THE PERCEPTION OF VOWEL SOUNDS

Peter Ladefoged, M.A.
Lecturer in Phonetics, University of Edinburgh

A thesis submitted for the degree Ph.D.

Also published as a supplementary report under the Ministry of Supply Extra-Mural Contract No. 7/GEN/1421 to investigate the specification of speech sounds by means of acoustic parameters.

EX BIBL UNIV EDINBURGEN

SUMMARY

This thesis discusses what is meant by "phonetic quality" in so far as vowel sounds are concerned. The first chapter points out that phoneticians implicitly describe speech sounds in two different ways: firstly in terms of phonetic quality which is in practice the attributes of a sound which convey information by virtue of being part of a socio-linguistic code; and secondly in terms of their personal quality, which consists of the attributes that convey information only about the speaker considered as an individual.

The second chapter discusses the historical development of our present descriptions of the phonetic quality of vowels. It is suggested that not only is there very little evidence for the traditional articulatory descriptions, but also that the acoustic descriptions of vowels in terms of their formants have never been shown to be adequate for all vowels.

The third chapter describes the collection and spectrographic analysis of a large number of sets of cardinal vowels spoken by twelve phoneticians. It appears that no precise statements about the acoustic correlates of phonetic quality can be made; but it is probable that vowels such as [e, ɛ, a, ɑ] are best specified in terms of the relation between the pitches of their first two formants and the pitches of the first two formants of other vowels of this general type spoken by the same speaker. This is not true of vowels such as [i, u, o, ɔ] . These vowels often cannot be conveniently analysed as having two or three formants; nor, even when a procrustean formant specification is possible, are they adequately specified in terms of formant frequencies (or pitches). An illustration is given of a pair of vowels spoken by one speaker in which the frequencies of the first three

formants are almost identical, but which are nevertheless heard as having quite different qualities (viz. [ɔ] and [o]), probably because their formant bandwidths are very different.

In chapter four it is shown that in so far as a formant frequency specification is adequate, it is the relative and not the absolute values of the formant frequencies which convey linguistic information. In an experiment using synthetic speech it was found that the same test sound was identified in different ways when it was preceded by different versions of an introductory sentence.

Chapter five discusses the psychological mechanism responsible for this process. It is suggested that previous theories accounting for similar processes are not fully adequate, in that they incorrectly consider a subject's adaptation to a series of stimuli to be independent of the response which he makes to the stimuli.

The sixth chapter discusses an experiment in which 18 phoneticians specified the vowels in ten words in a language which was unknown to them in terms of a standardised cardinal vowel diagram. It is shown that among phoneticians trained in the use of the cardinal vowel system, there is often a high degree of agreement in the specification of a vowel which has a degree of lip rounding similar to that of the nearest cardinal vowel; but many disagreements occur in the specifications of vowels with less familiar lip positions. Nevertheless it is concluded that, in view of the results detailed in the previous chapters, at present the best method of specifying many vowels is in terms of the traditional cardinal vowel system.

PREFACE

Parts of this thesis have already been published or submitted for publication elsewhere. Much of chapter one is contained in "The classification of vowels" Lingua, 5.2, (1956). Material from this article is also incorporated in chapter six. Chapters four and five are based on articles written in collaboration with D.E. Broadbent: "Information conveyed by vowels" which appeared in J.Acoust.Soc.Amer. 29.1, (1957), and "Vowel judgments and adaptation level" which is being submitted for publication as a Royal Society paper. Both authors are, of course, responsible for the final published forms of both articles. The first article, however was actually prepared by the present author; and the version which appears in chapter four is substantially as originally written by him and without Broadbent's later additions and alterations. On the other hand, the draft of the second article was prepared by Broadbent. This article has been considerably rewritten by the present author before being incorporated in this thesis: some of the points made in the original version have been omitted (particularlythose involving the use of advanced statistical analyses); and others have been substantially amended. The result must therefore be considered as the work of the present author, and Broadbent should not be held responsible for any of the statements that are made. Nevertheless the fact remains that the development of the original ideas expressed in chapter five was entirely due to him, and it might be best to consider the present author, in this chapter, as Broadbent's inadequate mouthpiece.

Several other people have also contributed a great deal to this thesis. Walter Lawrence took part in the discussions which resulted in the experiments described in chapters four and five; and it is of course solely due to his work in devising a suitable

speech synthesiser that these experiments became possible. J. Anthony was also particularly helpful in the preparation of the material for the synthetic speech experiments, and in discussions of problems of acoustic analysis and synthesis. In addition R. Sillito, Elizabeth Uldall and Peter Strevens offered a great deal of helpful criticism. However the greatest debt is obviously due to David Abercrombie. Many of the concepts elaborated here were originally suggested by him; and he has been critically concerned with the development of the whole thesis.

Finally I would like to thank all the phoneticians who took part in the experiments described in chapters three and six; it was an especially great privilege to be able to work with Professor Daniel Jones in this way.

Chapter 1. The phonetic quality of vowels.

Phoneticians are often faced with the problem of comparing vowels. Because of this they have elaborated theories whereby vowels are classified in certain ways according to their quality. In practice most of these theories work very well, and there is a high degree of agreement among the judgments of skilled phoneticians. It is often possible for a phonetician to describe a vowel in such a way that another phonetician who has not heard the vowel is nevertheless aware of the particular vowel quality that has been specified. However, the assumptions underlying the theories of vowel classification are seldom explicitly stated. This thesis will consider the ways in which phoneticians make their judgments; it will also consider the possibility of making instrumental measurements which would lead to a similar form of specification.

When a phonetician hears a vowel he is usually capable of allocating it to one of a number of "general human categories of sound". (Abercrombie 1954). Although phoneticians sometimes use a limited number of symbols to indicate broad areas of vowel quality, in detailed treatises they usually distinguish between many different types of vowel. Bloch and Trager (1942) give 42 different vowel symbols, and recognise the possibility of having to use additional diacritics "where it is necessary to indicate still finer distinctions." (op.cit.) Other phoneticians denote small differences between vowels by means of vowel diagrams and descriptive labels. Thus Jones (1956) distinguishes between many variant pronunciations of the vowels of English, some of them being closer together than any two symbols used by Bloch and Trager. Many other phoneticians have produced similar classifications. But what are the criteria on which all these

judgments are based? What are phoneticians really doing when they describe a vowel sound by allocating it to a certain box in their scheme of categories, or a certain point on their vowel diagrams?

Obviously this has nothing to do with the different vowels that occur in any one language. It is a general phonetic problem, and cannot be considered in terms of phonemes or diaphones. This point, however, has been overlooked by some investigators. Russell (1928) for instance, in his work on the relation between vowel quality and articulation, denotes the quality of the vowels by references to keywords, e.g., "vowel [E], pep". But as Jones (1956) has noted, there are many variant pronunciations of the vowels of English. Some speakers of Standard English have a vowel in the word pep which is the same as the vowel in the word pap as spoken by other speakers. Consequently we can regard "vowel [E], pep" only as a label giving a very rough indication of the vowel spoken by a particular subject. For many purposes this may be all that is required. But a more precise classification is necessary when discussing the nature of vowel quality. Consequently, Russell's research and conclusions are of only limited value.

Most other works in the field recognise that vowel qualities cannot be precisely specified in term of key-words. For example, Peterson (1952) who conducted tests with the vowels [ɪ] and [æ] took elaborate precautions to ensure that the vowels pronounced by different subjects were in fact phonetically identical. Consequently the resulting research is of much greater value. However, even a writer such as Joos (1948), who is usually well aware of the dangers of using phonemic

categories as a basis for research into phonetic quality, occasionally becomes confused. When discussing acoustic measurements of the second vowel in the word hotel as pronounced by three different speakers he says: "Evidently the three vowels were perceptibly different. Yet the phrase Where is a hotel was pronounced in exactly the same style by all three speakers. True, they do not speak the same dialect, but they all speak American English, and there is no American isogloss dividing varieties of /ɛ/ in hotel" [No evidence for this statement is cited] ... "this is not a question of phonemics: these three vowels are already phonetically identical, although acoustically distinct." (op.cit.) As we are given no other relevant evidence concerning the pronunciation of these three speakers, it is difficult to know what is meant by the statement that these vowels are phonetically identical. Joos does not say that he means that he has heard these vowels, and, as a phonetician, considers them to be identical. There may be no American isogloss dividing the varieties of /ɛ/ in hotel. But this vowel is certainly known to vary a good deal from speaker to speaker; and it does not seem useful to assume (as Joos appears to do) that these vowels have the same phonetic quality just because they belong to the same diaphone.

It is obvious that phonemic classifications do not provide a satisfactory basis for establishing phonetic categories. But it is essential for many purposes that there should be criteria for classifying vowels in general phonetic terms. Much of the teaching of phonetics rests on the premise that it is meaningful to talk about the similarity of vowels spoken by different people. A phonetician often has to consider whether a vowel he has pronounced is, or is not, the same as a vowel another person, such as an informant, has pronounced. This situation also occurs in

teaching the pronunciation of a language. It might be considered that on such occasions the teacher is merely trying to prevent the pupil from making a vowel sound which could be confused with that in a different word. But most teachers of pronunciation are not content with ensuring that their pupils are understood. They are usually trying to prevent them from speaking with any accent which they consider to be undesirable.

We see, therefore, that a phonetician has to be capable of distinguishing many more vowel qualities than there are in his own speech. It is not sufficient for him to be able to say that a given vowel sound is like one or other of his own vowel sounds. He must be capable of saying something about the degree of similarity. His doing so indicates that his own vowels are like known places on a map, and that he is making a phonetic judgment in stating the distance between one of his own vowels and a vowel pronounced by another speaker.

Phonetic similarity in this sense seems to be definable only in terms of judgment by the trained observer. Any listener can distinguish between certain vowel sounds; after a little training it becomes possible to distinguish between many more; a highly trained phonetician can differentiate between a very large, but nevertheless finite, number of vowels. Consequently, in phonetic research, two vowels can be equated if, and only if, a trained phonetician regards them as being the same.

So far in this discussion of similarity no reference has been made to the fact that the phonetician will sometimes equate two vowels that can be differentiated by any observer, however untrained. Thus vowels spoken by a bass and a soprano may be considered to be the same, although they are obviously

different sounds. The psychological processes underlying this identification are not fully understood. What probably happens is as follows. The listening phonetician assesses the sound; but he does not consider the sound as a whole. He focuses his attention on certain features of the auditory sensation, by reference to which he can determine the quality of the vowel. The rest of the auditory sensation is considered to be irrelevant when judgments concerning quality are being made. In a similar way, when judgments concerning pitch are being made the listener focuses his attention on certain other features of the auditory sensation, neglecting those features from which he could abstract information about the loudness and quality of the sound.

We know a considerable amount about the acoustic correlates of pitch and loudness. Thus we know that the pitch of a sound depends primarily on the fundamental frequency of the sound wave whereas the loudness is largely dependent on the amplitude. But we do not yet know so much about the acoustic correlates of that part of the auditory sensation from which vowel quality can be abstracted. But before we can discuss this question we must note that speech sounds can be equated in a way peculiarly their own. Many other sounds can be classified according to three separate factors: their pitch, their loudness, and their quality. Thus we can disregard loudness and quality and compare the pitch of two sounds, considering only whether they are on the same note or not; similarly we can compare their loudness irrespective of the other two factors; and we can also compare their quality irrespective of everything else. It might appear that in comparing the vowel sounds of a soprano and a bass we are assessing their quality. But it is not as simple as this; in the case

of speech sounds we can consider pitch, loudness and two sorts of quality, which we may label phonetic quality and personal quality. Thus when we say that two vowels are different we usually do not mean to imply anything about their pitch and loudness, nor about their personal quality, but only that they differ with regard to that one aspect of their quality which we term phonetic quality.

As a particular instance we may consider the case of vowels pronounced by speakers with recognisably different voices. We may find that both speaker A and speaker B are pronouncing what we recognise to be the same vowel; furthermore, these vowels may be identical in pitch and loudness; but they may nevertheless be distinguished as sounds, i.e., we could tell that they were not pronounced by the same person. Thus when listening to speech sounds observers can assess, with a high degree of convergence, four variables; loudness, pitch and two aspects of quality.

Most people cannot consider any other sounds in terms of four variables. It is only in the case of speech sounds that nearly everybody can say that they may be similar as regards one aspect of quality, but different in another.

This point has often been overlooked in discussions of vowel quality. There is a tendency to assume that variations in the personal quality of the vowels of two speakers are of the same kind as differences in phonetic quality. It is, of course, true that we can often distinguish between two speakers of a language by assessing the variations in the phonetic quality of their vowels. But the fact remains that it is also possible to find speakers whose vowels do not differ in what we wish to define as phonetic quality but who can nevertheless

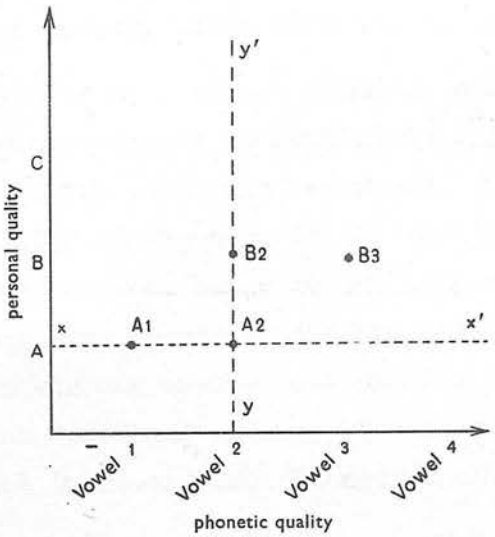


Figure 1.1
Differences in quality among speech sounds can be shown in two dimensions; one axis represents personal quality and the other represents phonetic quality. Thus the sounds designated by points on the line x—x′ have a personal quality in common; and the sounds designated by points on the line y—y′ have a particular phonetic quality in common.

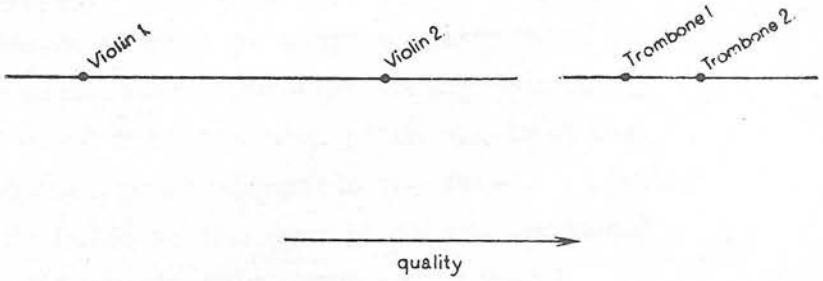


Figure 1.2 Differences in musical quality are in one dimension. The line representing trombone quality cannot meaningfully be placed above that representing violin quality. There is nothing in common between trombone 1. and violin 1.

be distinguished by another factor, their personal quality.

When we assess the quality of a speech sound we need a frame of reference which cannot easily be applied to the assessment of the quality of any other kind of sound. We cannot, for instance, consider differences in the quality of two violins playing the same note as being in any comparable with the differences between two people saying the same vowel. An expert violinmaker can easily tell violins apart by what might be called the "personal" quality of each violin. But this does not mean that there are two dimensions of quality for violins.

