as his/her sources.[10] Using some standard symbols, we can call the initial state of the learner S_0 (=UG for child acquisition). The child moves out of this state on the basis of some linguistic data generated by the person making use of the target system (we will call the target G_1). This data provides the acquirer with information about how they should modify their current knowledge state to get something closer to that of their target.

91. Let us call the first post-UG knowledge state constructed by the acquirer S_1, the second S_2, the third S_3, etc. In order for these knowledge states to be different from the anterior knowledge state, the acquirer must get some data which causes them to revise their current hypothesis (S_x), positing the new hypothesis (S_{x+1}).

[10] Let us make, provisionally, the assumption that the acquirer has a single target. You may be nervous about the fact that acquirers often seem to talk to a lot of different people, presumably each with their own grammar, and thus cannot be meaningfully said to have 'a source'. This is, I think, a non-issue. If exposed to systems which are producing distinct output, the acquirer will not learn a mixture of the two systems, but rather both. Ditto for four or five systems. They may ultimately continue to use into adulthood only a subset of these, but I don't think that matter is of any particular relevance. The *acquirer* does not stand in a descent relationship to the earlier grammar, one of their grammars does.

92. It is of some interest to remind ourselves at this point in the discussion that we do not call it *change* when we talk about the relationship between a target's knowledge state (G_1) and that of their 24-month old daughter. 'Adult English' doesn't change into '24-month old English', only to later change into the next stage of 'adult English', once the acquirer matures.

93. But there must be some point at which differences between the linguistic system of the target and that of the acquirer *do* count as change, because otherwise, there would be no change! Let us imagine, then, that there exist some point at which it becomes meaningful to say that grammatical acquisition is complete and that any differences which persist at that point do count as 'change'.

94. The 'stages' before that count as steps in the acquisition of the grammar, but not as grammars in the relevant sense, and that is why their relationship to their target is not *change*—change is a relationship between grammars.

95. The entity we have at the point where it becomes meaningful to call differences 'change' is the new grammar, G_2. The 'changes' that have taken place under this scenario are thus the differences between G_1 and G_2. Here it is graphically:

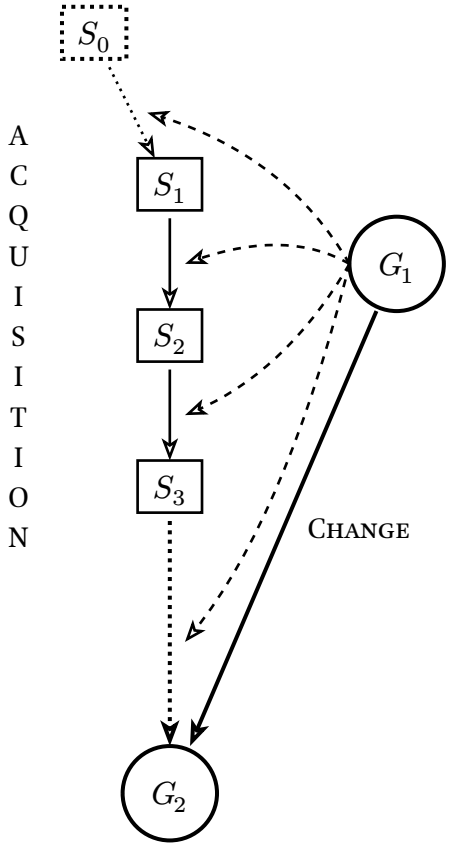

96. This conception of things has several interesting implications for how we conceive of change events. Take the following claim by a prominent diachronic syntax scholar. Lightfoot (1991:167) says that 'changes involving new parameter settings tend to take place much more rapidly than other changes…' The sociolinguistically-oriented approach to syntactic change (e.g., in the work of Tony Kroch) often asserts that it is studying 'change in progress', where the statistical development of a change from its beginning through its completion, often hundreds of years later, is graphed in a familiar S-curve shape. Can either of these ways of understanding 'change' be maintained, under our assumptions?

97. They cannot. The differences between G_1 and G_2 are the 'changes' which took place.[11] They did not exist before G_2 existed. Once G_2 existed, they did. Since everything before G_2 is an S-state, over which change is not defined, nothing that happened before then is 'change'.

98. G_2 is the point at which a 'grammar' has actually come into being (and thus 'change' could be defined). I assume that G_2 doesn't later change to G_3 then G_4 (although it may spawn construction of a new grammar in another speaker, which would appropriately bear those indices). The G_2 speaker can presumably learn additional languages (or dialects, we are of course just talking about grammars) after G_2 comes into being, but these are not *developments* of G_2—they are independent creations, which will be posited by the learner *alongside* G_2.

99. So what are David Lightfoot and Tony Kroch talking about? They are talking about the *spread* of the change through a population. Spread of course always takes time (it could be defined as diffusion across space in correlation with time). The spread of tea in a cup of hot water follows Kroch's S-curve. Let us call this spread 'diffusion', to keep it distinct from change proper.

100. Another nice (in my view) implication of this conception of change is the fact that we can see fairly readily why change is *inevitable*. The acquirer has no direct access to G_1, but is only exposed to noisy indirect data regarding the properties of G_1, presented in a non-systematic manner to his or her parser. From this incomplete and degenerate input, the acquirer is able to *approximate* G_1, but not perfectly replicate it. Any failure to replicate it will represent difference, and difference that persists to G_2 will be change.

