

# Variability in Phonetics

**Mark Tatham**

Reproduced from *York Papers in Linguistics* No. 4, pp. 47-53 (1975).

Copyright © 1975 Mark Tatham

---

I decided when preparing this talk not to concern myself with some detail, however frontier, of recent developments in phonetics. Those phoneticians among you may have had your fill of that at the recent 8th International Congress in Leeds (August 1975), and those non-phoneticians among you might perhaps have soon lost interest among the technical aspects. I am aware that linguistics has become so divergent that at a meeting of this kind any speaker is facing a mixed audience. The classic answer is to provide a survey paper — but somehow that also has its attendant dullness. What I'm going to try to do is take one central problem which phoneticians are concerned with because it continues to turn up wherever we investigate, and attempt to use that problem to illustrate a major turn that phonetics has taken in very recent years. Only recently have we in phonetics and phonology come to grips in anything like a systematic way with this problem. Specifically I am going to talk about variability.

Variability is a term used to cover several types of phenomenon in both language sound patterns and in phonetic realisation of those patterns. I shall try to cover the major types as I proceed, hoping that as I do so you will get a sense of the way in which phoneticians are thinking these days.

Variability is often used to refer to the fact that every 'repetition' of an utterance is different. Whatever care a speaker takes he cannot avoid producing these differences. This fact contrasts sharply, for example, with the comparative lack of difference when, say, a tape recording is replayed over and over or when a speech synthesiser produces an utterance. Indeed with speech synthesis it is often this lack of variability which contributes to the somewhat artificial sound you hear. Many demonstrations of synthetic speech deliberately add a degree of randomness, or what we call 'jitter' to the speech to make it sound more natural or more human.

This variability can be looked at on several different levels. It is not at all difficult to imagine that the actual soundwaves of utterances when repeated, could be different: differences of amplitude, for example, are obvious, particularly if they are gross, even to someone not particularly listening for them.

Differences in overall rate of delivery of an utterance might be likewise obvious. Less obvious, however, will be shifts in formant frequencies, subtle changes in fundamental frequency or minor phase relationship changes, across the sound spectrum.

We might also look at articulator movements — the timing of certain articulatory events might vary. Even the temporal relationships between gross or 'landmark' events such as the positive articulator contact required for a stop vary. The actual movement of the articulators in space also varies in repeated same utterances.

Moving a stage further back in the chain of events ultimately producing soundwaves, we know that muscle contraction associated with articulator movement also varies, as do the neural signals themselves which are responsible for the muscle contraction.

We all know that early attempts to come to grips with this phenomenon produced a number of different solutions, perhaps the most significant (and certainly the most coherent) of which was the theory of the phoneme. Depending on which definition you adopted the phoneme was a unit of some kind, abstract in nature — though not always explicitly said to be such — which grouped together or stood over - dominated a certain range of variations. That the phoneme covered a range of variations was never very difficult to agree upon.

Precisely defining the range to be covered, however, was always guaranteed to precipitate considerable argument, as was likewise the formal method of ‘extracting’ the phoneme. So there were difficulties arising from how ‘linguistic’ you wanted to be — did you care about the substantially different articulator control required for a palatal /l/ as opposed to a velar /l/ and so want to arrange your variations into two groups with separate phonemes, or did you want to point out that since palatal /l/ and velar /l/ were never used to differentiate morphemes in English they should be grouped into the same phoneme? Did you want to have the phoneme as a physical event varying each time it was used over a defined range, or did you see the phoneme as a label for the entire range — that is, as some unit distant from physical events by some degree of abstraction? And so on.

But all this is history and much metatheoretical water has been under the bridge since the preoccupation with such arguments. Some of us are glad to think that it is all over and some are even embarrassed by that earlier preoccupation — that view though is uncharitable. Also under the bridge we have had rapid development of data gathering methods associated with advances in phonetics laboratory technology. We now know by using various more or less sophisticated pieces of equipment that variations are really — as far as the instruments are concerned — quite considerable. We are now literally inundated with data of this kind.

Obviously we get nowhere if every now well documented event is taken in isolation. Every linguist knows why we must not indulge in such listings or cataloguings. The problem is to avoid these listings.