The problem is best explained by means of diagrams. Figure 1.1 on which speech sounds can be depicted, shows two dimensions: personal quality, and phonetic quality. (Pitch and loudness are considered as constants for this purpose). When speaker A pronounces vowel 1 at a given pitch and loudness we may represent it by a point as shown; similarly his pronunciation of vowel 2 at the same pitch and loudness is represented by a second point alongside the first. Speaker B pronouncing vowel 2 (also at the same pitch and loudness) is designated by a point immediately above A2. When B pronounces vowel 3 it is represented by a further point as shown. So in the diagram it is possible to depict different speakers (i.e., differences in personal quality) along one axis, and different vowels (i.e., variations in phonetic quality) along the other.

But it is not possible to diagram differences in the perception of the quality of different musical instruments in the same way. All such differences excluding pitch and

loudness can be shown on a straight line (figure 1.2). The fact that there are differences in quality between two violins merely means that there is a length of line corresponding to violin quality. By no means is it possible to represent two different dimensions of perceived quality of musical instruments. We may represent violin 1 playing a given note at a given loudness by a certain point, and violin 2 playing the same note at the same loudness by a second point alongside the first. If we now consider, say, two trombones playing the same note as the violins and at the same loudness, we shall be able to represent these sounds by two more points some distance along the same line. Even if it is felt that this is not possible, and that there has to be an entirely separate line for trombone quality, there can be no justification for placing this second line above the first. This would imply some kind of sameness about points along a vertical line, i.e., that there was something in common between trombone 1 and violin 1 , trombone 2 and violin 2 , etc. But this is not so; differences in perceived musical quality may be small or great, but they are all of one kind. It is only in the case of speech sounds that we habitually distinguish two dimensions of quality.

It should be noted that we are not saying that there is anything peculiar about vowel sounds considered as physical entities. It is only in the way that they are normally perceived that they differ from other sounds. We could no doubt learn to assess other sounds in terms of two kinds of quality. Indeed, it is possible that this is actually done by some observers. A musician considering organs could possibly tell organ 1 from organ 2, irrespective of the stops; in such a case it would be possible to represent the perceived

quality in two dimensions, the quality of being organ 1, or organ 2, being shown in one dimension, and the quality of the different stops (diapason, violin, trombone, etc.) being represented in the other. However, in practice musicians are not often concerned with this kind of assessment; whereas phoneticians are continually specifying speech sounds in terms of their phonetic quality. Nearly all phonetic theory relies on the tacit assumption that it is possible to recognise two kinds of quality. Nevertheless, this fundamental point is rarely considered and the different kinds of quality are seldom explicitly distinguished.

The peculiar position of speech sounds is due to their being habitually assessed as part of a means of communication. Every speaker has learnt to separate personal quality from phonetic quality as a result of his constant experience of the socio-linguistic system. Phoneticians, who are trained observers of socio-linguistic systems, have become highly skilled in their assessments of the different kinds of quality. We may conclude that when a phonetician equates vowels spoken by different voices he does so because in the appropriate socio-linguistic system there could be no difference of codified information conveyed by the differences in the sounds.

This does not mean that phonetic equivalence can be regarded as a product of phonemic classification. Spoken language conveys far more than is considered by most present day phoneme theories. The listener to speech apprehends not only the objective meaning of the words, but also indications of the speaker's attitudes and moods, his place of origin, his social status, etc. Much of this information has not yet been codified by phoneticians. But it is nevertheless

part of the message conveyed by means of the socio-linguistic system. Moreover some of it is usually taken into account by phoneticians who wish to specify the precise phonetic quality of a vowel. Anyone who is concerned with describing the sounds of a language, or with dialectology, or with teaching pronunciation, or with developing voice operated devices such as speech typewriters, in fact anyone who is concerned with a branch of phonetics which is not reducible to a simple phonemic solution, will inevitably have to consider phonetic quality in this way. Over and above the phonetic quality there are certain aspects of speech sounds that are personal to the speaker. They are mainly due to physiological features and are not part of any system. In themselves they convey no information except that the sounds were spoken by that particular speaker. Consequently they cannot be regarded as units in an accepted code.

This situation can be summed up by means of a diagram (figure 1.3) which shows the relation between the two main dichotomies. The first dichotomy is between the features of speech which have been learnt or acquired by the speaker (functional features); and the features of speech which are due to those physiological aspects of the vocal mechanism producing them which are not within the voluntary control of the speaker (organic features). The second dichotomy is between the features of speech which are characteristic solely of the speaker as an individual (personal quality); and those which are part of the socio-linguistic system (phonetic quality). The latter depends entirely on functional features of speech; whereas personal quality may be partly innate and partly acquired, and hence depends on both organic and functional features. (This viewpoint and the accompanything diagram will be further elaborated in chapter 4).

The distinction between phonetic quality (the attributes of the auditory sensation that enable phonetician to consider a speech sound as part of a socio-linguistic system) and personal quality (some other features of the auditory sensation) is one of the basic assumptions of phonetics.

Chapter 2 - The historical viewpoint

In order to get a clearer understanding of the problems involved in the description of vowels, it is advisable to consider how our present concepts of vowel quality arose. There are two main lines of development involved:- the articulatory and the acoustic. We will consider first the development of the description of vowels in articulatory terms.

One of the first writers to attempt to describe the position of the vocal organs during the pronunciation of vowel sounds was Robert Robinson (1617). His account is of great interest to us in that it includes a schematic diagram of the articulators. But its importance shouldnot be overemphasised, since it seems probable that Robinson's work had little or no influence on succeeding writers.

According to Robinson, the vowels are "framed by the placing of the tongue in sundry partes of the roofe of the mouth .... the first taking its beginning in the innermost part of the roofe or pallat, and so the rest continuing forward, each one orderly in his degree to the last place, being more neare to the outmost part of the roofe .... For the more manifest demonstration of the construction of the vowells, I have here devised and placed this ensuing figure .... [Figure 2.1] By the arch lyne AB, is represented the roofe of the mouth, by the point C, from whence the fiue seuerall lynes are drawne, is supposed the roote of the tongue it selfe, and by the seuerall angle of the same lynes vnder z nɛ ɔ e are supposed certaine eleuations and bendings of the tongue, which cause the fiue seuerall sounds called short vowells, for which the same characters z nɛ ɔ e

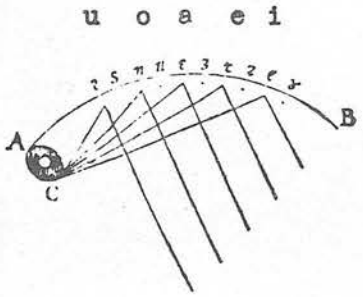


Figure 2.1 A diagram of the tongue positions for the short vowels from Robinson (1617). The IPA vowel symbols have been added above Robinson's symbols.

are framed". Op.cit. (It appears later that these characters correspond roughly to the phonetic symbols [u, o, a, e, i] .)

The next important articulatory description of vowel quality is due to Wallis (1653), who, in contrast to Robinson, greatly influenced all succeeding writers. His great book Grammatica Linguae Anglicanae (1653), to which the treatise De Loquela is prefixed, ran to five editions during his lifetime (1616-1703) and many other editions and psendo-editions were published subsequently (Lehnert, 1936). In the current work reference has been made to the sixth edition, published in 1765 in London. In addition the popular grammars of the eighteenth century were largely based on Wallis's work. Thus Greenwood (1711) says: "I have in this Book taking in everything that was Material from Dr. Wallis, but he writing for Foreigners and in Latin, I have not pursued his method". Similarly Brightland's Grammar of the English Tongue (1711) is to a great extent a translation of Wallis's work.

Wallis, in the section entitled "De sonorum formatione", describes vowels in terms of a pair of oppositions: high as opposed to low tongue position; and rounded as opposed to spread lip position. The original description in Wallis (1653) does not include a diagram of any kind. But the restatement in Brightland's (1711) shows a table of the form reproduced here as figure 2.2. All the descriptions of this date are difficult to interpret. But it seems that these authors were thinking in terms of a (somewhat muddled) two parameter system.

Wallis's contemporary and rival, William Holder, did not advance the description of vowels to any extent. He used a uni-dimensional system for classifying vowel quality, though noting that sometimes extra factors had to be taken into account

(Holder 1669). But his descriptions are inclined to be ad hoc specifications, which cannot be considered as a unified system.

A more interesting contemporary account of the formation of vowels is given by Isaac Newton who said in his notebooks (c.1665):

"The fore pte of ye mouth straitned by (drawing in y^{e} lips &) contracting ye middle pte of y^{e} tongue to y^{e} rough of y^{e} mouth but y^{e} hinder parte at the throate being widned maketh y^{e} pronuntiation of y, being more dilated at y^{e} rough it maketh i, more still at y^{e} rough but straitned at y^{e} throate makes e still more straitned at ye throat dilated at y^{e} rough and ye lipps & chaps a little opened makes a, more still y^{e} lips & chaps wide open make o, more still ye lipps a little thrust out & contracted makes ω , more still makes u, y^{e} throate and lipps most straitned & lipps thrust out most makes w.

Soe y^{e} greatest cavity in y^{e} mouth being first made in y^{e} throate & thence by degrees moved towards ye lipps further from y^{e} larinx causes y^{e} pronunciation of ye vowells in order y i e a o ω u w. The filling of a very deepe flaggon wth a constant streame of beere or water sounds y^{e} vowells in this order w, u, ω, o, a, e, i, y." (The symbols are the available printed forms nearest to those of Newton's original handwriting).

By the first part of the nineteenth century the description of vowels had become standardised in a form only slightly different from that originally put forward by Wallis in 1653. Thus Wheatstone (1837) describes vowels as departing from "aw" as in folly in two series. "In the first (series), the external aperture remains open, and the internal cavity gradually diminishes by the successive alterations of the tongue. In the

second the tongue remains depressed, but the aperture of the lips is gradually diminished. There is also an intermediate series of vowel sounds, obtained by different elevations of the tongue when the lips are partially closed; these though abounding in many foreign languages, are not used in our own".

This description which is typical of many of the time, suggests a diagram of the form shown in figure 2.3. A similar scheme is used as late as 1863 by Helmholtz who uses the diagram shown in figure 2.4 (keywords taken from footnotes in Ellis (1885) translation).

The mid-nineteenth century traditional scheme was radically altered by A.M. Bell. In his early work Bell (1849) had used the current form of description in which vowels were classified in terms of two series "labials" and "linguals" plus an intermediate series "labio-linguals". But when he was working out his system of "visible speech" (Bell 1867), he realised that this form of description was inadequate. He describes very poignantly how he was haunted by the impossibility of placing the vowel in Sir into the existing categories.

The system of vowel classification devised by Bell (1867) was new in many ways; but in particular in that it appeared to describe the position of the tongue in two dimensions. The height of the tongue was described as being high, mid, or low; and the highest point of the tongue was described as being at the front of the mouth, at the back, or "mixed" (i.e., with both the front and the back of the tongue raised). This resulted in a scheme in which there were nine "cardinal" tongue positions. (This is, incidentally, the first use of the word "cardinal" in the description of vowels). Bell usually discusses these positions in terms of a 3 x 3 arrangement

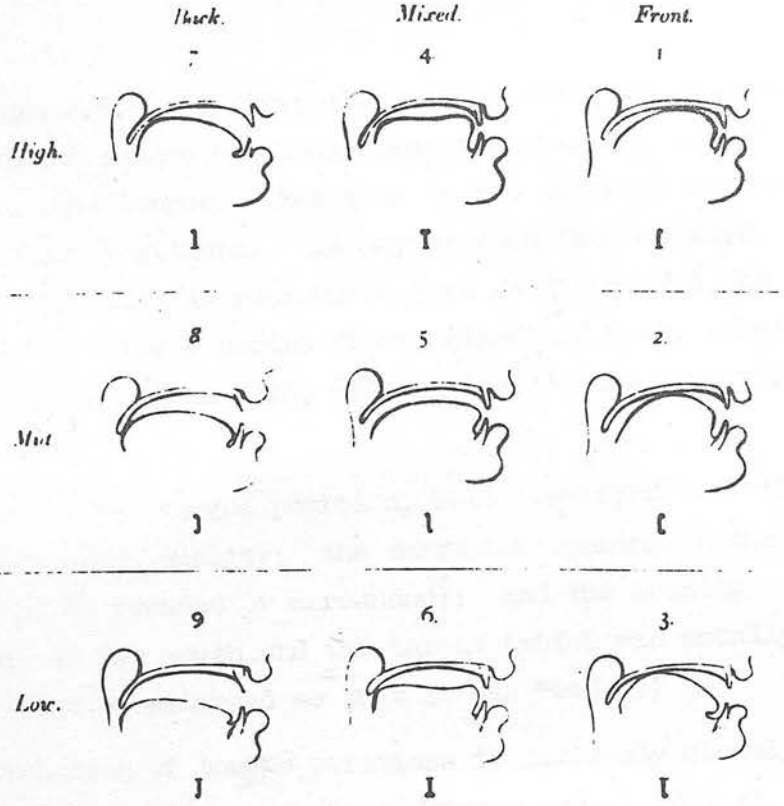


Figure 2.5 A display of vowel articulations from Bell (1867).

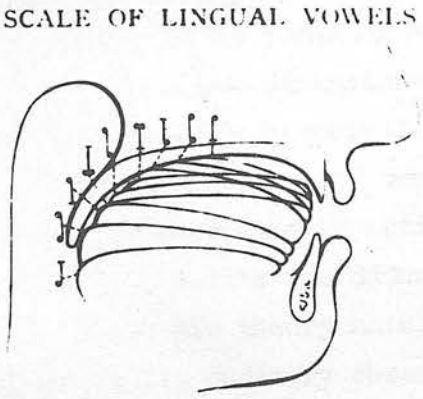


Figure 2.6 A combined view of the tongue positions in figure 2.5, also from Bell (1867).

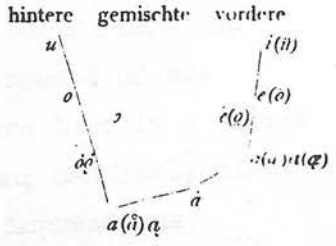


Figure 2.7 A representation of the tongue positions of French vowels from Passy (1887).

as shown in figure 2.5. At first glance this gives the impression of a two-dimensional scheme of classifying the position of the highest points of the tongue. But this is not actually how Bell considers the tongue positions. As may be seen from another of Bell's diagrams, which is reproduced here as figure 2.6, he considers each of the three series front, mixed and back, separately, and does not apply the terms high, mid and low, in the same way to each of them.

In addition to the tongue position, Bell described two other factors affecting vowel quality: the degree of opening of the lips (which could be rounded or unrounded); and the opening between the back of the mouth and the throat (which was usually "primary" but could be enlarged so that it was "wide".)