101. One other important point needs to be established before we move on. Imagine a series of acquisition events, as in the figure below:

[11] I will continue to include by the general term 'the grammar' both the computational system and the computationally relevant lexical features made available by the lexicon. The 'encyclopedic lexicon'—open class words such as nouns, adjectives, verbs, etc., can be acquired throughout the lifetime of the speaker, and thus are *not* included in the 'termination point' G_2.

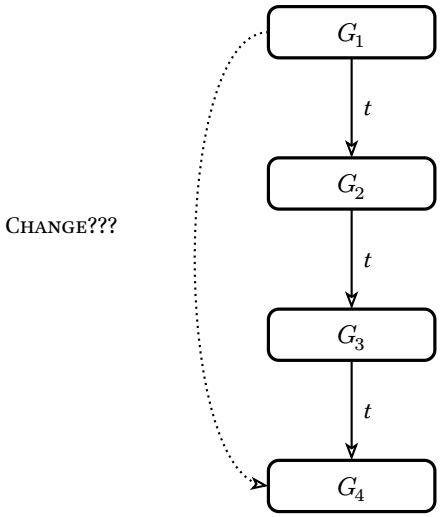

102. Is the relationship between G_1 and G_4 to be characterized as *change*? Imagine, to be concrete, that G_1 had an [p^h] in some word and G_4 had an [h] (and that we don't know what G_2 and G_3 had). Can we say that p^h 'changed into' h?

103. This question is well known in the traditional historical linguistics literature, usually posed of more 'dramatic' differences such as IE **dw-* and Armenian *jerək-*.[12] Can **dw-* 'change into' *jerək-*?

104. Remember that we posit concepts, such as *change*, to assist us in the scientific pursuit of trying to model the natural world. The definition of *change* we should adopt is not given by the facts: instead, the set of phenomena which our theory of what is and what is not a possible change event is going to need to account for is given by our definition. When we ask what we want, in our model, to label 'change' we are asking what definition is most likely to prove productive in developing a scientific account of the world.

105. It seems clear, given those considerations, that it is not useful to call the differences between G_1 and G_4 (or, for that matter, between G_1 and G_{400}) 'change'. While the dynamic factors which play into the differences between G_1 and G_2 can be made coherent within the context of a restrictive theory of *language acquisition* (given input X, Y, Z what kind of grammar does the learner construct?), the properties of G_4 bear no such relationship to those of G_1. If there are phenomena which our theory of language acquisition preclude, we can restrict the relationship between G_1 and G_2, and thus, if we define 'change' over such a domain, we can restrict the set of possible changes (given a set of initial conditions). But such a restriction will not limit the properties of G_4 given those of G_1.

[12] Cited cases include **dwoHu* > *erku* 'two' and **dweHros* > *erkar* 'long'. See Hock (1991: 583-9) for an informal discussion.

106. The only limits G_1 imposes on G_4 arise from an understanding of the restrictions on the relationship between G_1 and G_2, between G_2 and G_3, and between G_3 and G_4. That is, the set of a possible differences between G_1 and G_4 is a direct function of the set of possible 'changes' (in the restrictive, single-generation direct inheritance sense) between each of the steps between G_1 and G_4. These we get from our definition of *change*. No *additional* restrictions beyond these hold of the differences between G_1 and G_4, thus adding this relationship to the set considered for a definition of *change* merely weakens our capacity to make a restrictive theory of that phenomenon.

107. In actual practice we will not have been present to directly record any change event, in the sense required by our definition of 'change'. I do not see this as a problem: it is rather a normal condition of scientific research. The constrained definition of change doesn't influence what kind of data we get, but only the types of explanations we can offer for the observed data.

108. For example, in the case of IE *dw- to Armenian *jərək-*, which stand in correspondence in some relevant sense, we can ask the question of whether their relationship is regulated by our concept of 'change'. In my view, it is not and could never be (this is an *impossible* change). If the elements stand in correspondence, we can deduce that there must be *intermediate stages* in the development of this correspondence. We can then formulate a hypothesis about what those intermediate stages might be. This is standard practice in historical linguistics, of course, but it is often undertaken without clarity about how it relates to a coherent notion of 'change'.

109. For example, we should note that under this kind of definition, some issues which figure prominently in the traditional historical linguistics literature become quite irrelevant. For example, is change 'gradual' or 'abrupt'? Well, since it doesn't exist until the moment G_2 comes into being, and does exist at precisely that moment, it is always 'abrupt'. There can be no 'gradual' chain shift events (such as the Great English Vowel shift), taking several generations to complete. These are simply sets of independent 'change' events.

110. For a more recent theoretically-sophisticated diachronic analysis that seems to run afoul of this conception of things, below is a quote regarding a syntactic change in word order in the history of English from Roberts (2021: 173). The details of the change and the roles of the factors adduced need not bother us in this context: I am interested only in how the concept 'change' plays out:

> As they [Taylor and Pintzuk 2015: 238-9] point out, FOFC [the Final-over-Final-Condition—mrh] dictates that the change in TP is a precondition for the change in VP, but that this does not mean that the first change had to be completed before the second change began.

111. Here and in many places in Roberts (2021) we find that 'change' is used to cover some notion other than that being advocated here. Of course in many instances in the book, we also find the 'change' concept precisely as advocated here, explicitly so on pp. 426-428. On those pages Roberts uses the term (attributed to me by Lightfoot, apparently) 'mirage of gradualness'. He then gives some reasons why this mirage 'is very apparent to us in our observation of the phenomena'. But in the quote above (and in numerous places where the time course of change plays a role in that book) the linguist's *mirage* is being causally invoked in the explanation of the change event in the natural world — and this, I fear, is not possible.