First attempts could be described as no more than attempts to discover thresholds. By this I mean deciding how much of a variation is to count as a variation, So we had ideas like — if you can’t hear it, then forget it. Can you hear the difference between a palatal /l/ and a velar /l/? — Answer ‘yes’; so we must hold onto that variation and let go of the smaller unhearable variation which exists in the repetition of palatal /l/ and velar /l/. This was generally the technique employed more recently in so-called automatic speech recognition. Computers were programmed to react to or ignore certain variations in the soundwaves of speech in order to decide what was being said. Programming consisted of setting thresholds. It does not or should not take many tries to discover that such an approach is inappropriate.

Twenty years or so ago some other branches of linguistics became inverted — as has subsequently phonetics. We are properly no longer concerned with discovering, by whatever means, patterning in the human being’s speech, or speaking more technically, in the output of the device. We are concerned with how that output has come about.

This, of course, was a relatively new idea for phoneticians and it is no exaggeration to say that they were gripped by the novelty. Most of them, however, quite missed the point and failed to realise that inverting the approach meant abandoning the security of reality.

Reality was, perhaps still is, a concept very dear to the phonetician’s heart. My laboratory contains expensive and sophisticated equipment — it even boasts a computer. That makes me, it might seem, a master at discovering reality. If some of the facts of speaking still elude me then this is due to the shortage of equipment or technological shortcomings, one might imagine.

It was perhaps this belief more than anything which held us back. For it was this belief which caused us to bury ourselves in variability data and caused us to fail to understand properly the notion of abstraction. The theory of speech production — which was now seen to be what we were aiming for — was to be real. No one made the mistake of misunderstanding the meaning of theory or of model, but we did somehow believe in a vague idea of factual accuracy.

Phonology, which had become detached from phonetics, had not made this mistake. Phonology had fortunately subscribed to the idea that its theory should be about the information necessary to give a potential utterance sound shape rather than be about actually how that was done by the human being.

It was not, of course, until phoneticians realised the gross incompatibilities between their speech production models and the model of phonology that it was understood that the phonetician's idea of the supremacy of instrumental data was inappropriate.

I have dwelt on this history in order to emphasise the way in which the development of phonetics has been largely out of phase with the development of other branches of linguistics and particularly to show how somewhere along the line phonology became separated from phonetics — not just because the areas were becoming so large that no one researcher could be both phonologist and phonetician, but because, due ultimately to differences in the method of collecting data, incompatibility was inevitable. The two components are now again coming into phase.

Being, at this point, up to date in the history, I can now enter a more controversial area. Compatibility with the phonology is essential if either component is to say anything at all to the other.

It is not simply that in the last fifteen years or so the two components did not fit neatly together — an appropriate readjustment component might, however crudely, have taken care of that. The problem was that there was incompatibility of theoretical approach. Compare this with the fact that compatibility of theory had, for example, existed in the older English school of phonology/phonetics. It is by means of theoretical compatibility that data and evidence in the two components can be matched: attempting, as perhaps was done earlier, to contrive this matching of data collected from quite different theoretical perspectives, can be disastrous.

Let me illustrate this by returning to my central theme of variability. We have all become acquainted with the binary *vs.* *n*-ary argument which seems to continue *ad nauseam*.

Feature specification of phonological segments seems adequate and elegant, using a binary system of notation. Working toward the phonetic output of the phonology, however, such a specification seemed to become progressively less adequate to the point where some phonologists — even if they did not develop the idea much — spoke of converting at this level to the *n*-ary system, which of course kept to the established idea that the output of the phonology should embody only the minimum information necessary to enable pronunciation — or, better, to enable the relatively autonomous phonetic component to operate.

That the idea of *n*-ary feature specification should arise is traceable to concern by phoneticians over the observed variability at the surface.

Phonetics, however, has changed. We can now see that variability has been quite wrongly modelled. The switch of emphasis from 'How do we speak?' to 'What do we need to know to speak' has led us to wonder just to what extent variability is programmed in the system — and, if it is programmed at all, just at what point the necessary information is required.

Let me continue my palatal/velar /j/ example. We know that /j/ is used in English to distinguish morphemes; we know that we quite systematically use palatal and velar /j/ when speaking even though their use as such does not, as far as production goes, contribute any further to keeping distinct morphemes that would otherwise be confused. We further know that phonetically each of the pair is articulated differently — and, more importantly, systematically differently. Palatal /j/ has a front articulation adjacent to a front vowel and a retracted articulation adjacent to a back vowel and so on.