Bell's tabulation of tongue positions is obviously closely allied to many modern methods of classifying vowels. But we should note that this is not a measure of its validity. For two hundred years before Bell another system had been used. Now, ninety years later, in which nearly all considerations of vowels have been in terms of the position of the highest point of the tongue in a two dimensional system, we are apt to believe that this system of description is based on known facts. But our modern descriptions of vowels are not the result of experimental observation of articulations, but are largely a direct adaptation of Bell's two dimensional tabulation; and Bell, like Wallis, based his theory mainly on subjective impressions. Although Bell's auditory observations were far better than those of the earlier works, even a cursory glance at figure 2.6 shows that his knowledge of articulatory positions was not in fact much greater.

With the later work of Bell we enter into an era when

phoneticians had the avowed intention of doing one thing; but actually succeeded in doing another. They set out to describe the positions of the vocal organs during the production of different sounds; but their success in this enterprise was partial. They did, however, succeed in providing categories with which to describe their auditory impressions.

In Bell's classificatory scheme 36 vowel qualities were specified. Sweet (1890) modified and elaborated Bell's system so that he was able to specify 72 vowel qualities. All these vowels are given articulatory descriptions; but it is often difficult to understand what Sweet meant. In particular his doubling of Bell's nine tongue positions by the addition of a further nine "shifted" positions, and his modification of Bell's "primary" and "wide" to "narrow" and "wide" seem especially incomprehensible. Sweet had a very good ear; and, whatever he may have thought he was doing, he was probably only devising categories which would accomodate the auditory distinctions which he could hear.

Both Bell and Sweet had declared that they were specifying only a limited number of tongue position solely in order to simplify their descriptions. Passy (1887) considered this to be too inexact, and accordingly devised the diagram of the French vowels shown in figure 2.7. This is probably the first example of this kind of vowel diagram; and certainly, in devising it, Passy was moving further towards the representation of auditory qualities rather than articulatory positions. At the time when this diagram was published Passy had no exact knowledge of the tongue positions of French vowels; and when he represented them as being at various distances from one another, he was almost certainly assessing what he heard rather than what he felt.

The next step was taken by Jones (1917) who instituted the well-known cardinal vowel system in which, in the first place, eight vowel qualities are specified, two in terms of tongue positions and the other six in terms of what was then called accoustic (and what we now call auditory) criteria. But it should be noted that although Jones uses this system to describe auditorily perceived qualities, it is apparent both in his conversation (see chapter 3) and in his published works, that he considers that a point on a cardinal vowel diagram actually specifies an approximate tongue position. Thus the legend beneath the vowel diagram in the eleventh (1956) edition of his English Pronouncing Dictionary reads "A diagrammatic representation of approximate tongue-positions of average English vowels compared with those of Cardinal Vowels. (The dots indicate roughly the positions of the highest point of the tongue)". But, in fact, although one knows the position of the highest point of the tongue for some of Jones's cardinal vowels, there is no published data about his tongue positions during the pronunciation of either the other cardinal vowels, or any of the vowels of English. It is probably incorrect to consider that points on vowel diagrams describe tongue positions for any speaker, even approximately; and in any case it should be noted that Jones makes fairly precise, and not approximate statements about the quality of many vowels.

The only published X-ray data showing the tongue positions in a complete set of cardinal vowels is that of S.Jones (1929), No measurements are given, and the quality of the reproduction makes it difficult to make an accurate appraisal of the tongue position. But it is readily apparent that the tongue positions are very different from the theoretical description of the articulations of the cardinal vowels. In addition, it seems that the tongue does not move in a series of even approximately

equidistant steps when a set of cardinal vowels is pronounced. Figure 2.8 shows the distances on the published photographs between the highest points of the tongue in the first four and the second four cardinal vowels; the tongue has such a different shape for the front and for the back vowels that it is meaningless to compare cardinal vowel number four with number five. There is no scale attached to the photographs, but it is apparent that out of the three intervals between the first four cardinal vowels, and the three intervals between the second four cardinal vowels, only two are equal; and some are so different from the others that it is difficult to see how phoneticians could persist in considering that the tongue moves in a series of approximately equidistant steps.

The first large scale attack on articulatory descriptions was launched by Russell (1928) who even went so far as to declare that "phoneticians are thinking in terms of acoustic fact, and using physiological fantasy to express the idea". (op.cit.) This is perhaps a little over-stated, since the "acoustic facts" are even now not entirely agreed. But it does seem probable that some of our auditory impressions of vowel quality may be more simply correlated with acoustic measurements rather than with articulatory data. This being so, it is advisable to consider the development of our present day acoustic theories of vowel quality.

The first major contribution to the acoustics of speech was made by Willis (1829). He decided on an original plan of operations: "namely, neglecting entirely the organs of speech, to determine, if possible, upon the usual acoustic instruments, what forms of cavities or other conditions, are essential to the production of these sounds". (op.cit.) His experiments involved

the use of a reed in a pipe with a variable effective length. Thus he produced two tones simultaneously, that of the reed, and that of the particular harmonic amplified. Possibly because his experiments took this form, he came to the conclusion that there were two acoustic features for each vowel sound: the pitch on which it was said (corresponding to the pitch of the reed); and its own characteristic note (corresponding to the resonance tone of the pipe). In his own words "... a given vowel is merely the rapid repetition of its peculiar note". (op.cit.)

This theory was taken up and elaborated by Helmholtz (1867). He found that the vowels A, O, U (i.e, [a, o, u]) had a single resonance as indicated by Willis, but that E, I. Ö, Ü (i.e., [e, i, ø, y]) were characterised by two fixed resonances each.

A few years later A.G. Bell (1879) found two resonances for each of his father's (A.M. Bell's) cardinal vowels. After this, for the next 65 years, apart from the further observations of other vowels by Lloyd (1891), Paget (1923), Crandall (1925), [redacted] Fletcher (1929) and others, the acoustic theory of vowel quality did not advance to any extent. During all this time it became more and more common to consider that every vowel was characterised by two resonances or "formants". But it is noteworthy that the evidence presented was often capable of other interpretations. Nearly all the investigators presumed that because some vowels were clearly characterised by two formants, all other vowels must be capable of a similar specification. As a result, somewhat procrustean procedures were often involved in the location of formants.

After 1945 a number of writers began to look at the established theories in a new way. The first publication of a relation between the formant frequencies and the traditional form of vowel diagram is by Essner (1947), who showed that, if the frequency of the first formant was plotted against the

frequency of the second formant, "Nous constatons que l'ensemble des voyelles considérées est inscriptible dans un quadrilatère, et ceci rapproache cette représentation acoustique de la représentation actuellement classique de Daniel Jones et Lloyd James"(op. cit.) Seemingly independently this point was also made by Joos (1948); and members of the Bell Laboratories published an article in which "it is shown that movements of the major resonances in the voiced sounds of speech may be represented by traces in a three dimensional graph". (Potter and Peterson (1948)).

Both Joos (1948) and Potter and Peterson (1948) acknowledge their recognition of the importance of formant frequencies as being in part due to the development of the sound spectrograph (Potter et al., 1946). This was the first instrument to make large scale acoustic analysis a practical possibility; and this had important consequences. Previously all the observations of formant frequencies had been based on measurements of the acoustic characteristics of the vowels of a few individuals. But now, after the analysis of a large number of vowels, it became apparent that vowels which were considered to be phonetically equivalent did not necessarily have the same acoustic characteristics.

In an article discussing the theoretical issues involved, Peterson (1951) pointed out that "four aspects of a sustained vowel" can normally be recognised: fundamental pitch, loudness, voice quality (which he considered to depend primarily on the physiological structure of the vocal chords), and phonetic value (which he considered to depend mainly on the cavity shapes). This point of view is closely akin to that developed in the introductory chapter to this thesis; but the two should

not be confused. Peterson regards both voice quality and phonetic value as terms for physiological specifications. This viewpoint enables him to say in a later article: "The phonetic value of a speech sound, of course, is independent of language and meaning". (Peterson 1952). As opposed to this in the present thesis vowel quality is considered as having two aspects: personal quality which conveys information about which individual is speaking; and phonetic quality which is a term used to describe the aspects of a sound that a listener can interpret as conveying information in a socio-linguistic system. This phonetic quality is not considered to be independent of "meaning" (cf. Firth, 1957).

Peterson's experimental approach to the problem, however, is somewhat similar to that to be described in chapter 3 of this thesis. In his theoretical article (Peterson 1951) he suggests: "One approach to the study of phonetic value would be to select a group of vowels which had been judged to have essentially the same phonetic value, but which differed in the other three aspects: those physical factors or relationships which were found to be common to such a set of samples would then define their phonetic value". Later he put this procedure into practice, recording samples of the vowels /I/ and /æ/ in such a way, that "A group of matched vowels was obtained for men, women and children". As a result of his measurements of these vowels he states that "The front vowels could be rather readily identified by observing the positions in frequency of the peaks of the first three formants". But he does not state exactly what functions of the first three formant frequencies could be used as an objective measure of phonetic

values.

Peterson's work is valuable, but limited in that he provides data about only two vowel qualities. Another phonetician who has given prominence to the difficulty of separating phonetic and personal quality is Joos (1948); and his work though in some ways more valuable, is even more limited in that he presents hardly any evidence for his conclusions, claiming that the research was done "under conditions which forbade publication at the time (during the Second World War) and required leaving behind almost all data but what could be retained in the memory". (op.cit.)

Joos came to the conclusion: (which, considering the lack of data, is better regarded as a hypothesis) that the phonetic quality of a vowel depends on the relationship between the formant frequencies for that vowel and the formant frequencies of other vowels pronounced by that speaker. A necessary part of Joos's theory is that whenever a listener to speech has to identify a vowel without the benefit of any clues from the context, he utilizes whatever knowledge he has of the speaker's formant frequencies in other words. Even when the vowel which the listener is hearing is quite unlike any that he has ever heard that speaker produce before, he nevertheless focuses his attention not on the absolute values of the frequencies of the formants, but on the relations between those frequencies and the general ranges of frequencies which he expects to be characteristic of the speaker. Thus on this theory the phonetic value of a vowel depends on the way in which its acoustic structure fits into the pattern formed by the acoustic structure of other vowels produced by the same speaker.

Joos's theory is undoubtedly of paramount importance as will be seen in subsequent chapters of this thesis. But it should be

remembered that when originally put forward it was supported by very scanty evidence; and it should certainly not be regarded as proven for all vowels.

Chapter 3. Experiments involving cardinal vowels

In order to investigate the way in which vowel quality might be specified in instrumental terms, it is convenient to have a corpus of sounds which competent observers consider to be identical in phonetic quality. As we have noted, the vowels occurring in normal speech cannot be used for this purpose. There is little to be learnt about differences in vowel quality from measurements of the vowels in a number of English words spoken by different people, simply because words spoken by different people seldom contain vowels which are considered (by phoneticians) to be phonetically equivalent. I have never yet met any group of potential subjects who normally pronounced a set of words containing a number of phonologically different vowels in such a way that competent phoneticians consider that all the vowels in the corresponding lexical items have the same phonetic quality. The only vowel sounds which are always said to be identical in quality irrespective of the speaker are cardinal vowels.

As we saw in the previous chapter, there have in the past been various systems involving 'cardinal vowels'. But undoubtedly the best known at the present time is that devised by Daniel Jones. In this thesis, the phrase 'cardinal vowel(s)' is used as a technical term. Whenever the phrase occurs without further qualification it is intended to designate any sound(s) produced by Daniel Jones and stated by him to be cardinal vowel(s), or any sound(s) produced by any other speaker which are considered by competent observers to be equivalent in phonetic quality to the corresponding cardinal vowel(s) produced by Daniel Jones. By competent observers we mean phoneticians

who are thoroughly conversant with the exact quality of the original cardinal vowels as a result of prolonged training and instruction by Daniel Jones, or by one of his pupils.

It should be noted that in this thesis the phrase 'the cardinal vowels' is not used to designate any set of sounds produced as a result of merely following certain published specifications. Daniel Jones has made it quite clear that: "The cardinal vowels cannot be learnt from written descriptions; they should be learnt by oral instruction from a teacher who knows them". (Jones 1956). The implication of this is that the written descriptions are to some extent 'imitation labels' (Pike 1943). But this, of course, is irrelevant to the fact that the cardinal vowels are precisely determined vowel qualities.

This point of view is not understood by many American writers. The Haskins Laboratory group, for instance, say of one of their speakers - "a phonetician whose native language is French. Naturally his conception of the color of the cardinal vowels might differ from that of Daniel Jones". (Delattre et al. 1952). This seems to imply either that the vowel qualities of the cardinal vowels will depend on the native language of the speaker, or that they are intended to occupy areas rather than specific points in the vowel continuum; but, from our point of view, neither of these implications is correct. The same writers go on to say: "However we should guess that the differences, if any, are very small, since for 11 of the 16 vowels the I.P.A. offers French vowels as guides to pronunciation".(ibid) But Daniel Jones has expressly stated that it is possible to give only "some very rough indication of the values of the cardinal vowels by means of key words". (Jones 1956).

Other American writers display a similar lack of knowledge of the cardinal vowel system. Thus Hockett (1955) is obviously not very well aware of either the practical use or the theoretical background of the system. After having discussed the "traditional misunderstanding of 'IPA'" in equating phonological units with cardinal vowels (which of course, IPA phoneticians do not do) he goes on to state that it is not surprising that some of the French vowel phonemes (sic) are similar to cardinal vowels, because "the French vowel system played a major role in the development of the original "cardinal vowel" theory". (op.cit.)

However, these aberrations need not concern us unduly, since the more rigorous point of view outlined above is shared by most British phoneticians. Since in this thesis the cardinal vowels are considered to be not sets of sounds produced according to certain specifications, but sounds judged by competent observers to have certain phonetic qualities, it follows that sets of cardinal vowels pronounced by a reliable phonetician trained by Daniel Jones will be identical in phonetic quality, (but will probably not be identical in personal quality) with sets of cardinal vowels pronounced by any other reliable phonetician who has undergone similar training. As we have noted, they are the only sounds which have precisely determined qualities; consequently no other sounds are as useful in experimental investigations of what phoneticians mean by phonetic quality.

In January 1954 recordings were made of a number of sets of cardinal vowels spoken by different phoneticians on various pitches. The procedure was as follows. All the recordings were made in the studio of the Phonetics Department of the School of Oriental and African Studies, University of London. (Thanks are due to Professor J.R. Firth for placing

the studio at our disposal, and to Mr. N.C. Scott for making all the necessary arrangements). This studio is a pleasant medium-size room, in which it is possible to converse without being over-conscious of the surroundings. It is reasonably soundproof, but some noises and vibrations from other parts of the building can be heard occasionally.

The recordings were made on a Ferrograph tape recorder (Edinburgh University Phonetics Department, Ferrograph No. 6) which was checked before and after the experiment, and was found to have a flat frequency response from 60 - 9,000 cps ± 2 dbs. On some of the tapes made at S.O.A.S. there is a slight print through which just audible; but even considering this, there is a signal to noise ratio of better than 40 dbs.

During the recording sessions I sat alongside the recording machine in one corner of the room, Professor Daniel Jones sat at the other end of the room alongside a small table on which stood the microphone (a Reslo Ribbon). The phonetician who was being recorded sat in another chair alongside the table where he could converse with Professor Jones.