112. When we, as scientists, recognize explicitly how our pretheoretical, everyday experience of the world interferes with a scientifically useful conception of what we are seeing by inducing mirages, it is very useful to be as clear as possible about how easy it is to fall victim to them — as Roberts is. But the goal of being sensitive to the presence of these induced mirages is to eliminate them from our formal analyses, not incorporate them!

113. The implications of this conception of 'change' for our Comparative Western Micronesian table play out importantly not in the table itself, which is a static picture, but in the process whereby we infer sound change events from that table. The changes themselves must be *minimal*, in the sense that they are consistent with single-generation non-convergent acquisition events (possibly a chain of such things). To see how this plays out we need to deal with a few more definitional matters!

1.3.3. Phonetics and Phonology

114. One of the questions I raised about our table concerns the status of the lexical data within it. It is, as I mentioned, alleged to be in phonemic form, but should it be?

115. Before we can really answer that question, we are going to need to understand what I mean by 'phonemic' or 'underlying' representation, and what I mean by 'phonetic' or 'output' representation. It is possible, perhaps even likely, that my assumptions about these matters differs from your own. As I mentioned above, it is not important whether you accept my own (obviously correct) assumptions about these matters or persist in using your own more delusional ones — what is important is that whatever work you do be consistent with the assumptions you make use of. To state the matter less jocularly, I think most work in historical linguistics, regardless of what its authors say, makes assumptions like those I will propose here, because I do not think that competing assumptions converge on a coherent basis for analyses. But again, this may just be my lack of imagination.

116. The phonological computation operates over a set of stored representations. These representations are structured feature sets (bundled, for example, sequenced, and perhaps hierarchically configured).[13] These representations are the *phonemic* or *underlying* or *input* representations.

[13] I'll assume — though as I mentioned before, in this case almost certainly counterfactually — that something like our current feature set is the correct one.

117. I'll assume serial, rule-based computation over those features represents the 'phonological derivation'. Nothing hinges on this derivation being serial (though I know of no coherent phonological theory that lacks serialism) or rule-based.

118. These computations modify the values, presence/absence, bundling, hierarchy, and sequencing of these features. When the derivation has completed it will have produced a representation. This representation is in the same *representational alphabet* (set of phonological features) as input representations are. This is the *phonetic* or *surface* or *output* representation.

119. Such representations can be passed on to other computational and eventually electromechanical systems of the human body so as to produce articulating behavior and, under appropriate conditions, an acoustic stream. The process of getting from the phonetic representation to this behavior/acoustic output involves converting the output representation into some forms which are not in the same representational alphabet as we started with (e.g., muscle commands) — such a conversion of representational form I will call, following Hale & Reiss 2008, *transduction.*

120. I will assume that the phonology (and the grammar) cease to play a role in this process once it has generated and passed on the output representation to other systems. It does not interfere with the transduction, nor with the physical motions of the articulators, nor with the acoustic stream as it emerges from the body's interaction with its physical context. These latter processes can perturb and in part obscure for a perceiver the properties of the output representation.

121. A direct implication of these assumptions is that there is no such thing as 'language specific phonetics'. A 'language' is a grammar, and the grammar ceases playing a role once the output representation is computed. Every aspect of the computation of that representation that could affect its implementation *linguistically* is phonology.

122. I say *linguistically* because of course the properties of your body definitely affect the acoustic realization of the phonetic representation, as do matters such as speech rate, emotion, etc. There is a great deal of evidence that such factors are non-linguistic: e.g., you do not pronounce the words you have stored with the acoustic properties of the voice of the person from whom you learned that word, nor at some particular speech rate that was in use when you learned the word. You store a representation that abstracts away from these acoustic properties of the input stream. A linguistic representation.

123. So, this is what I am asking about when I ask whether the data in our table is or should be phonemic or phonetic: should that data consist of the long-term stored representations that form the input to phonological computation, or the generated output representations that represent the output of phonological computation?

124. It is important to note that in the post-1950's world it is traditional to list the daughter language forms in such a table, and in general discussions of sound change, in their *taxonomic phonemic* form. Indeed, in modern discussions of the Comparative Method, argumentation is explicitly offered as to why this is a good idea, or perhaps even necessary. In a book on the Comparative Method, for example, Fox (1995) writes the following:

> … consider what is meant by "sound change." In purely phonetic terms this can only mean that a particular articulatory gesture is replaced by another, whether in all cases where it occurs in the pronunciation of the words of the language, or just under certain conditions. From the point of view of structuralist phonology, however, this change must be evaluated differently according to whether or not it results in a change in the phonemic status of the sounds involved. If it merely affects how the allophones of a phoneme are articulated, then the change is, phonologically speaking, superficial, and indeed irrelevant; only if it affects the system of phonemes itself, or the grouping of sounds into phonemes, can it be considered a genuinely *linguistic* change. (Fox 1995:38)

125. He goes on to explicitly discuss the significance of this observation regarding change for the Comparative Method:

> What are the implications of this for the Comparative Method and for the reconstruction of protolanguages? First, it means that since phonemes, rather than sounds, are regarded as the basic units of pronunciation, our methods of reconstruction must be geared towards determining the phonemes of the protolanguage rather than its sounds. In fact, determination of the latter is regarded as rather unimportant… (Fox 1995:42)

126. We always want to be careful when we're told that some systematic data about the world is 'unimportant', 'superficial', or 'irrelevant' without careful explication as to why that might be the case. I'll argue below that phonetic information has none of those properties, and is, in fact, critical to work that we do in comparative linguistics.