The experiment to determine whether the variation is voluntary or involuntary is not too difficult and its result shows us that this variation, though systematic, is not intended. The variation we have discovered does not originate in the brain, it seems, but is quite simply associated with the mechanical constraints on tongue movement. Such constraints are expressible by rules which have, in this case, nothing to do with linguistics, and which when operating to modify the intended articulation of palatal /j/ or velar /j/ result automatically in the variation.

Because of the way the phonetician is constructing his theory he now asks the question do the speaker know about these constraints in the narrow sense of 'know' used in linguistics. The answer is 'yes', and I'll tell you how we know that.

Consider that in English we have two voiceless stop consonants /t/ and /k/ which are articulated by bringing the tongue positively in contact with either the extreme front or the extreme back of the palate. The intended articulation of these two targets is modified by the mechanical constraints I mentioned just now — resulting in retracted /t/-sounds adjacent to back vowels like /a/ and fronted /k/-sounds adjacent to front vowels like /i/.

However, we could predict that in a language with, say, three or four palatal stop consonants, each with a different target, the operation of the mechanical rule would produce fronted central palatal consonants [c«—] which were further forward than retracted front palatal consonant [t—>]. Such a situation would, it could be hypothesised, be perceptually intolerable since a decoding confusion would result. Accordingly such languages severely limit the range of front-back variation that occurs. In other words, signals are sent from the brain which are intended to counteract the mechanical constraints. The point of all this is that such signals could not be sent unless we knew the constraint. The deliberate modification of the constraints is a function of the linguistic intention for the articulation to be kept apart for perceptual reasons and is a function of the mechanical constraint itself.

What I want you to notice, of course, is that empirical observation of what we do when we speak is contributing to the formulation of a phonetic model which is a statement of what we need to know when we speak rather than what we do when we speak.

There is another kind of, at first sight, random variation in speaking that I'd like to mention. There have been a number of experiments designed initially to discover whether the tense/lax phonological feature used to distinguish, say, /p/ and /b/, at the output of phonology, was realised phonetically in terms of the degree of contraction of the muscles controlling the prime articulator — in this case the lips. Using the relatively new and fairly elaborate technique of electromyography to detect the amount of electrochemical activity in the muscles it has been claimed by some researchers that the tense/lax distinction cannot be observed in the articulation itself.

It is not this finding which I wish to dwell on — that is relatively well known — but the detail of the data obtained. Notwithstanding the statement I just made about the lack of a phonetic correlate of the tense/lax difference between /p/ and /b/ in the phonology, if I articulate a /p/ and then a /b/, I can guarantee that the amount of contraction in the lip muscles will not be the same. It may well have been more for /p/ than for /b/. If so, then I might have said that there was 'reality', in some sense, to the abstract idea tense/lax because /p/ is specified as [+tense] — if the other way around I would have concluded from the negative correlation that tense and lax should be reversed in the phonological specification of this segment or that [+tense] means less contraction and that [-tense] or lax means more contraction. However, if I say another /p/ and another /b/ I discover that the result might be the other way around. And if I go on saying /p/s and /b/s I will notice that sometimes /p/ has more contraction than /b/, sometimes the other way around, and that sometimes the difference is great and sometimes small. My earlier conclusion that lip contraction for /p/ and /b/ is the 'same' would need explanation, therefore. It is, of course, based on a statistical test applied to many articulations of /p/ and /b/. There are two points here: to what extent is my result 'same' based on reality and to what extent is this variation or spread of individual results under our control?

I don't propose to go into the notion 'same' except to point out that it was this very problem that bedevilled the phonemicists. The results are the same, of course, on some abstract level and the formal device used to determine sameness and to define the abstract level has been the statistic — on no actual occasion of utterancing /p/ and /b/ are they the 'same' in the real world.