The following phoneticians took part in the experiment with Professor Jones: G.F. Arnold, J. Carnochan, Miss Chapellaz, D. Fry, A.C. Gimson, Miss Henderson, J.D. O'Conner, J. Pring, N.C. Scott, Miss Tooley, J.D. Trim, and Mrs. Whitley. At the time of the experiment each of these phoneticians had been on the staff of the Phonetics Department of either University College or S.O.A.S., London, for at least five years. Thus the performers in this experiment were all very experienced phoneticians who were fully conversant with the cardinal vowel system and were accustomed to specifying vowel qualities in a wide variety of

languages; and the ultimate judge of the quality of each vowel was Professor Daniel Jones.

Each of the subjects was first of all rehearsed by Professor Jones in the performance of the eight primary cardinal vowels, and then pronounced a number of sets of vowels on different pitches, each set being criticised and discussed by Professor Jones and by the subject. Recordings were made of each subject saying at least five sets of vowels; in addition a great deal of the comment and discussion was recorded. The original tapes (nine reels, each with 30 minutes of recording on the top track in the C.C.I.R. sense) are stored in the record library of the Phonetics Department of the University of Edinburgh.

One or two points of interest arose during the making of these recordings. In the first place it became apparent that, despite the published specifications, Professor Jones was considering the eight primary cardinal vowels as being formed of two relatively independent sets, the first four constituting one set, and the second four the other. Time and again he would make comments of the form "The first four were all right, but let's have the back set again". It was continuously noticeable that when a subject had any difficulty in correctly performing one of cardinal vowels two, three, six or seven, Professor Jones would usually rehearse his pronunciation not of that vowel in isolation, but of the vowel in the context of the appropriate sub-set. On the other hand cardinal vowels one, four, five and eight were often regarded as qualities which could be practised in isolation. Thus it seems that these vowels may have been regarded as the end points of two independent sub-sets. This is, of course, not suggested in any way by the original descriptions

of the cardinal vowels. Cardinal vowels one and eight are end-points in the series (although no use is made of this in defining the quality of cardinal eight), and cardinal five is defined as a fixed point in the series; but cardinal four is not defined as an end-point in any way. (It is interesting to note, however, that the late Ida Ward, a former pupil, colleague, and close associate of Professor Jones has published a description of the cardinal vowels in which it is stated that the tongue positions of cardinal vowels one, four, five and eight "Give the four 'corners' of the cardinal vowel figure". (Ward 1945).)

There were also indications that most of the subjects thought that both Professor Jones's performances of cardinal vowel number six, and the pronunciation of this vowel which he expected from them, were not in accord with their idea of the quality indicated by the published description of the cardinal vowel system. Nearly all the subjects had more difficulty with this vowel than with any of the others. Mr. Scott (who was one of the most senior of the subjects) verbalised what was apparently the common difficulty by saying "I feel I am coming forward when I please you more" (Tape 1, half way through). Later on Professor Jones and Mr. Scott had the following conversation which was recorded (Tape 2, near the end).

Professor Jones: What did you think of my number six?

Mr. Scott: Well ... as you see ... I'm not very capable of ...

Professor Jones: No, I know, but nor am I ..... I mean I've always had difficulty with it.

(Mr. Scott then goes on to say that it is as he remembers it on the record; and then restates his difficulties with his own number six, namely, that he feels he is coming forward when he pleases Professor Jones more).

Professor Jones: You don't think mine is a little . . forward do you?

Mr. Scott: Would you mind making it?

(Professor Jones says [a, ɔ, o, u]).

Mr. Scott: Well ..... I'm not sure it .... gives me something of a forward impression ..... which may be entirely my ear ..... wrong calibration.

None of the other phoneticians who took part in the experiment made any remarks of this kind to Professor Jones; but several of them, in private conversation with myself, said that they agreed with Mr. Scott's point of view.

On the nine original tapes there were a very large number of attempts at pronouncing sets of cardinal vowels (as well as a great deal of discussion). At the time when the recordings were made, Professor Jones noted that many of the sets were incorrectly performed, and should be disregarded. The remaining 92 sets (all of which Professor Jones at the time of their performance, had considered to be "Excellent", "model sets", "very good", "good", "fair", "reasonable", or of which similar words had been used) were copied onto another tape, henceforth referred to as "Tape A". The copy was made

by playing back the original recordings on the tape recorder on which they were made and re-recording the relevant parts on another Ferrograph with a frequency response of 60-9,000 cps ∓ 2dbs. As a result of the copying there were slight losses in the higher frequencies, but the copied recordings are within ∓ 1 db of the original recordings up to 6,000 cps.

Copy tape A was played back on 13 April 1954 at University College, London, to a group consisting of Professor Jones and most of the phoneticians who had been recorded. This meeting was not very profitable, the only outcome being that Mr. Pring declared that he was not satisfied with his own performance and accordingly he requested that the recordings of his cardinal vowels should not be considered further.

A subsequent visit in the following month to Professor Jones at Gerrard's Cross was far more profitable. On this occasion Professor Jones was able to listen to tape A and carefully consider each set of vowels in turn. His written comments are listed in table 3.1. It should be noted, however, that these comments cannot always be taken at their face value. They must often be considered in conjunction with the set to which they are applied. Thus Gimson, set 7 is noted as "g" (i.e. good); but it was apparent from Professor Jones's additional oral comments that he meant "good - for a set pronounced in falsetto".

Table 3.1 Cardinal Vowel Tape A

	Comments by Prof.D.Jones	Additional notes by PL	Prelim Analysis	Final Analysis
ection 1				
iss Henderson				
Set 1	ɔ? otherwise good (ɔ not far enough from o).	clicks		
1a	on A ɔ↓? a-? e↓ ?		X	
2	g. ɔ OK	falsetto		
3	ɛ↓ ɔ OK	clicks, not good recording		
4	v.g. except for wobble on ɔ	Wow effect, quaver in voice	X	X
5	i↓ ɛ↓? ɔ OK			
rofessor D. Jones				
Set 1	g. a flat D (Nearly C#)	clicks on [e] [u]		
2	C. not v. good (near may lower limit for this)			
3	Flat E g.		X	X
4	F slightly flat f.		X	

	Comments by Professor D.Jones	Additional notes by PL	Prelim Analysis	Final Analysis
Mr. Carnochan				
Set 1	g	flat second 4 vowels	X	X
2	g ɔ↑	print through	X	
3	g OK	a - ɑ break	X	X
4	g	bass	X	X
5	F/fairly g	bass	X	
6	a↓	high no [i]		
7	c g	high	X	
Mrs. Whitley				
Set 1	v.g. i on A, others on A↓		X	
1a	g.		X	
2	ɔ otherwise g.	faint [u]		
3	e↑ otherwise v.g.			
4	v.g. ɔ v.g.		X	

		Comments by Prof. D. Jones	Additional notes by PL	Prelim Analysis	Final Analysis
Mr. O'Connor					
Set 1		g. ɛɹ ɔɹ			
	2	g. ɛɹ ɔɹ			
	3	g. ɔɹ oɹ			
	4	g. ɔ OK oɹ			
	5	g. o OK I think		X	X
	6	g. oɹ			
	7	v.g.		X	X
	8	v.g.	bass	X	X
	9	g. (4 missing)	bass, no [a]		
Mr. Scott					
Set 1		v.g.		X	
	2	v.g.	[ɔ] partially erased	X	
	3	v.g.	models	X	X
	4	v.g.		X	X
	5	v.g.		X	X
	6	v.g.	bass	X	
	7	g.			

		Comments by Prof. D. Jones	Additional notes by PL	Prelim Analysis	Final Analysis
Section 2					
r. Trim.Set	1	v.g. model		X	X
	2	g. ɛ↓		X	X
	3	v.g. 8 missing			
	4	g.		X	X
	5	ɔ↓ u↑			
	6	g u↑		X	
	7	g		X	
	8	v.g.	print through	X	
	9	g		X	
	10		falsetto		
Miss Chapallaz					
set	1	fair i↑ o↑			
	2	fair ɔ↑			
	3	g i↑ u↓			
	4	i ɪ↑ ɔ↑ not v. clear			
Miss Tooley					
Set	1	fair i↑ o↓ u?			
	2	g exc. for o↓	[ɛ] wavery		
	3	g up to ɔ, o u weak (poor recording)	bad recording of back vowels fading	X	X
	4	g weak u	fading on [u]	X	X
	5	ɛ↑ ɔ o u not clear	too high pitch		

		Comments by Prof.D. Jones	Additional notes by PL	Prelim Analysis	Final Analysis
Mr. Arnold					
Set	1	g. ɔ ɪ	Print through Noise on [i]	X	X
	2	g. ɔ ʔ ə u +			
	3	v.g.		X	X
	4	v.g.		X	X
	5	g. exc. for ɔ ə			
	6	g. exc. for ɔ ə			
	7	not v.g.	falsetto, noisy		
Mr. Gimson					
Set	1	v.g.	clicks [ɛ a]	X	X
	2	v.g.		X	X
	3	v.g.		X	X
	4	v.g. ɑ + ?		X	
	5	g exc. for ɑ ə	low (bass)		
	6	g.		X	
	7	g.	falsetto		

	Comments by Professor D.Jones	Additional Notes by PL	Prelim Analysis	Final Analysis
Mr. Pring				
Set 1	g. a^3 ?	request all sets not to be considered further		
2	ε τ			
3	Hesit^{ns}			
4	good a "			
5	g. except for e τ			
Dr. Fry				
Set 1	g. sulcaliz^{n}		X	
2	g.		X	
3	g. sulcaliz^{n} esp. noticeable in No.7		X	X
4	g.		X	
5	g.		X	
6	g.		X	X
7	v.g.	bass	X	X
8	g. No.1 missing	bass, lower limit		
9	v. sulcal		X	

		Comments by Prof. D. Jones	Additional notes by PL	Prelim Analysis	Final Analysis
Professor D. Jones					
	Set 5	g. husky		X	
	6	g.		X	X
	7	v.g.		X	X
	8	g.		X	
	9	falsetto			
	10	husky			
	11	husky and wobbly			

The 53 sets on Tape A noted in the column headed "Preliminary analysis" were analysed on a Sound Spectrograph. However this analysis was not made with sufficient care to permit accurate measuring and, although some useful preliminary insights into the nature of vowel quality were obtained, the results had to be discarded. The 31 sets noted in the column headed "Final analysis" were also analysed with the aid of a 48 channel spectrograph (Fant, 1958c). This analysis was also not suitable for accurate measuring.

The final acoustic analysis was, however, made of these same 31 sets. These sets were selected so that in the tape for final detailed analysis, there would be for each of the eleven phoneticians at least two, and, if possible, three sets of vowels which were considered by Professor Jones to be good complete sets not in the extreme pitch ranges.

The final acoustic analysis of all the vowels in these 31 sets was made with the aid of a modified Sonagraph (a type of sound spectrograph manufactured by the Kay Electric Co.). From the point of view of the current research, the most important modifications were:

(1) As well as the normal Sonagraph frequency scale, in which 1" = 2,000 cps, two additional scales were added. In the one (called the expanded frequency scale) 1" = 1000 cps, and in the other (the ultra-expanded frequency scale) 1" = 500 cps. A new switch enabled any of the three scales to be selected. The expanded and the ultra-expanded scales are particularly useful for making accurate frequency measurements of components in the lower part of the spectrum.

(2) The potentiometer (P6) which controls the position of the base line was brought to the front panel. Thus the base line position could be easily adjusted for each spectrogram. When examining low frequency components (below about 350 cps) the base line was raised so that it was well above the bottom of the display. In this condition a mirror image of the lower part of the display is formed below the base line.

(3) An extra switch and resistors were put in which enable the base line to be raised about $1\frac{3}{4}$", i.e., to a point about half way up the paper. Using this switch it is possible to make two spectrograms one above the other.

(4) As an aid in using the three modifications described above, a small neon striking at 90 volts was connected to the stylus. This neon lit up whenever there was a voltage on the stylus sufficient to mark the paper; in the case of the base line such a voltage occurs even when the drum is not revolving. Consequently, in order to find or alter the position of the base line, the stylus could be put in the required position and the potentiometer mentioned in (2) above adjusted until the neon lit.

(5) A calibrating unit was devised on the model of that described by Peterson and Barney (1952). This unit was constructed by Mr. J. Anthony. An Advance

Signal Generator was used in conjunction with the unit, which changed the square wave output of the Advance into a series of short duration pulses. A wave form of this latter kind consists of the fundamental and many harmonics with nearly the same amplitude. The output of the calibrating unit was fed into the Sonagraph record amplifier immediately after each item which was to be analysed. Thus every spectrogram was individually calibrated. In the current work the frequencies in the calibrating sound were always part of the series 400, 800, 1200, 1600, 2000, 2400, 2800, 3200, 3600 and 4000 cps.

With the aid of this modified Sonagraph the frequency of a point on the display could be determined with an accuracy of ∓ 2.5% at the lower end of the ultra-expanded scale (i.e., ∓ 5 cps at 200 cps) and ∓ 1% at the upper end of the expanded scale (i.e. ∓ 40 cps at 4000 cps).

The main object of the acoustic analyses was to measure the frequency of each of the first three formants for all the 248 vowels. At least two and often three or four spectrograms were made for each vowel. It was found that wide band spectrograms (i.e, the kind of display made using an analysing filter with a bandwidth of 300 cycles, in which the component frequencies of a sound are shown as a function of time, the intensity of the components being shown by the degree of blackness of the trace) were usually the most convenient way of locating formants and measuring the centre frequencies (i.e., the centre of the darkest area in the appropriate region). This kind of display enables the

investigator to examine a considerable length of each sound, and thus to see what features of the acoustic pattern are fairly constant throughout the duration of the sound. The first formants of [i, e, ɛ, a] and the first and second formants of [ɑ, ɔ, o, u] were usually measured using this kind of display made with an ultra-expanded frequency scale (1" = 500 cps); and the second and third formants of [i, e, ɛ, a] and the third formants of [ɑ, ɔ, o, u] were usually measured using this kind of display made with an expanded frequency scale (1" = 1000 cps). Narrow band sections (i.e., the kind of display made using an analysing filter with a bandwidth of 45 cycles, in which the frequencies of the components of a sound at a given instant are shown as a function of their intensities) were also often made; they were especially helpful in showing whether a formant could be said to be present or not (i.e, for the purposes of this study, whether there were any components within the appropriate frequency ranges which were within 35 dbs of the component which had the highest amplitude in that sound.) Ultra-expanded narrow band spectrograms were used for determining the fundamental frequency of each set of vowels.

So far, no operational definition of a formant has been given. There are many definitions in the literature, but most of them are not very precise. Thus Ladefoged and Broadbent (1957) talk of "formants or regions of the auditory spectrum in which there is a relatively large amount of spectral energy". And somewhat earlier in the current upsurge of acoustic phonetic research, Dunn (1950) says that: "... the different vowels have associated with them different frequency regions, in which the sound is more intense than elsewhere in the spectrum.