127. But let me note that in a workshop on *theoretical historical linguistics* arguing that phonetics is important is a little strange. Didn't we just argue in some detail that we are interested in *grammars*, which are long-term stored aspects of linguistic knowledge. Aren't phonetic representations fleeting, transitory constructs? If 'change' is a difference between grammars, what do phonetic representations have to do with it? And why are they, in my view, important?

128. One way to approach these questions is to ask what difference it might make. That is, how would our table be different if we used *phonetic* forms, rather than phonemic. It is particularly interesting to ask this question when it comes to data that could be relevant to our understanding the relationships between the languages, and to our understanding of the relevant sound changes.

1.3.4. Five Cases

129. There are five cases arising out of our PWMc table that I'd like to discuss: recall that the original table was created by taking the phonemic representations provided in a mix of *orthographies* and putting them in corresponding IPA symbols. No detailed attention was paid to the actual phonetics of the segments in the individual languages — in general none is provided in Bender et al., so scholars making use of their extremely valuable collection of data, will not have access to any such information when interpreting the table.

130. Our goal here is not to walk through what the optimal reconstruction and phonological history of the Western Micronesian languages actually is: that would entail a consideration of *all* available data and a much more extensive discussion, which would teach you — at great pains on your part — about something you probably don't care very much about. The *methodological* point is just to show you that it is often quite clear that the use of phonetic, rather than phonemic, data will influence our conception of the phonological history of the languages under discussion, and in a positive way.

131. The 'correspondence sets' to be discussed are given in the table below. They concern (1) the development of PWMc *l and *n, (2) the development of PWMc *r, (3) the development of PWMc final vowels, and (4) the development of PWMc *f. We will not worry about the Marshallese-internal palatalization and velarization (and rounding, which has not come up in our data) of *l and *n by adjacent vowels: Proto-Marshallese can, I think, be safely assumed to have had simple *l and *n for these segments.[14]

PWMc	Chk	Pul	Wol	PuA	Sns	Mrs
*l	n	l	l	n	r	$l^{ɤ}$,l^{j}
*n	n	n	l	n	n	$n^{ɤ}$,n^{j}
*r	r	r	r	l	l	$r^{ɤ}$
*VCV#	VC#	VC#	VCV#	VCV#	VCV#	VC#
*f	f	f	f	θ	f	j

132. Let's start with the developments of PWMc *l and *n. As you can see from the correspondence sets, several languages show only a single outcome for these segments: Chk and PuA have only /n/ and Wol has only /l/. In the languages that distinguish between the two proto-segments, all show /n/ as the outcome of PWMc *n and either /r/ or /l/ as the outcome of PWMc *l. Interchange between *l* and *n* is widespread in Oceania (and elsewhere, as is interchange between *l* and *r*.)

[14] By contrast, the development of *r to Mrs $r^{ɤ}$ is an unconditioned, regular development.

133. The merger of *l and *n in half of the languages surveyed is perhaps a bit surprising. The change feels non-trivial, but it clearly cross-cuts the standard classification of these languages into a 'Chuukic' subgroup (all the languages except Marshallese) and a non-Chuukic group (only Marshallese, in the data provided, but there are several other languages in this group). It is worth noting, though I have not presented you with any relevant data, that *ll (i.e., geminate *l*) shows up throughout the Chuukic group as /nn/ (including in Woleaian and Sonsorolese).

134. The 'phonemic' data given in Bender et al. would lead one to believe that there is a shared innovation of PWMc *l into /n/ in Chuukese and Pulo Annian.

135. But here's Goodenough & Sugita's (1980: xiv) description of Chuukese /n/: 'a voiced, dental or alveolar, nasal continuant tending, unless double, to denasalize between vowels, where it sounds more like an alveolar flap'. The intervocalic denasalization also takes places in 'close sandhi' connection with preceding vowel-final or following (underlyingly) vowel-initial words.

136. And here's Oda's (1977: 12) description of Pulo Annian /n/: '[the distinction] between a geminate and a single *n* is a matter of nasal vs. flap'. She posits a rule (1977: 16) that says '[t]he alveolar nasal *n* is denasalized and becomes a flap [ɾ]... when it occurs singly.' When geminate, it remains a nasal.

137. And finally, here's Sohn's (1975: 14) description of Woleaian /l/:

> Woleiain *l* is a **flap** sound more or less like the *r* in British *very*.

138. Sohn also notes (1975: 15) that when *l* would be geminated by phonological processes in Woleaian, it surfaces as [nn].

139. Given this information, the phonetic data for the *l and *n data sets looks rather like this (I've adduced the geminate data for completeness sake):

PWMc	Chk	Pul	Wol	PuA	Sns	Mrs
*l	ɾ ~ n	l	ɾ	ɾ	r	lˠ,lʲ
*ll	nn	nn	nn	nn	nn	lˠlˠ,lʲlʲ
*n	ɾ ~ n	n	ɾ	ɾ	n	nˠ,nʲ
*nn	nn	nn	nn	nn	nn	nˠnˠ,nʲnʲ

140. It is quite easy to see that the use of phonetic data has changed fundamentally the relationships between the languages in the table. In particular, the connection between Chuukese, Woleaian and Pulo Annian, originally showing an *n* : *l* : *n* correspondence, now looks much closer (*ɾ* ~ *n* : *ɾ* : *ɾ*) and seems to invite an inference where the Chuukese allophonic alternation may be an intermediate step between the PWMc situation and Woleaian and Pulo Annian. Whatever the ultimate analysis (which would require phonetic data for the other Chuukic languages), that the phonetic data will play a key role in determining a proper diachronic understanding is, I think, clear. The phonemic data, as presented in Bender et al., misleads the non-specialist in critical ways.