The variation within, say, /p/ is what is interesting. Could it be the case that we, perhaps for the sake of variety, deliberately vary the articulation? That may be, but adopting such a view as a preliminary model would not be all that sound. Alternatively, could it be that the variation is unintentional and that all we do is make sure it falls within a certain range — rather as earlier we made sure that our palatal consonants fell in a certain range; only this time

that range is governed by the fact that the lips must close and close sufficiently to support the increasing air pressure in the oral cavity? The hypothesis is that we govern the range of variability. A demonstration of this can be seen in another experiment.

In French there is a set of back rounded vowels contrasting with two sets of front vowel, one round and the other spread or *non*-round. Consider the single back vowel [u] and the pair of front vowels [i] and [y] — [u] contrasts with the others on the front/back feature, whereas [i] contrasts with [y] on the rounding feature, the tongue position being substantially similar. The acoustic distinction between [y] and [i] derive from the lip rounding.

In articulating these vowels, therefore, to avoid perceptual confusion it is essential with [u] to get the tongue right and with [y] to get the lip rounding right. Lip rounding for [u] is not critical. Suppose we look closely at lip rounding in [u] and [y] in French. We might expect an articulation with more precision for [y] than for [u]. Assuming that more precision means less variability we could hypothesise that over repeated articulations of [y] there would be a narrower range of variation than there would be with [u]. And indeed this turns out to be the case — [y] is articulated in French with more precision on the lip rounding parameter than [u] has. To obtain the narrower range of variation associated with greater precision it is necessary for the speaker to know the characteristics governing the variation which would have occurred without additional intended precision.

I have been using these examples of experiments in phonetics to illustrate not only the sort of data we can collect but to give some idea of the direction in which the theory of phonetics is moving. We are clearly concerned with what a speaker knows. But I have actually been talking about what he needs to know: this perhaps slight nuance reflects the way the phonetician goes about his model building.

The model I have been using dwells on the idea of information — just as does the model in the other components of the grammar. I do not know whether phonetics requires greater attention to neuro-mechanical constraints than semantics does to pure neural constraints. Some of the logical operations which may turn out to be universal at the semantic level may well give us some indication of neural constraints and begin to define some areas of limitation in the brain. The mechanical constraints in speaking which do turn out to be universal tell us something about the physiological properties of the vocal apparatus — but, of course, we don't need a linguistic model to tell us that.

What is interesting linguistically is the way the brain takes account of the mechanics of the system — sometimes letting things like variability have full rein, and at other times deliberately restricting (but never eliminating) the variability. That is — the way the brain takes account of this phenomenon in providing a soundwave encoding of language. Clearly a major constraint on the phonology is provided by the neural processing limitations of the brain itself, but clearly also a further major constraint derives from knowledge of the properties of the speaking apparatus and the possibilities of its control.

The phonetic model I am talking about is therefore a model of the knowledge a speaker has of the vocal apparatus — including its limitations — and of its control. You could imagine that speaking would at least involve knowing either the target articulations required for realisation of phonological segments or the soundwaves associated with these segments. Furthermore it involves also knowing which muscles are associated with which articulators and what would be the degree of contraction necessary to achieve articulator movement to the required position and at what time such contractions were to occur. It would also be necessary to know something of the phonological system and how it relates to the phonetics — to establish, for example, that here is a language with four palatal stops therefore requiring somewhat more precise control of the tongue than would normally be necessary. Such increased precision would require knowledge of the mechanical rules governing the variability in the first place. And so on.

The phoneme problem has of course been solved by deciding there is no problem. That is, by modelling the data in such a way that the problem does not arise. By inverting the model and being concerned initially with the minimum necessary information of a phonological

nature required for distinguishing morphemes, passing then to the modification of this information to include systematic variations of a *non-phonetically* constrained nature (like palatal and velar /l/s) we avoid altogether the problem of how to group such variants. If we stop at that point by stating that we have included all the linguistic content we wish voluntarily to give a potential utterance, we delimit the phonological component. Phonetics is concerned with realising the demands of the phonology and to do so must operate on an input coming from that phonology. It might do so in the way I have outlined.

These illustrations I have been using show the extent to which phonetics has become aware of the use of abstraction. It is fairly clear that early phonemics had grasped that some kind of abstraction was necessary — to rid oneself of the problem of tackling variability; it is equally clear that the phoneticians or phonologists of the time were not entirely aware of what it means to be working in the abstract. As I said earlier, it is difficult to rid oneself of the ‘realities’ of the phonetics laboratory.