The name "formant" has been applied to these regions ....." But Potter and Steinberg (1950), who also begin with a similar specification, go on to give "a more precise meaning to frequency or formant position". They say that they do not have a complete answer, but suggest that "the ear deals with something akin to effective pitch centers of loudness of the energy concentrations ..... Expressed formally, the central or formant frequency is given by the following relation:

$$F = \frac{\sum w_i f_i}{\sum w_i}$$

where f_i = frequency of the i th component, and w_i = a weighting factor ... Ai/Ao where Ai is the amplitude of the ith component, and Ao is the amplitude of the dominant or maximum component". This is a somewhat long but precise definition of a formant frequency. Nevertheless, as the authors themselves admit, it is not always easy to apply in practice. Other investigators (Peterson and Barney 1952, Fant 1958a) have also found it difficult to locate the centre of a formant.

In the course of the examination of the current spectrographic data the following difficulties were encountered:

1. The centre of formant one is difficult to locate when it is low in frequency. In all vowels there is usually a great deal of energy at the fundamental frequency. When a formant is within one and a half octaves of the fundamental it is more difficult to specify its centre frequency, because in these circumstances usually only the frequency components



Figure 3.1 Trim, set 1, [u]; expanded scale narrow band sections.

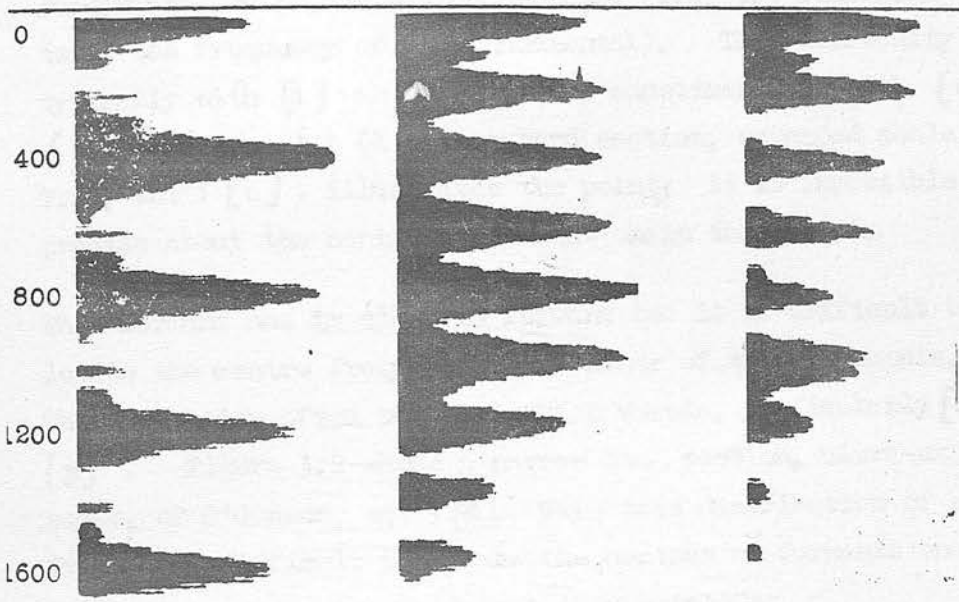


Figure 3.2 O'Connor, set 7, [a]; ultra-expanded scale narrow band sections

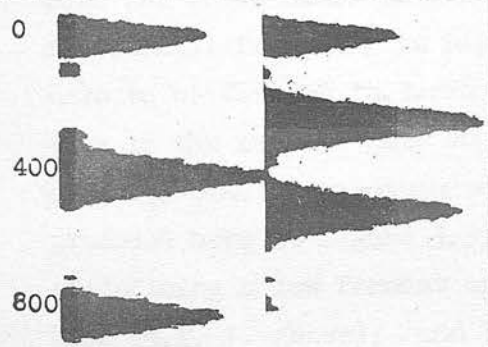


Figure 3.3 Tooley, set 4, [e]; ultra-expanded scale narrow band section

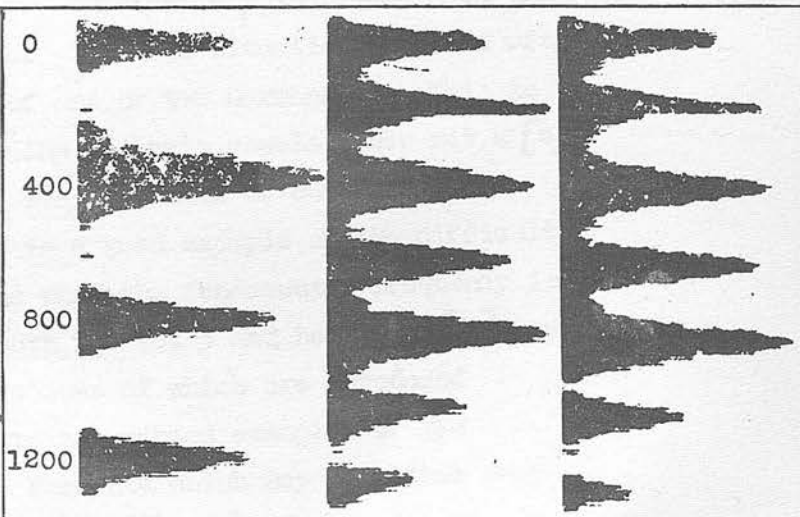


Figure 3.4 Tooley, set 3, [a]; ultra-expanded scale narrow band sections

higher than the assumed centre frequency decrease in amplitude; there is thus no peak in the spectrum, and it is difficult to know whether to specify the formant centre as being nearer to the fundamental or to the second harmonic (i.e, the component with twice the frequency of the fundamental). This difficulty occurs typically with [i] and [u] ; but sometimes also with [e] and [o] . Figure 3.1 (a narrow band section, expanded scale, of Trim, set 1 [u]) illustrates the point; it is impossible to be precise about the centre of formant one in this vowel.

2. When formant one is close to formant two it is difficult to locate the centre frequencies of either of these formants. This situation often occurs in back vowels, particularly [a] and [ɔ] . Figure 3.2 shows a narrow band section, ultra-expanded scale, of O'Connor, set 7 [a]. With this distribution of energy it is very difficult to locate the centres of formants one and two, unless it is presumed that they coincide.

3. Both the above difficulties are considerably increased when the fundamental frequency is high. In this situation formants often have to be defined in terms of one or two harmonics. This is true in the case of many of Miss Tooley's vowels; her set 4 [e] (a narrow band ultra-expanded scale section of which is reproduced here as figure 3.3) is a good example of the difficulty of defining a low formant one when the fundamental frequency is high (c.f. 1. above); and both her set 3 and her set 4 [a] (narrow band ultra-expanded scale sections of which are reproduced here as figure 3.4) illustrate aggravated examples of the difficulty of specifying two formants which may be either very close together or coincident (c.f. 2 above).

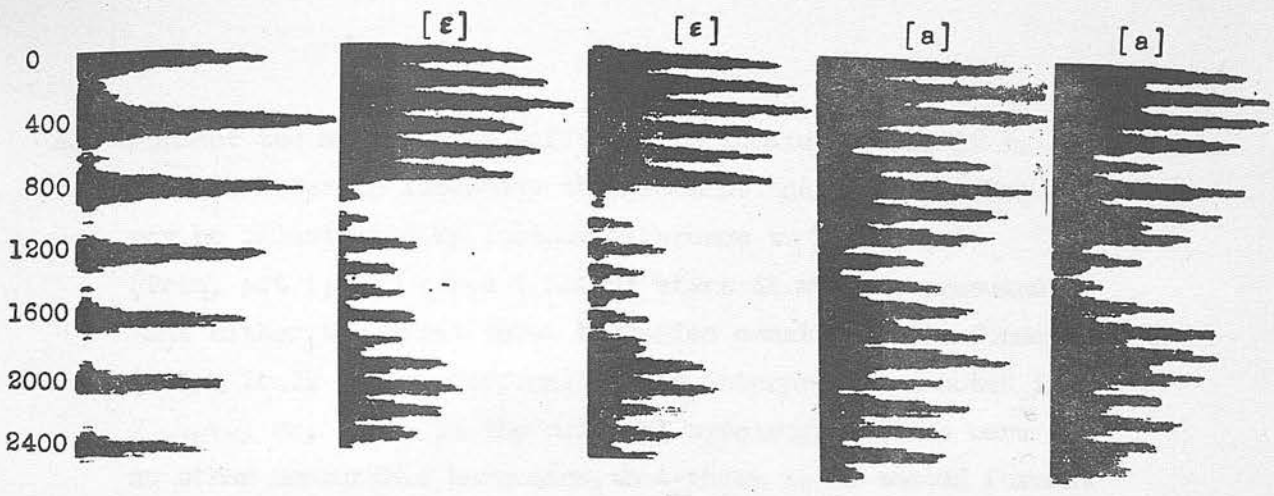


Figure 3.5 Jones, set 6, [ɛ] and [a]; expanded scale narrow band sections

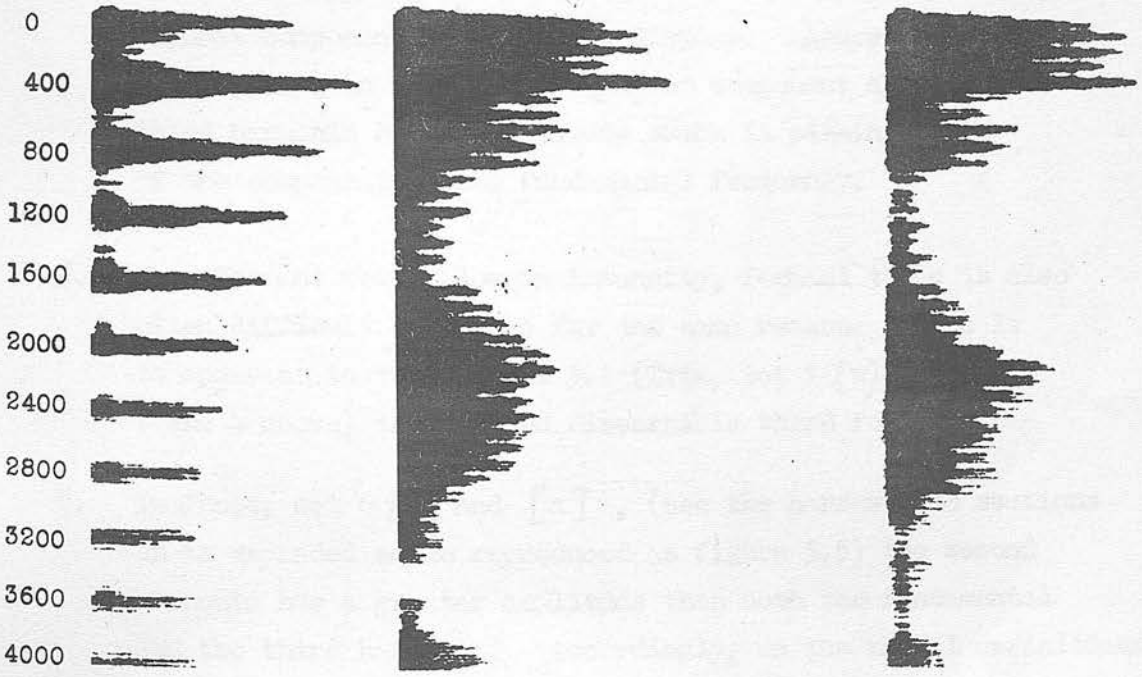


Figure 3.6 O'Connor, set 8, [e]; expanded scale narrow band sections

4. Formant two may also be difficult to locate because it is so much lower in intensity than formant one. This point may be illustrated by further reference to figure 3.1 (Trim, set 1, [u] - see 1 above) where it must be presumed that either the first three harmonics constitute two formants (which leads to the difficulties in interpretation noted in 2 above) or, since in the original spectrogram there were no other measurable harmonics, that there is no second formant with sufficient intensity for its frequency location to be measured. All the narrow band sections reproduced in this thesis have an intensity scale such that 1 mm = 1 db, and a frequency response which is level (∓ 1 db) from below the component with the lowest frequency to above the highest component shown on the display. Accordingly we can say that in Trim, set 1 [u] no component above the third harmonic has an amplitude which is within 35 dbs of the component at the fundamental frequency.

5. When formant two is low in intensity, formant three is also often difficult to locate for the same reason. Thus it is apparent that in figure 3.1 (Trim, set 1 [u]- see 1 and 4 above) there is no discernible third formant.

6. In Jones, set 6 [ɛ] and [a] , (see the narrow band sections on an expanded scale reproduced as figure 3.5) the second harmonic has a greater amplitude than both the fundamental and the third harmonic. Accordingly, on the normal definitions of a formant, there are formant peaks in these regions. But these peaks are in very unusual positions for the first formants in open vowels; they are, in fact, rather close to the first formants in [i] and [u] . Moreover there

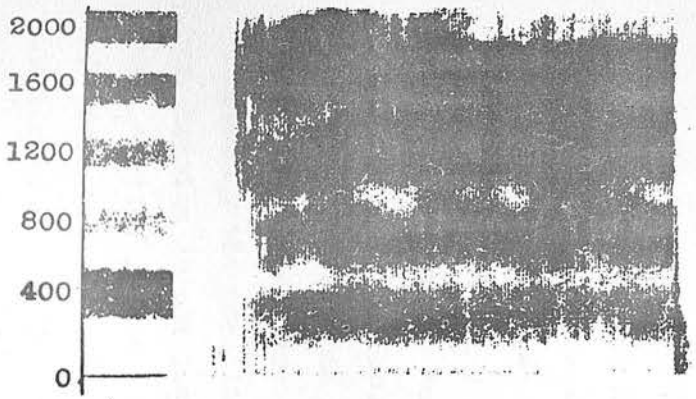


Figure 3.7 Whitley, set 4, [a]; expanded scale wide band spectrogram.

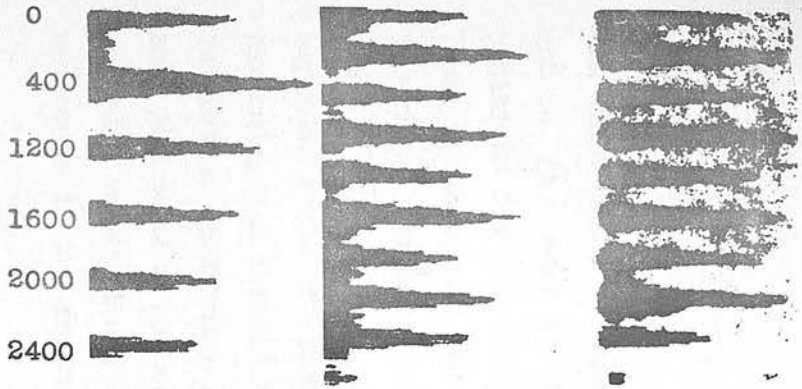


Figure 3.8 Whitley, set 4, [a]; expanded scale narrow band sections.