141. Regard the data for PWMc *r, I am interested in the Marshallese form. As I hinted at above, for *l and *n we get three outcomes (only two show up in our table, so I only mentioned two so far) based on the backness and roundness of adjacent vowels: a palatalized ('front') outcome (l^j, n^j), a velarized ('back') outcome ($l^ɤ$, $n^ɤ$), and a velarized and round outcome (l^w, n^w). Marshallese also has the three segments *r^j, r^ɤ,* and *r^w*. Interestingly, however, the palatalized Marshallese segment *r^j* arises from the PWMc palatal stop **c* only, regardless of adjacent vowel qualities for backness and roundness. And the round and non-round velarized rhotics arise only from PWMc **r*, which never palatalizes (but shows up round in the context of round vowels). Why are the Marshallese rhotics so different in their history from the lateral and dental nasal?

142. That the PWMc palatal stop **c* might show up as an invariably palatalized segment in Marshallese is not terribly surprising, I suppose. But from where do we get the invariant velarization of the descendants of PWMc **r*?

143. It looks from the phonemic data that this is purely Marshallese problem, but when we consider the phonetic data more carefully, the problem becomes more interesting.

144. Sonsorolese shows /l/ as its reflex of PWMc **r*. There is nothing unusual in this: interchange of *l* and *r* is common throughout Oceania (and elsewhere). But here's Quackenbush's description (1968: 47) of Sonsorolese /l/: 'a voiced laterally-released prevelar stop somewhat like the *gl* in *aglow*.'

145. Here's Capell's (1969: 17) description of that same segment: '[t]he *l* is very peculiar… [i]t is at once palatal, like the *ll* in Continental Spanish, and velar.'

146. That is, there is a strong velar component to the obviously somewhat strange Sonsorolese phoneme /l/. Given this, our task in trying to figure out why Marshallese has the change of **r* to *r^ɤ* may look very different than the table implies: if the Sonsorolese velarization is to be connected to that in Marshallese, which seems plausible if not absolutely required, the important additional phonetic information about the proto-language segment may make our diachronic story for Marshallese much more straightforward. What should at any rate be clear is that there is *potentially relevant* evidence that is being obscured by the fact that the table contains only phonemic data.

147. Turning next to the word-final vowel data we can see that final postconsonantal vowels in polysyllables were lost everywhere except Woleaian, Pulo Annian, and Sonsorolese. This includes both the non-Chuukic language Marshallese (and several other non-Chuukic languages we have not considered here) and the Chuukic languages Chuukese and Puluwat (and several other Chuukese languages not considered here). Given the subgrouping, which I think is very likely to be correct and is supported by a wealth of other evidence, we need to accept that the loss of final vowels is a relatively trivial innovation, since Marshallese shows what must be a 'parallel independent development' relative to the loss in Chuukese and Puluwat.

148. None of the language show a reduction of unstressed vowels to schwa, so it is of some interest to ask how trivial the loss of non-reduced final vowels actually is. Our account would be helped if there was something about the final vowels in Woleaian, Pulo Annian, and Sonsorolese that made that parallel independent development more trivial.

149. And there is. In Woleaian, Pulo Annian, and Sonsorolese final short vowels in polysyllables are all *phonetically voiceless*. That is, the forms of the word for 'cry' in those three languages is phonemically /taŋi/, as Bender et al. give it, but *phonetically* [taŋi̥]. And the phonetic forms of the 'fly' word, given phonemically as /laːŋo/, /naːŋo/, and /raːŋo/, are phonetically [ɾaːŋo̥], [ɾaːŋo̥], and [raːŋo̥].

150. It is easy to see that the parallel and independent loss of final voiceless vowels is going to be a far more trivial development than is such a loss for unreduced, voiced final vowels. All of the non-Chuukic Western Micronesian languages have lost these final vowels, so we can clearly reconstructed 'cry' not as **taŋi*, as everyone does, but as **taŋi̥*. But only if we reconstruct not *phonemically* (because the voicelessness of vowels in this context is predictable, even in the proto-language, and thus non-phonemic), but phonetically.

151. Let me turn finally to the development of PWMc **f* in Pulo Annian. The reflex of this segment is given as θ and the discussion of Oda (1977) is pretty non-specific about its actual phonetics. This development has recently captured the interest of Patrick Honeybone (2016), in his consideration of the question of whether **f* > θ represent a possible or an impossible sound change.

152. I won't have time to consider all aspects of Honeybone's interesting paper, but let me limit myself for the moment to his Pulo Annian discussion. Ultimately, to avoid Pulo Annian being an exception to the claim that **f* > θ is not a possible sound change, Patrick makes two suggestions. The first is that in some sense the θ of Pulo Annian isn't *really* a voiceless interdental fricative, but rather an 'unvoiced alveolar fricative with a flat or "slit") cross-sectional tongue shape.' He transcribes this segment as θ̪, and concludes: 'Given all this, it seems that the change that caused us to consider this case is not, in fact, of the type f > θ. Rather, it seems to be f > θ̪.'

153. I am a little uncomfortable, as much as I like phonetic detail, with side-stepping this data with such subtle phonetic argumentation in the absence of any access to speakers or even recordings. And it appears that the potentially still living 10 speakers have abandoned Pulo Anni and disappeared into nearby villages. One question that the analysis certainly gives rise to is by what argument we can safely conclude that if f > θ̪ **is** a possible change, f > θ should not be. But Patrick presents an alternative diachrony that for the most part avoids these issues.

154. The second, and ultimately adopted, proposal gives a particular relative chronology for the development from Proto-Chuukic to Pulo Annian that involves a sound change which merged *s and *f as pre-PuA *t. This is a change whereby (Honeybone 2016: 342) 'the two PCk fricatives … merged with the least marked PCk obstruent (i.e. t), retaining their laryngeal state.'

155. The ordered steps proposed are specifically: (1) t > d before *a*, (2) s, f > t, (3) c > s, (4) t > θ̲, (5) d >t. It is apparent that there is no change of *f to θ in this scenario, though there is the 'single-step' change of *f to t, which I personally find just as uncomfortable and thus just as likely, if not moreso, to be an impossible change.

156. Here are the correspondence sets for the descendants of Proto-Chuukic from Bender et al. regarding the segments that are implicated in Patrick's scenario:[15]

PWMc	PCk	Mrt	Chk	Pul	Crn	Crl	Stw	Wol	PuA	Sns	Uli
*f	*f	f	f	f	f	f	f	f	θ	f	f
*t before a	*t	s	s	h	h	s	s	t	t	t	t
*t elsewhere	*t	∅	∅	∅	∅	∅	∅	s	θ	ð	s
*T	*s	s	s	h	h	s	s	s	θ	ð	s
*s	*d	t	t	t	t	t	t	t	t	t	θ
*c	*c	š	č	r̥	r̥	š	r̥	š	s	s	c

157. Not surprisingly given the table, the developments here are quite a mess (see Sohn et al. 1977 for a very unsatisfying discussion, which, however, reveals how tangled the affair is), and better phonetic data could be useful.[16] But these developments are some of the clearest for mapping out the internal structure of the Chuukese languages, and they show that Pulo Annian and Sonsorolese are pretty closely related, as I think everyone has always suspected (probably on geographical terms alone). With the exception of the treatment of *f and possibly the voicing of the interdental fricatve they show the same developments. No one else has an interdental fricative for *t before non-low vowels or for PCk *s. They probably subgroup of Woleaian, the three languages being the only ones which agree in the total merger of PWMc *t before *a* and PWMc *s (PCk *d). The three also agree in the merger of PWMc *t not before a low vowel with PWmc *T (PCk *s), though here Ulithian comes along for the ride.

158. Against this data we can evaluate the chronology of the sound changes posited by Honeybone (2016). The first change involves a merger of PWMc *t before a with PWMc *s (PCk *d). We can see that this change can only have happened in the history of Woleaian, Pulo Annian, and Sonsorolese, all other Chuukic languages keeping their reflexes of these two segments distinct. So Honeybone's first change tells us all subsequent developments concern this subgroup.

[15] I have followed Patrick in using *s for the segment Bender et al. give as *T at the Proto-Chuukic level. I have also inserted θ for the orthographic <d> Bender et al. give for Pulo Annian. Jackson (1983: 17) lists the IPA for this segment as θ. I'll also ignore here the 'sporadic' ∅ reflexes of PCk *t in all Chuukic languages.

[16] The reconstructions for Proto-Chuukic by Bender et al. are pretty poor, in my view. Every language treats PWMc *t before *a* differently than PWMc *t before non-low vowels, so these segments must be reconstructed as distinct in Proto-Chuukese. Of course, the two Proto-Chuukese segments may have been allophones, but I think we've seen enough evidence by now to realize that that does not matter. The reconstruction of *d, if taken literally, for a correspondence set that doesn't include any *d* is not possible.

159. Honeybone's second sound change merges PWMc *T (PCk *s) and PWmc *f with the instances of PWmc *t not before *a*. Concerning PWMc *T the Pulo Annian merger with *t not before *a* finds perfect agreement in Woleaian and Sonsorolese. The problem is that those languages do *not* include *f in that merger, indicating that that development of *f is a separate change. This is not necessarily fatal to Honeybone's scenario: after all, he only needs *f to *eventually* merge with the merger product of *T and *t before non-low vowels — it need not be a single-step merger as he has depicted it. The motivation he stated for the change will be gone, be could still apply to a later shift of *f to *t.

160. Honeybone's third change, of c to *s* is found only in Pulo Annian and Sonsorolese. Since it doesn't interact at all with the f to t change Honeybone still needs, the hypothesis still stands. However, the next change is the t to θ one, and now we can't wait for f > t. So we insert (4a) f > t, (4b) t > θ.

161. But now, since (4a) is Pulo Annian only, (4b) can no longer get our Sonsorolese developments. We will need an independent change of t > θ in that language. This is still not fatal, of course, but it is undesirable.

162. Finally, Honeybone's fifth sound change is d > t. We have had to hold on to this development for a long time, because we needed to merge *f into our *t* segment for change (4b). This change, however, is found in every Chuukese language except Ulithian. Indeed, if Ulithian θ arises via an independent change of *t* to θ, which could follow this development, the d > t change probably dates to Proto-Chuukic. Of course it could be an independent development in some or all of the languages — it would have to be one in Pulo Annian if Honeybone's chronology is to work. Again, this seems very cumbersome.