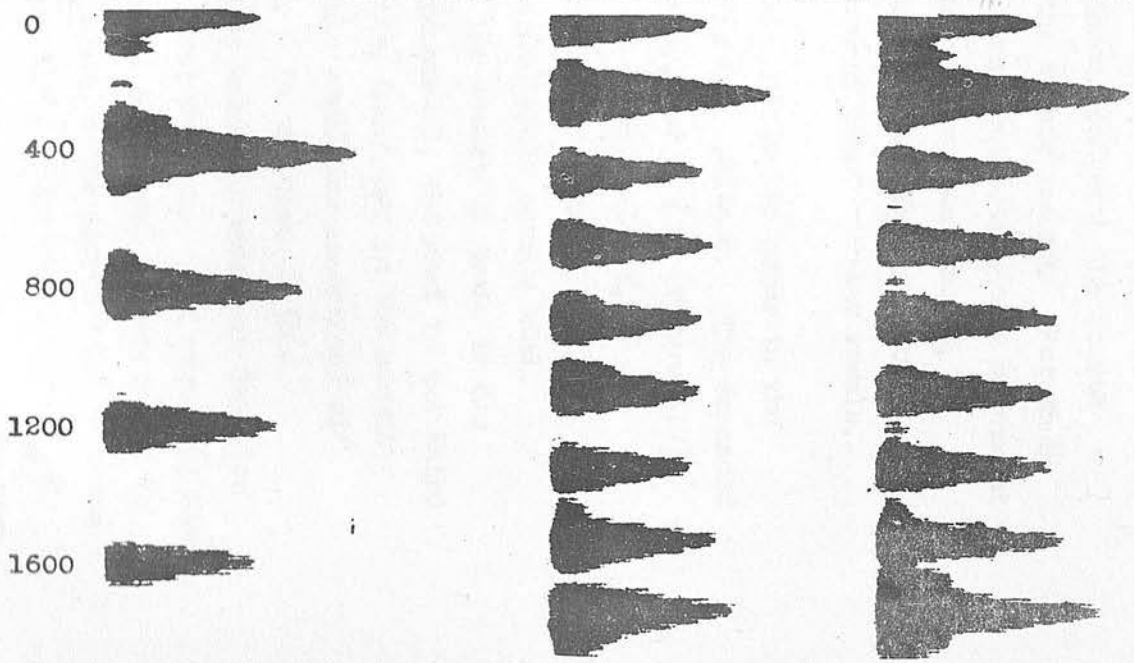


Figure 3.9 Whitley, set 1, [a]; ultra-expanded scale narrow band sections.

are other peaks of energy (in [ɛ] near the fourth harmonic and in [a] near the fifth harmonic) which are in more usual positions for these vowels. It would seem, therefore, that the lower peaks should not be considered as formants. Similar "spurious" formants occur in other vowels. Thus in O'Connor set 8 [e] (see figure 3.6 for narrow band sections on an expanded scale) formant one is associated with the fourth harmonic, and formant two with the twenty second. But there is another peak which might well have been defined as a formant near the seventh harmonic. This peak can be ignored, or regarded as spurious, only through knowing that it does not occur in the formant structure of other similar vowels.

More extreme examples of this difficulty occur in the analyses of [a] as pronounced by Mrs. Whitley. The expanded scale wide band analysis of Whitley, set 4, [a] (figure 3.7) shows that there are four bands of energy below 2,000 cps (i.e., in the region in which there are normally only two formants). The associated expanded scale narrow band sections (figure 3.8) show that the lowest of these is due to the high intensity of the fundamental, the next to the third harmonic, while the next two peaks, which are in the normal positions for the formants of this vowel, are associated with the fifth and seventh harmonics. In the case of this vowel it is therefore possible to make a procrustean decision concerning the position of the first formant. But even this kind of decision is not possible for the same vowel as pronounced by Mrs. Whitley in set 1 (ultra-expanded scale narrow band sections of which are reproduced in figure 3.9). The first six harmonics cover the frequency region up to 1,300 cps. Somewhere in that region, probably near the fourth or fifth harmonic, there

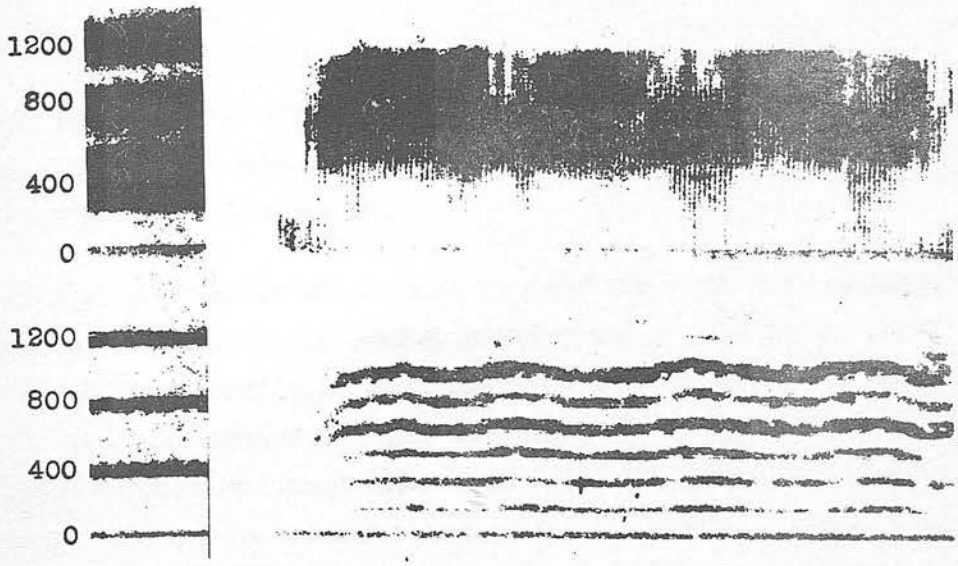


Figure 3.10 Scott, set 4, [a]; expanded scale wide band (upper display) and narrow band (lower display) sonagrams.

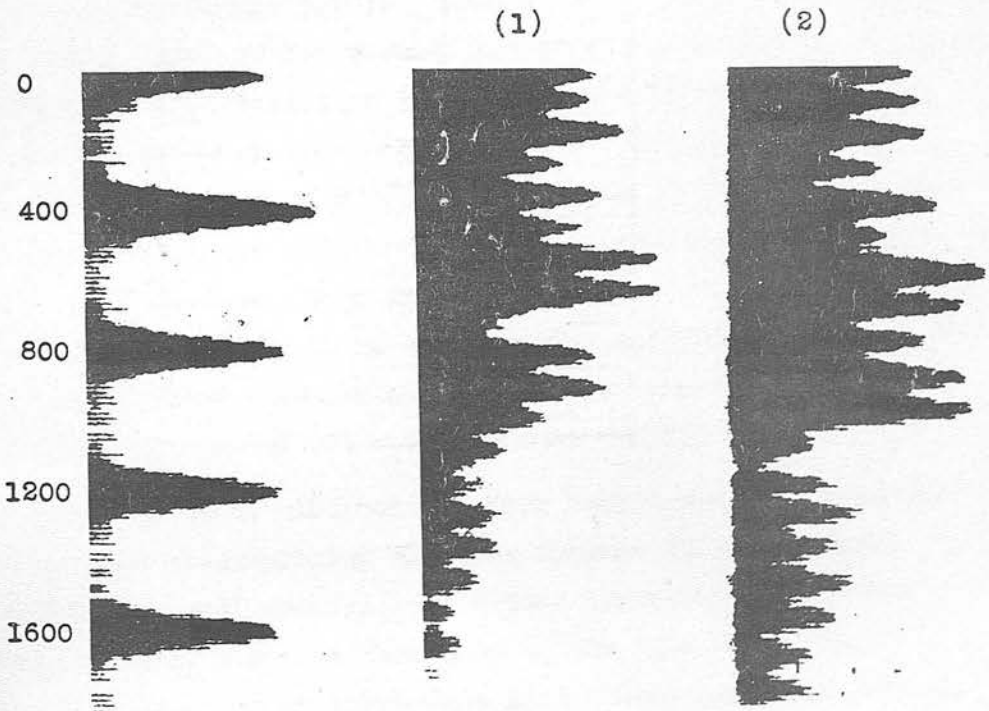


Figure 3.10a Scott, set 4, [a]; ultra-expanded scale narrow band sections.

would be for other similar vowels a first formant. But for this particular vowel there seems to be no way of deciding exactly where the first formant is.

7. The final difficulty which must be considered is that although these are theoretically steady state vowels, they do in fact often vary in quality. Most of these variations are so small as to be imperceptible. But sometimes there are quite considerable audible variations which have been considered irrelevant (or for some reason have not been attended to) by Professor Jones when assessing the quality of the vowels. An example is Scott, set 4 [ɑ] (see the expanded scale wide band and narrow band analyses reproduced in figure 3.10). During the production of this vowel there was a considerable variation in the breath force. This is associated with a variation in both the fundamental frequency, and the intensity of the second formant. Analyses of this vowel made at different moments would yield different results. These variations are illustrated by the ultra-expanded narrow band sections reproduced in figure 3.10a. In section (1) the fundamental frequency is 92 cps, and the second formant peak is at 880 cps, whereas the values derived from section (2) are: fundamental frequency 100 cps; formant two 950 cps.

When all these difficulties have been taken into account, it seems almost surprising that the formant theory of vowel quality is so well established. This impression is further substantiated by the fact that much of the data about the formant frequencies of vowels has little auditory significance; it is most unlikely that a formant of comparatively low intensity (e.g. F2 in [i] in any of the Jones sets, or F3 in [u] as pronounced by most subjects) contributes anything to the auditorily

perceived quality. There are, however, three reasons why we should continue to describe vowels in terms of their formants. Firstly, when we consider spectrographic analyses of an individual's conversational utterances, a pattern develops and it is comparatively easy to pick out the movements of the formants from one sound to another; consequently in these circumstances specifications in terms of formants make a convenient way of characterising the pattern. Secondly it has been claimed (Fant 1956) that specification of the formant frequencies is sufficient to enable predictions to be made of the formant levels, and also of the spectrum envelopes; so, irrespective of the auditory effect (or lack of it) of a particular formant, specifying its frequency and that of the other formants completely may characterise the sound. Although the formant frequencies may be hard to locate, once they have been found they might provide the best simple specification of the sound. Thirdly, as has been shown by Lawrence (1953), if a speech synthesiser is programmed in terms of formant frequencies and other features of the acoustic characteristics of speech, it may produce intelligible speech of a reasonable quality. This is, of course, an excellent pragmatic reason for regarding the formant frequencies as among the most important features of the acoustic characteristics of speech.

However, the lack of a definition of a formant remains; and very probably it is not possible to provide a single definition which is of use to the investigator who is examining spectrographic data (as opposed to a definition which is useful to the research worker, synthesising speech or the physiologist/acoustician who can specify the properties of the vocal tract). As Fant (1958a) says "A single formula, for instance, providing

a center of gravity measure of a formant from all harmonics of any importance, is not recommendable since the systematic differences between various types of formants are too large". In other words, unless one knows, at a particular moment, which "type of formant" one is dealing with, it is impossible to know how to handle the acoustic data. Indeed, on the basis of the experience of analysing all the acoustic data discussed in this thesis, it would seem that without knowing the sound which is being investigated, and without some previous knowledge of approximately where the formants of such a sound might be expected to be, it is impossible to make valid measurements of the formant frequencies. The knowledge of approximately where the formants might be expected to be is partly the result of the experience of examining many similar sounds and is partly derived from experience in synthesising sounds. Thus we know that in figure 3.7 (Whitley, set 4 [a]) the formants are in the places indicated partly because these places are the ones which are in the usual ranges for the frmants of [a] , and partly because if we tried to synthesise a sound as similar as possible to this one, it would have to have something like these formant frequencies. We cannot say exactly what the formant frequencies would have to be in order to make the best synthetic match to the original sound, because we have no procedure for testing when such a match has been achieved. At the moment our synthetic vowels are not indistinguishable from real ones; and until we know more about the reasons for this we cannot get phoneticians to make up their minds as to which of two similar synthetic sounds is most like a somewhat different natural sound. Most phoneticians will declare which of two synthetic sounds is the better match to a natural sound only when the two synthetic sounds are widely different in character (e.g.,

one with formants in the correct positions derived from figure 3.7, and the other with formants in the spurious positions).

In summary, therefore, the procedure for determining formant frequencies used in the current research consists of (1) listening to the sound and estimating from experience of analysing and synthesising similar sounds, the possible parts of the spectrum in which the formants might be located (2) examining spectrographic analyses and finding the centre frequencies of the regions within those parts that have a relatively high intensity. The essential weakness of this procedure is its circularity - the necessity of having to prejudge the answer before examining the acoustic data. This fault is especially serious in some of the judgments of the third formant, concerning the position of which there is at present no agreed data. However, the procedure adopted is, in the present state of our knowledge, the only possible way of overcoming the difficulties in locating formant frequencies which have been noted above; no other procedure has been formulated which is equally useful in the analysis of spectrographic data.

The results of this analysis of the 31 sets of cardinal vowels are given in table 3.2. A question mark indicates that there was no ascertainable peak (within the expected range) which was less than 35 dbs below the largest peak in that sound. A dash (e.g., in Mrs. Whitley set 1 [a]) indicates that the formant centre could not be found for some other reason (see page 44). On many occasions for the vowel [a] the only possible interpretation of the data in accord with the procedure stated above is that

formant one has the same frequency as formant two; accordingly there is often only a single figure in these two formant columns for this vowel.

The data in table 3.2 show different frequencies. But it is known that the perception of equal intervals of pitch cannot be exactly correlated with either equal intervals or equal ratios of frequencies; and it seems probable that in the same way the perception of quality differences is not simply related either to the formant frequency intervals or to the formant frequency ratios. The unit for the measurement of pitch is the mel, and the revised mel scale devised by Stevens and Volkman (1940) has been shown by its authors and others (Koenig 1949, Munson and Gardner 1950, Fant 1958a) to have a wide applicability. Accordingly the data of table 3.2 have been converted into mels as shown in table 3.3. This conversion was achieved with the aid of a graph showing the relation between frequencies and mels [illegible] from the data of Beranek (1949).