163. I'd like to propose what I think is a significantly simpler method for getting us what we want, leveraging the fact that we would like to insert as much phonetic detail as possible into our 'correspondence' table. Patrick attacked this problem at the θ end, but I would like to attack it at its root, with the *f.

164. Bender et al. are quite clear about where in Proto-Oceanic our *f comes from. It arises from POc *p. But not from any POc *p. When POc *p sat before the back vowels *i and *u it was lost.[17] So all descendants of POc *p that feed into PMc *f (and thus PWMc *f and PCk *f) were before potentially palatalizing vowels.

165. If we look at the outcomes of POc *mp we see a split on the way to PMc: before round vowels this gave the labiovelar *pˣ,[18] before non-round vowels it gives PMc *p. That *p, however, is reflected in Marshallese without conditioning as *pʲ That is, it seems not implausible that the prenasalized labial stop developed into a velarized stop before round vowels and a palatalized one before non-round vowels — a sensible enough development. Precisely this has been proposed by Lynch (2003) for Micronesian languages, with excellent evidence that it happened in other Oceanic branches as well.

[17] In fact, I think it may have developed into a glide, but I'll leave that matter to one side.

[18] Often written *pʷ in the Oceanic linguistics literature.

166. Imagine the development of non-prenasalized *p was parallel, but that the *fʷ which appeared before round vowels was lost (or, more likely in my view, became a glide). It follows that the segment that we are reconstructing as *f was in fact that palatalized sister of this segment: *fʲ.

167. The dentalization of palatalized labials is a well-known and well-established phonological development (see, e.g., Ohala 1978). It is explicitly invoked by Lynch (2003) as part of his justification for reconstructing the Proto-Oceanic 'plain' labials as having been, in fact, phonetically palatalized. The dentalization of *fʲ gives us exactly what we need: a θ.

168. So the Pulo Annian data tells us that *fʲ > θ is a possible sound change. It tells us nothing about whether a non-palatalized, plain *f > θ is. As Honeybone also concluded, for different reasons.

169. If the linguists who created the correspondence set Honeybone was working from, and had worked out the historical phonology of the Micronesian languages, had done so including as much phonetic information as was available to them, the fact that it was a palatalized *fʲ that had developed into Pulo Annian θ would have been apparent to us all at the outset. And it would have made the explanation for the change event clear from the get-go.

170. And precisely why it is helpful to include phonetic detail in the table is completely clear: the phonemicization **hides** linguistically-relevant information from us. The generation of linguistic output forms — phonetic representations — is what the grammar does. The grammar constructed by the acquirer must include the machinery to generate such forms, and that machinery is thus part of the 'knowledge of language' the grammar is supposed to represent.

171. Where does the phonemicization process 'hide' knowledge about how to generate phonetic output representations? In the rule system, of course. Phonetic representations differ from phonemic representations *only* in that they have been subjected to the modifications imposed by the phonological system. That system is learned, and transmitted during the acquisition process, and thus *linguistic* (Fox notwithstanding). But it is completely absent from our 'correspondence tables': that has only forms, no rules. Should we have a similar table with rules in it instead of word forms?

172. I think not, but let's first talk for a moment about the relationship between the various components of the phonological system and the history of our research on change in the following way:

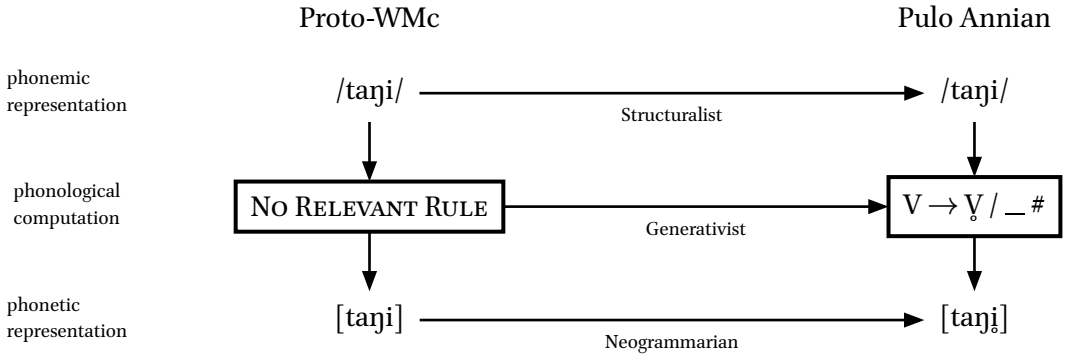

173. The Neogrammarians had no concept of the 'phoneme', nor of phonological computation. As such, their analyses were dependent upon the only level of representation they posited: the surface phonetic one.[19]

174. With the rise of structuralism, the phoneme-centric analyses that we are most familiar with became the norm in historical linguistics as well (though it took time, as theoretical innovations usually do, to be fully embraced). The phonetic information was stripped from the forms under discussion, in my view, as should now be apparent, much our detriment.

175. With the generative revolution the recognition of the importance of the rule system became clear. Oddly, this development was associated with simultaneous reduction in the significance of both phonemic and phonetic representations. Language change was 'rule addition, rule loss, rule reordering' and the like. Phonemic representations were held to be 'remarkably stable' over long spans of time (think of the SPE era underlying representation for English 'filth': /fūliθ/). Because they basically did not change (though of course they must have changed *sometimes*!), they were not important. Because the phonetic representations were fully determined by the URs + the rule system, they were not important either.