Table 3.2

Subject	Set	ə				ɛ			æ			ɑ		
		F0	F1	F2	F3	F1	F2	F3	F1	F2	F3	F1	F2	F3
Arnold	1	176	220	2490	3330	345	2320	2810	665	1860	2540	945	1680	2725
	3	134	295	2370	3200	365	2150	2470	555	1810	2470	900	1630	2640
	4	146	235	2420	3380	365	2240	2600	575	1840	2540	905	1550	2630
Carnochan	1	163	180	2150	2800	354	2110	2730	600	1860	2520	835	1440	2360
	3	118	200	2400	3720	320	2090	2710	610	1850	2340	890	1660	3060
	4	95	275	2200	3200	360	2180	2880	570	1850	2540	875	1480	2340
Fry	3	158	325	2270	3170	400	2125	2710	650	1875	2525	765	1440	2680
	6	195	230	2370	3270	365	2130	2720	575	1910	2550	725	1530	2650
	7	85	370	2210	2980	430	2150	2775	620	1960	2450	880	1510	1690
Gimson	1	130	285	2535	3550	400	2240	2980	660	2150	2800	995	1690	2690
	2	145	275	2440	3540	410	2255	3000	665	1980	2750	1010	1690	1980
	3	165	325	2370	2930	400	2150	2800	620	2000	2770	890	1880	2700
Henderson	1a	240	240	2480	3340	435	2560	2960	535	2150	3200	885	1600	3270
	4	320	320	2460	3350	340	2530	?	610	2590	3230	1220	1790	2800
Jones	3	145	190	2350	3550	345	2110	2600	490	1850	3400	905	1640	2100
	6	145	245	2480	3650	315	2250	2750	615	2110	3550	840	1750	2280
	7	165	190	2475	3600	305	2200	2720	645	1910	3550	925	1800	2250

a			ɔ			o			u		
F1	F2	F3	F1	F2	F3	F1	F2	F3	F1	F2	F3
770	945	2610	635	?	2750	365	675	2955	220	?	3230
795		2640	490	?	2740	355	?	3000	310	?	3330
880		2630	540	?	2550	395	?	3040	295	?	3280
680	970	2520	450	765	2370	350	760	2330	260	?	?
820		2820	520	870	2600	340	660	3550	210	?	3640
675	1030	2440	535	950	2360	330	700	2400	320	?	3600
770		2850	475	?	2640	325	?	2580	210	?	2370
715		2850	540	?	2650	335	?	2500	265	?	2300
685	938	2720	515	725	2610	400	600	2400	290	600	2400
740	1025	2930	485	850	2820	430	800	2540	235	?	3450
690	1035	2880	590	910	2900	380	785	2660	285	?	3500
670	990	2920	550	850	2900	490	850	2780	275	?	3080
625	965	3250	630	945	3240	390	715	3160	240	655	?
875		3120	595	900	3330	355	620	2820	320	?	2600
800		2740	550	?	2050	330	?	3550	170	?	3600
880		2250	625	?	2100	380	870	3200	240	?	?
720		2800	640	?	2450	320	660	3500	190	650	650

Subject	Set	i				e			ɛ			a		
		F0	F1	F2	F3	F1	F2	F3	F1	F2	F3	F1	F2	F3
O'Connor	5	187	235	2250	2980	540	2190	2700	670	1920	2710	1005	1680	2700
	7	200	235	2260	3240	400	2130	2600	595	1910	2600	970	1690	2920
	8	94	190	2220	2900	380	2100	2400	565	1820	2430	955	1590	2350
Scott	3	154	230	2490	3250	345	2150	2700	630	1840	2350	875	1650	2650
	4	172	265	2500	3300	330	2250	2690	670	1910	2520	800	1600	2520
	5	100	330	2420	3230	405	2240	2725	690	1880	2350	855	1580	2530
Tooley	3	220	245	2760	4050	445	2600	3300	625	2320	3150	1070	1700	3090
	4	248	250	2875	3630	375	2470	3110	640	2340	3150	1170	1750	3170
Trim	1	198	210	2380	3780	365	2250	2820	585	2075	2675	970	1750	2800
	2	200	285	2490	3470	360	2250	2760	685	2100	2635	825	1620	2625
	4	142	305	2410	3250	435	2300	2950	675	2040	2560	880	1800	2600
Whitley	1	217	220	2475	3700	420	2400	2825	615	1900	2900	-	1720	3110
	1a	200	260	2450	3630	390	2390	2800	590	2120	2800	1100	1620	3250
	4	257	265	2500	3580	400	2480	2950	685	2200	2900	1180	1700	3050

a			ɔ			o			u		
F1	F2	F3	F1	F2	F3	F1	F2	F3	F1	F2	F3
	900	2550	550	830	2370	500	820	2330	220	?	?
	880	2550	510	975	2450	410	770	2280	250	?	2340
650	840	2660	455	705	2600	380	690	2310	165	555	?
630	970	3220	630	920	2940	305	?	2690	250	?	?
675	950	3000	545	800	2830	345	660	2670	220	435	?
640	930	3000	485	775	2980	390	650	2620	250	?	3175
	-	3360	325	785	3060	400	?	3700	-230	?	3960
	1000	3360	375	875	3120	435	?	3740	270	?	3850
	690	3000	550	?	2830	380	?	3250	285	?	3600
	700	2875	500	?	2775	400	?	2775	320	?	3575
	845	2860	470	850	2820	400	650	3320	260	725	3640
640	1040	3440	400	?	3225	350	850	3230	235	651	3760
560	1025	3840	495	?	3120	340	?	3160	220	?	3920
680	980	3400	495	820	3420	495	?	3380	260	800	3700

Table 3.3

Subject	Set	i				e			ɛ			a		
		F0	F1	F2	F3	F1	F2	F3	F1	F2	F3	F1	F2	F3
Arnold	1	270	325	1770	2070	455	1700	1900	750	1485	1795	960	1395	1865
	3	215	405	1720	2030	475	1620	1760	655	1460	1760	930	1370	1835
	4	230	340	1745	2085	475	1660	1820	670	1470	1795	935	1320	1830
Carnochan	1	255	275	1620	1895	455	1605	1870	695	1485	1905	885	1260	1720
	3	190	300	1735	2180	430	1595	1860	700	1480	1710	920	1380	1985
	4	155	385	1645	2030	470	1635	1925	665	1480	1795	910	1285	1710
Fry	3	250	435	1675	2020	505	1610	1860	735	1490	1785	825	1260	1850
	6	295	335	1720	2050	475	1615	1865	670	1510	1800	795	1315	1840
	7	135	480	1650	1960	535	1620	1885	710	1530	1750	915	1305	1855
Gimson	1	205	395	1790	2135	505	1660	1960	745	1620	1895	995	1400	1855
	2	230	385	1750	2130	515	1670	1965	750	1545	1875	1010	1400	1850
	3	260	435	1720	1940	505	1620	1895	710	1550	1885	920	1490	1860
Henderson	1a	350	345	1765	2075	540	1805	1955	635	1620	2030	920	1350	2050
	4	430	430	1755	2080	450	1790		700	1815	2040	1140	1450	1895
Jones	3	230	290	1715	2135	455	1605	1820	590	1480	2090	935	1375	1690
	6	230	355	1765	2160	425	1670	1880	705	1605	2135	885	1430	1680
	7	260	290	1760	2145	415	1645	1865	730	1510	2135	945	1455	1670

a			ɔ			o			u		
F1	F2	F3	F1	F2	F3	F1	F2	F3	F1	F2	F3
830	960	1825	720		1875	475	755	1950	325		2040
850		1835	590		1875	465		1965	420		2070
915		1830	640		1800	500		1980	405		2055
760	980	1785	555	825	1720	460	825	1705	370		
870		1900	620	910	1820	450	745	2135	315		2160
755	1020	1750	635	965	1720	440	775	1735	430		2145
830		1910	580		1835	435		1810	315		1720
790		1910	640		1840	445		1775	370		1690
765	960	1865	615	795	1825	505	695	1735	400	695	1735
810	1015	1940	590	895	1900	535	855	1795	340		2100
770	1020	1925	685	935	1930	485	845	1845	395		2115
750	990	1940	650	895	1930	590	895	1890	385		1990
715	975	2045	715	960	2040	495	790	2015	350	740	
910		2000	690	930	2070	465	710	1900	430		1820
855		1875	650		1575	440		2135	265		2145
915		1670	715		1600	485	910	2030	350		
790		1775	725		1750	430	745	2115	290	735	2160

Subject	Set	i				e			æ			a		
		F0	F1	F2	F3	F1	F2	F3	F1	F2	F3	F1	F2	F3
O'Connor	5	290	340	1670	1960	640	1635	1860	750	1515	1860	1005	1395	1860
	7	300	340	1670	2040	505	1615	1820	690	1510	1820	980	1400	1940
	8	155	240	1650	1930	485	1600	1735	660	1465	1745	970	1345	1715
Scott	3	245	335	1770	2045	455	1620	1860	715	1470	1715	910	1380	1840
	4	265	370	1775	2060	440	1670	1855	750	1510	1785	855	1350	1785
	5	160	440	1745	2040	510	1660	1865	770	1490	1715	895	1340	1790
Tooley	3	325	355	1880	2260	550	1820	2060	715	1700	2015	1045	1405	1995
	4	360	360	1920	2155	485	1760	2000	725	1710	2015	1105	1430	2020
Trim	1	300	315	1725	2200	475	1670	1900	680	1585	1850	980	1430	1895
	2	300	395	1770	2105	470	1665	1880	765	1600	1835	875	1360	1830
	4	220	415	1740	2045	540	1690	1950	755	1570	1805	915	1455	1820
Whitley	1	320	325	1760	2175	525	1735	1785	705	1500	1930		1410	2000
	1a	300	365	1750	2155	495	1730	1895	685	1610	1895	1065	1360	2045
	4	365	370	1775	2140	505	1765	1950	765	1645	1930	1115	1405	1985

a			ɔ			o			u		
F1	F2	F3	F1	F2	F3	F1	F2	F3	F1	F2	F3
	930	1800	650	880	1720	600	870	1705	325		
	915	1800	665	980	1750	515	830	1680	360		1710
735	885	1845	560	780	1820	485	770	1695	260	655	
715	980	2035	720	940	1945	415		1855	360		
755	965	1965	645	855	1905	455	745	1845	325	540	
725	950	1965	590	835	1960	495	735	1830	360		2020
		2080	435	845	1985	505		2175	335		2240
	1000	2080	485	910	2000	540		2185	375		2220
	770	1965	650		1905	485		2045	395		2145
	775	1920	600		1885	505		1885	430		2140
	890	1915	575	895	1900	505	735	2060	365	795	2160
725	1025	2100	505		2040	460	895	2040	340	735	2190
660	1015	2210	595		2000	450		2015	325		2230
760	985	2090	595	870	2095	595		2085	365	855	2175

Table 3.4

Formants	i	e	ɛ	a	ɑ	ɔ	o	u
F1	-	-	-	1	(16)	-	-	-
F2	-	-	-			12	12	23
F1 or F2	-	-	-	1	16	12	12	23
F3	-	1	-		-	-	-	7
F1 or F2 or F3	-	1	-	1	16	12	12	26

The number of occasions on which it was impossible to locate a formant or formants in the 31 examples of each cardinal vowel

The first point to note about these results is the number of vowels which cannot be adequately described in terms of their formant frequencies. This information is summarised in table 3.4. It may be seen that an acoustic analysis in accordance with the procedure which has been outlined provides in nearly every case sufficient information about the front vowels (although, as has been noted, sometimes rather procrustean methods are necessary to extract the information). But this is not so for the back vowels. In 16 (out of the 31) examples of [ɑ] it was impossible to distinguish formant one from formant two, and accordingly the centres of these two formants were said to coincide. This solution is not one which could be readily adopted by any of the existing devices for the automatic extraction of formant frequencies; nor does it fit in with the Acoustic Theory of Speech Production (Fant, 1958), which shows that it is impossible for two or the resonances of the vocal tract to coincide. But from the point of view of acoustic analysis it is the obvious solution. However it is also possible to make an arbitrary decision that whenever two formants appear to coincide one shall be considered to be 75 cps (or 50 mels) higher than the coincident value, and the other shall be considered to be 75 cps (or 50 mels) lower. This course has been adopted in some of the work which follows.

Of the other back vowels, both [ɔ] and [o] often (39% of the time) could not be analysed in terms of three formants, because it was impossible to locate the position of formant two; and the complete analysis of [u] was usually (84% of the time) impossible. In the case of these vowels as with [ɑ], it would seem unlikely that any formant tracking machine would be any more successful in locating the formant frequencies. There is no existing

machine which can either (a) use built-in data to enable it to select for a particular vowel the regions of the spectrum in which there are likely to be formants, or (b) search these regions over a range of 35 dbs (or more if there is no automatic volume control ensuring that each incoming vowel has the same maximum level) and locate the centres of each formant.

When considering the analysis which was undertaken, and the results shown in table 3.4., it must be remembered that none of these cardinal vowels are in any way unusual sounds. They were produced under particular experimental conditions; but each of these qualities might quite easily have occurred in the normal everyday speech of an individual. The sounds are more difficult to analyse than those of ordinary conversation in that each of them has a constant fundamental frequency; consequently the formant centres may not be defined so clearly as they might be when their component harmonics are changing all the time. But on the other hand they are easier to analyse than the vowels in words, in that their quality remains invariant for an appreciable time; in ordinary speech some vowels do not last for more than three or four vibrations of the vocal cords, so that although the fundamental frequency may be varying, the formant centres may not be well defined.

The results show in table 3.4. are a measure of the formant finding difficulties which have been detailed previously (pp 41 - 44). It would seem that, at any rate as far as the back vowels are concerned, specifications in terms of formants are not practicable for the phonetician qua acoustic analyser, although they may be convenient for him qua synthesiser or qua research worker interested in the relations between vocal tract configurations



and the resultant wave forms. In other words, the configuration of the vocal tract may be usefully specified by stating the "formant" or resonant frequencies (Fant 1958 a, b, & c) and such a specification may be what is required by a speech synthesiser in that it may consitute a comprehensive statement about most of the other acoustic attributes of the sound, such as the relative formant levels or the overall spectrum envelope (Fant, 1956). But it will often be very difficult for a phonetician qua acoustic analyser to provide such a specification.