176. In short, to a large extent our increasing phonological sophistication has worked to the detriment of our understanding of diachronic phonology. We keep getting it wrong. And the fix is not obvious. The most obvious step might be to include phonological rules in the correspondence tables — as I mentioned above. The rules would then 'correspond' in the way segments do, and could be reconstructed based on that correspondence.

[19] That does not entail that they had the *correct* theory of that level of representation for any particular language; they certainly did not in many cases.

177. For example, many Indo-European branches have daughters which have a rule of final obstruent devoicing, including some archaic branches (probably, though there are issues) like Sanskrit. But the German rule, the Russian rule, the Wakhi rule and the Sanskrit rule have nothing to do with one another: the 'final obstruents' subjected to the rule have come to be final by different paths and do not stand in any sort of correspondence. Proto-Indo-European may or may not have had a rule of final obstruent devoicing, but if it did, the fact that German has such a rule is a completely unrelated fact. If we are going to put rules in the table, we need a mechanism for figuring out when such rules stand in correspondence and when they do not.

178. I don't think it will be useful to development such a notion as 'rule correspondence', however. Just as the fact that English sytems 'have a [p]', and Latin systems 'have a [p]', and Sanskrit systems 'have a [p]' does not allow us to conclude that the proto-language 'had a [p]', the same is true for rules. There is a quite different way to get at the 'phonological system' of the protolanguage.

179. The correct approach to the problem — and I think this represents the general practice of many historical linguists already, though they have not made it explicit — seems to me to be to reconstruct the phonetic output forms of the lexical items in the table (some may be morphologically complex). This will give us an inventory of proto-language phonetic representations over which phonemic and morpho-syntactic analyses can be performed. From such analyses we get a grammar and lexicon of the proto-language.

180. It is somewhat ironic that the push to be **more** theoretical, in our approach, moves us **away** from the focus on the rule system and (phonemic) lexicon that has characterized modern historical linguistics. But I have to shown you that if we adopt an I-language understanding of precisely what a 'language' is, and an acquisition-based theory of 'change', we begin to the right way to approach diachronic phonology, bringing **all** of our modern analytical machinery — UG, the nature of phonological computation, phonological acquisition, and the representational phonetics of the human mind — we can once again begin making progress.

References

Bender, B. W., Goodenough, W. H., Jackson, F. H., Marck, J. C., Rehg, K. L., Sohn, H.-m., Trussel, S., & Wang, J. W. 2003. Proto-Micronesian Reconstructions—1. *Oceanic Linguistics,* 42(1), 1–110.

Capell, A. 1969. *Grammar and Vocabulary of the Language of Sonsorol-Tobi* (Oceania Linguistic Monographs No. 12). University of Sydney.

Fox, A. 1995. *Linguistic Reconstruction: An introduction to theory and method.* Oxford University Press.

Goodenough, W. H., & H. Sugita. 1980. *Trukese-English Dictionary.* Memoirs of the American Philosophical Society, Vol. 141. American Philosophical Society.

Hale, Mark. 2015. The Comparative Method: Theoretical issues. In C. Bowern & B. Evans (Eds.), *The Routledge Handbook of Historical Linguistics* (pp. 137–151). Routledge.

— — —. 2007. *Historical Linguistics: Theory and Method.* Blackwell Textbooks in Linguistics, 21. Malden, MA: Wiley-Blackwell.

Hale, M., & Reiss, C. 2008. *The Phonological Enterprise.* Oxford University Press.

Honeybone, P. 2016. Are There Impossible Changes? θ > f but f ⊁ θ. *Papers in Historical Phonology,* 1.

Jackson, Frederick H. 1983. *The Internal and External Relationships of the Trukic Languages of Micronesia.* PhD dissertation, University of Hawai'i at Mānoa.

Lichtenberk, Frantisek. 1994. Reconstructing Heterogeneity, *Oceanic Linguistics* 33.1: 1–22.

Lynch, J. 2019. The Bilabial-to-Linguolabial Shift in Southern Oceanic: A subgrouping diagnostic? *Oceanic Linguistics,* 58(2), 292–323.

Oda, S. 1977. *The Syntax of Pulo Annian: A nuclear Micronesian language* (Unpublished doctoral dissertation). University of Hawaii.

Ohala, J. J. 1978. Southern Bantu vs. the World: The case of palatalization of labials. *Annual Meeting of the Berkeley Linguistics Society,* 4, 370–386.

Quackenbush, E. M. 1968. *From Sonsorol to Truk: A dialect chain* (Unpublished doctoral dissertation). University of Michigan, Ann Arbor.

Sohn, Ho-Min 1975. *Woleaian Reference Grammar.* University Press of Hawaii.

Sohn, Ho-Min, Anthony Tawerilmang, Isaac Langal, and Celestine Yangilmau.

1977. "Consonant Shifts and Subgrouping in the Sonsorol-Ulithi-Woleai

Chain." *Language Sciences* 44: 19–24.

Taylor, Ann and Susan Pintzuk. 2015. Testing the theory: Information structure in Old English. In *Information Structure and Syntactic Change in Germanic and Romance Languages*, Kirstin Bech and Kirstine Gunn Eide (eds), 53–78. John Benjamins.