So far we have followed Fant (1956) in the view that whenever we can specify the formant frequencies of a vowel, such a specification constitutes an account of all the principal spectral attributes of the sound. But of course it does not follow from this that a three formant specification can itself be regarded as a statement of the principal auditory attributes of the sound. From a three formant specification of a sound it may be possible to calculate the relative intensities of the formants but until such a calculation has been made, a straightforward statement of the formant frequencies of a sound is not an adequate account of the auditory properties of the sound. Thus among the auditory attributes of a sound having formants at 330, 2250 and 2690 cps (i.e., Scott, set 4, [e]) is a quality feature dependent on the fact that there is energy at each of the formant frequencies. But there is unlikely to be a quality feature of this kind in a sound with formants at 390, 650 and 2620 cps (i.e., Scott, set 5, [o]). In the first case, the second and third formants have approximately the same intensity as the first formant (as could perhaps be predicted through knowledge of their frequencies); they therefore make a significant contribution to the auditory quality. But in the

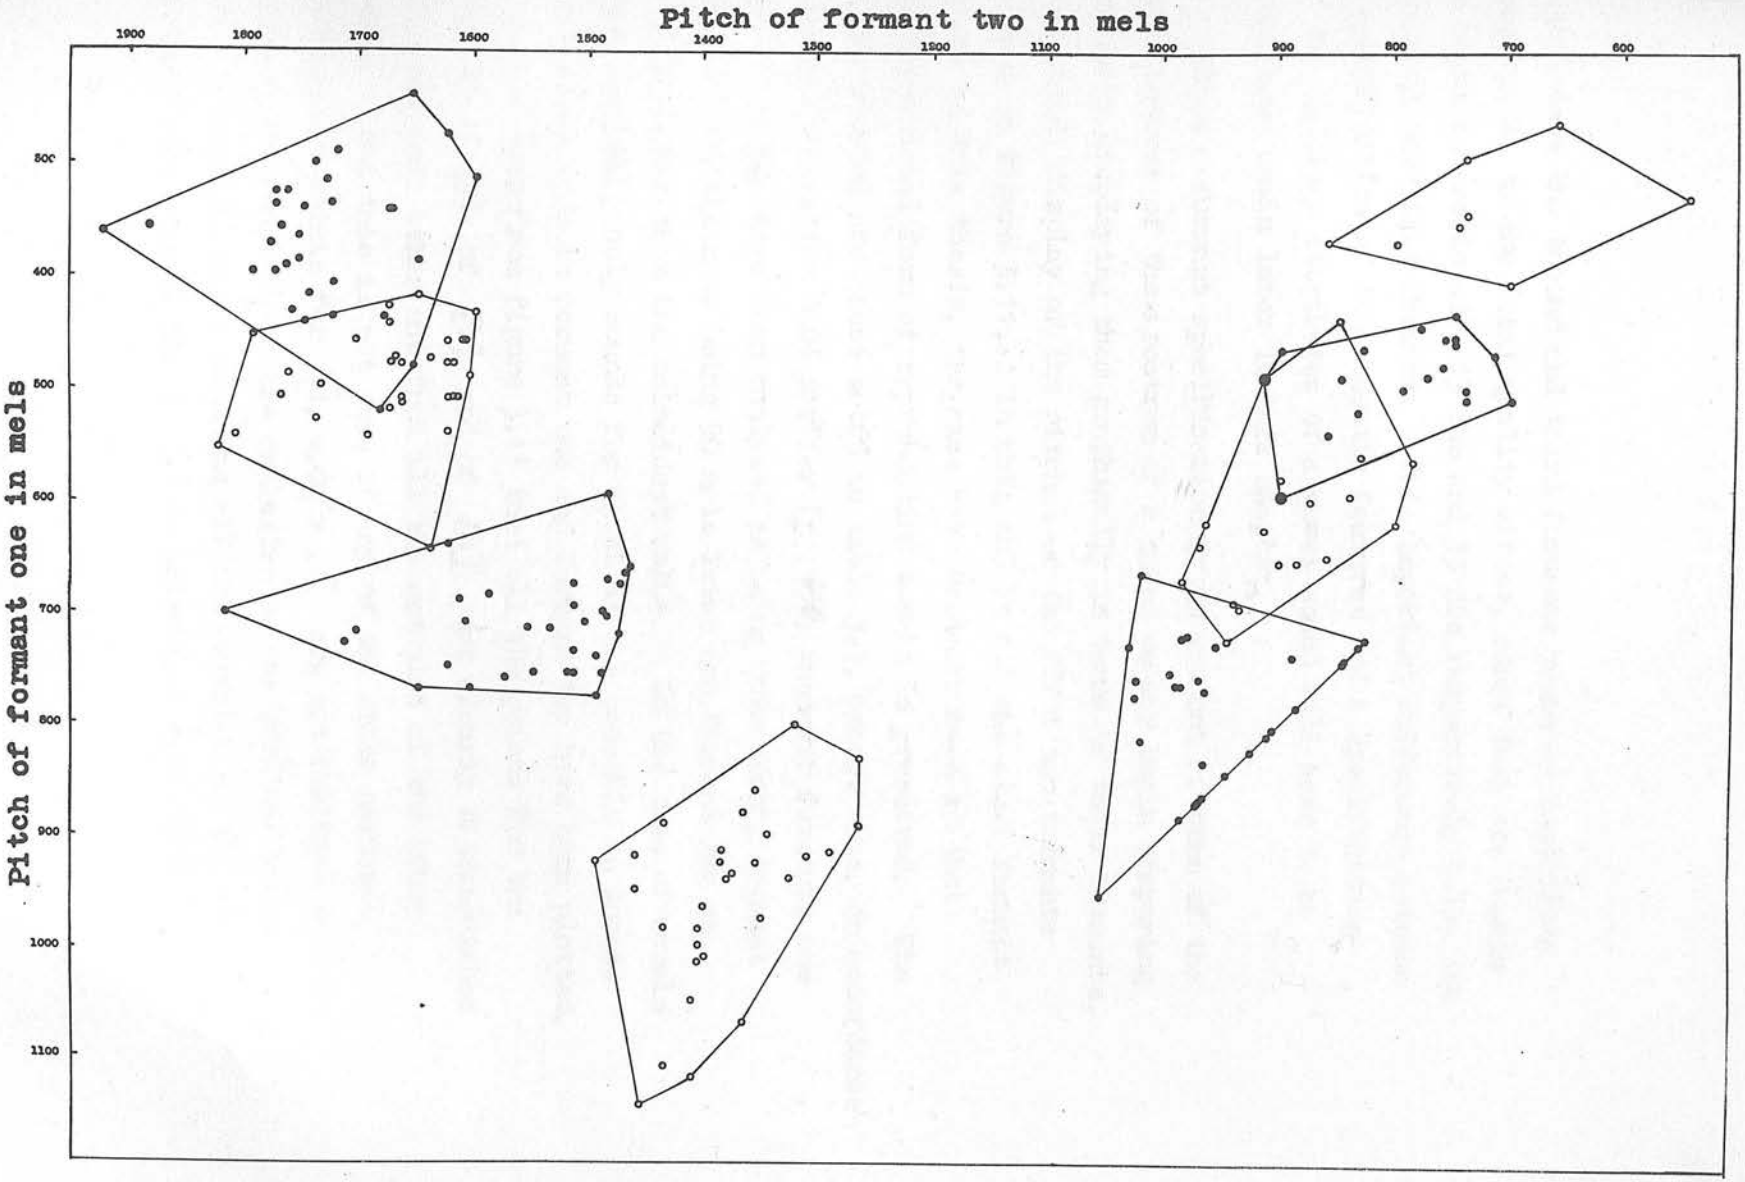


Figure 3.11 The pitch of the first two formants of a number of cardinal vowels.

second case the second and third formants probably contribute very little to the total quality effect, since they are (again as might be predicted), 19 dbs and 33 dbs respectively below the level of the first formant. This important difference between a specification of the acoustic features and a specification of the auditory attributes of a vowel sound will have to be considered again later in this chapter.

Since a formant specification is an account of some of the main features of the spectrum of a sound we may begin comparing vowels by displaying them graphically in terms of their formants. A graphical display of the pitches of the first two formants is shown in figure 3.11. In this and in all the other formant charts in this thesis, the axes have been arranged so that the traditional form of representing vowels is preserved. The values plotted are those noted in table 3.3, except that, in accordance with the suggestion made earlier (p. 50) whenever formants one and two of [a] have been analysed as being coincident, formant one has been taken as being 50 mels lower and formant two as 50 mels higher than the coincident value. In the case of vowels other than [a], only sounds for which it was possible to locate the centres of both formant one and formant two have been plotted. It may be seen from figure 3.11 that all the points for the examples of [ɛ] of [a] and of [u] are clearly distinguished both from each other and from all the examples of the other vowels. But this is not true of any of the other cardinal vowels; the points for [i, e, a, ɔ, o] are not confined to areas on the chart which are exclusive to one cardinal vowel. Thus the minimum area containing all the examples of [i] also contains seven out of the 31 [e]'s which have been plotted;

and the minimum area for [e] contains five out of the 31 examples of [i] . The back vowels are even more intermingled; the area for [ɑ] contains three out of the 19 plotted examples of [ɔ] ; that for [ɔ] contains six out of the 19 examples of [o]; and that for [o] contains five examples of [ɔ].

Most of this confusion may be said to be due to the spread of [i] and of [ɔ] . Thus the area for [e] is reasonably compact; and the areas for [ɑ] and [o] , while slightly less compact, are nevertheless in the positions which might be expected. The spread of the plotted values of [i] may be partly due to inaccuracies resulting from the difficulty (see page 41) of specifying the position of formant one in such a vowel; and that of [ɔ] may be due to the inaccuracies arising from the difficulties (see page 42) in specifying the positions of two formants when they are close together. But the inaccuracies of analysis are certainly not such that they can be used as a complete explanation of why, for example, two of the [ɔ] sounds have precisely the same first and second formant frequencies as two of the [o] sounds. (The coincident points are indicated by slightly larger filled circles in figure 3.11).

However, before discussing this matter any further, it is profitable to look at another point of interest in figure 3.11. It may be seen that the distributions of the examples of [ɛ, a, ɑ] , all of which have a comparatively high formant one, are triangular in form. In the case of [ɛ] the main variations are in the values of formant two; for [a] they are mainly in the values of formant one; and for [ɑ] which, as has been explained, is rather a special case from the analytical point of view,

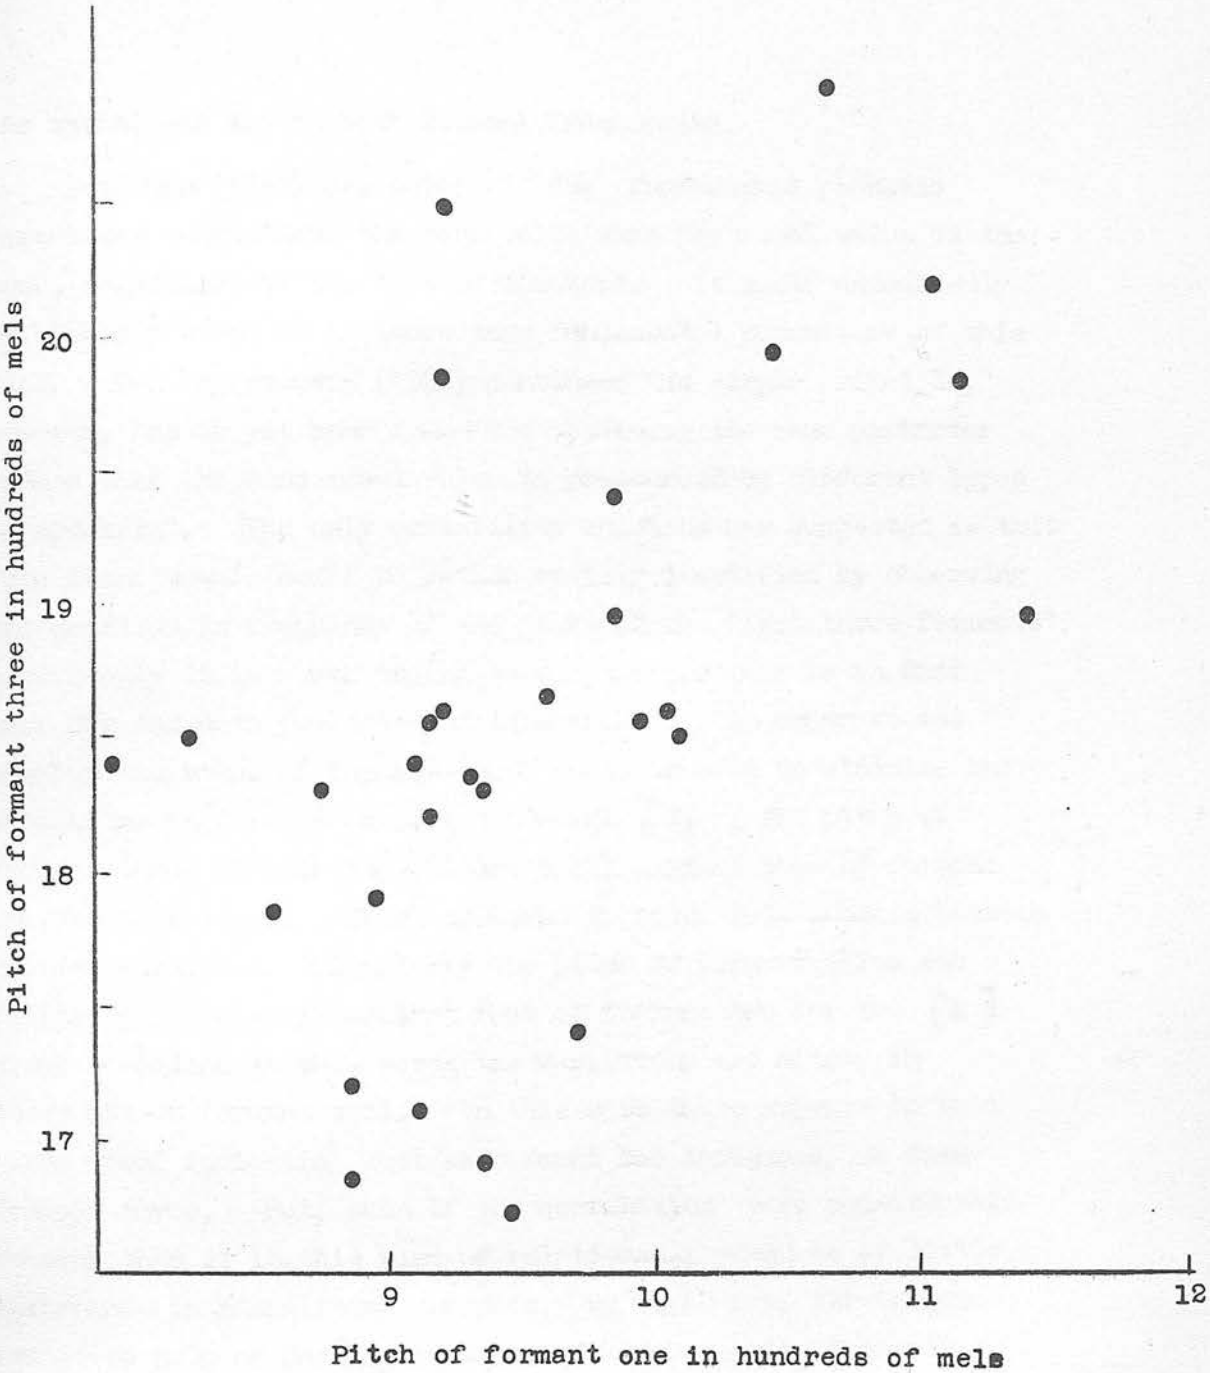


Figure 3.12 Formants one and three for the vowel [a]

the variations are in both formant frequencies.

Peterson (1952) has said: " The fundamental phonetic parameters should have the same value when the vowel value is the same, regardless of the type of speaker". It would undoubtedly be highly convenient if there were fundamental parameters of this kind. But as Peterson (1952) continues "No simple principle, however, has as yet been found for obtaining the same parameter values when the same vowel value is pronounced by different types of speakers". The only possibility which he has suggested is that "The front vowels could be rather readily identified by observing the position in frequency of the peaks of the first three formants". Accordingly it is worth trying to see whether this is in fact true for the data presented in this thesis. In order to see whether the value of formant three could be used to minimise the variations in formant one for the vowel [a] , the pitch of formant three was plotted (figure 3.12) against that of formant one for this vowel. There appeared to be no relationship between the two variables. Similarly the pitch of formant three was plotted (figure 3.13) against that of formant two for the [ɛ] vowel (since in this vowel the variations are mainly in the pitch of formant two). In this case there appears to be a vague trend indicating that as formant two increases, so does formant three. But, even if the correlation were considerably greater than it is, this kind of relationship would be of little assistance in normalising the data; we could not, for instance, use it to help us decide whether a given sound with formants whose frequencies were in the neighbourhood of the mean values for [ɛ] was in fact a cardinal [ɛ] or a slightly "centralised" vowel pronounced by a speaker whose cardinal [ɛ] had a formant

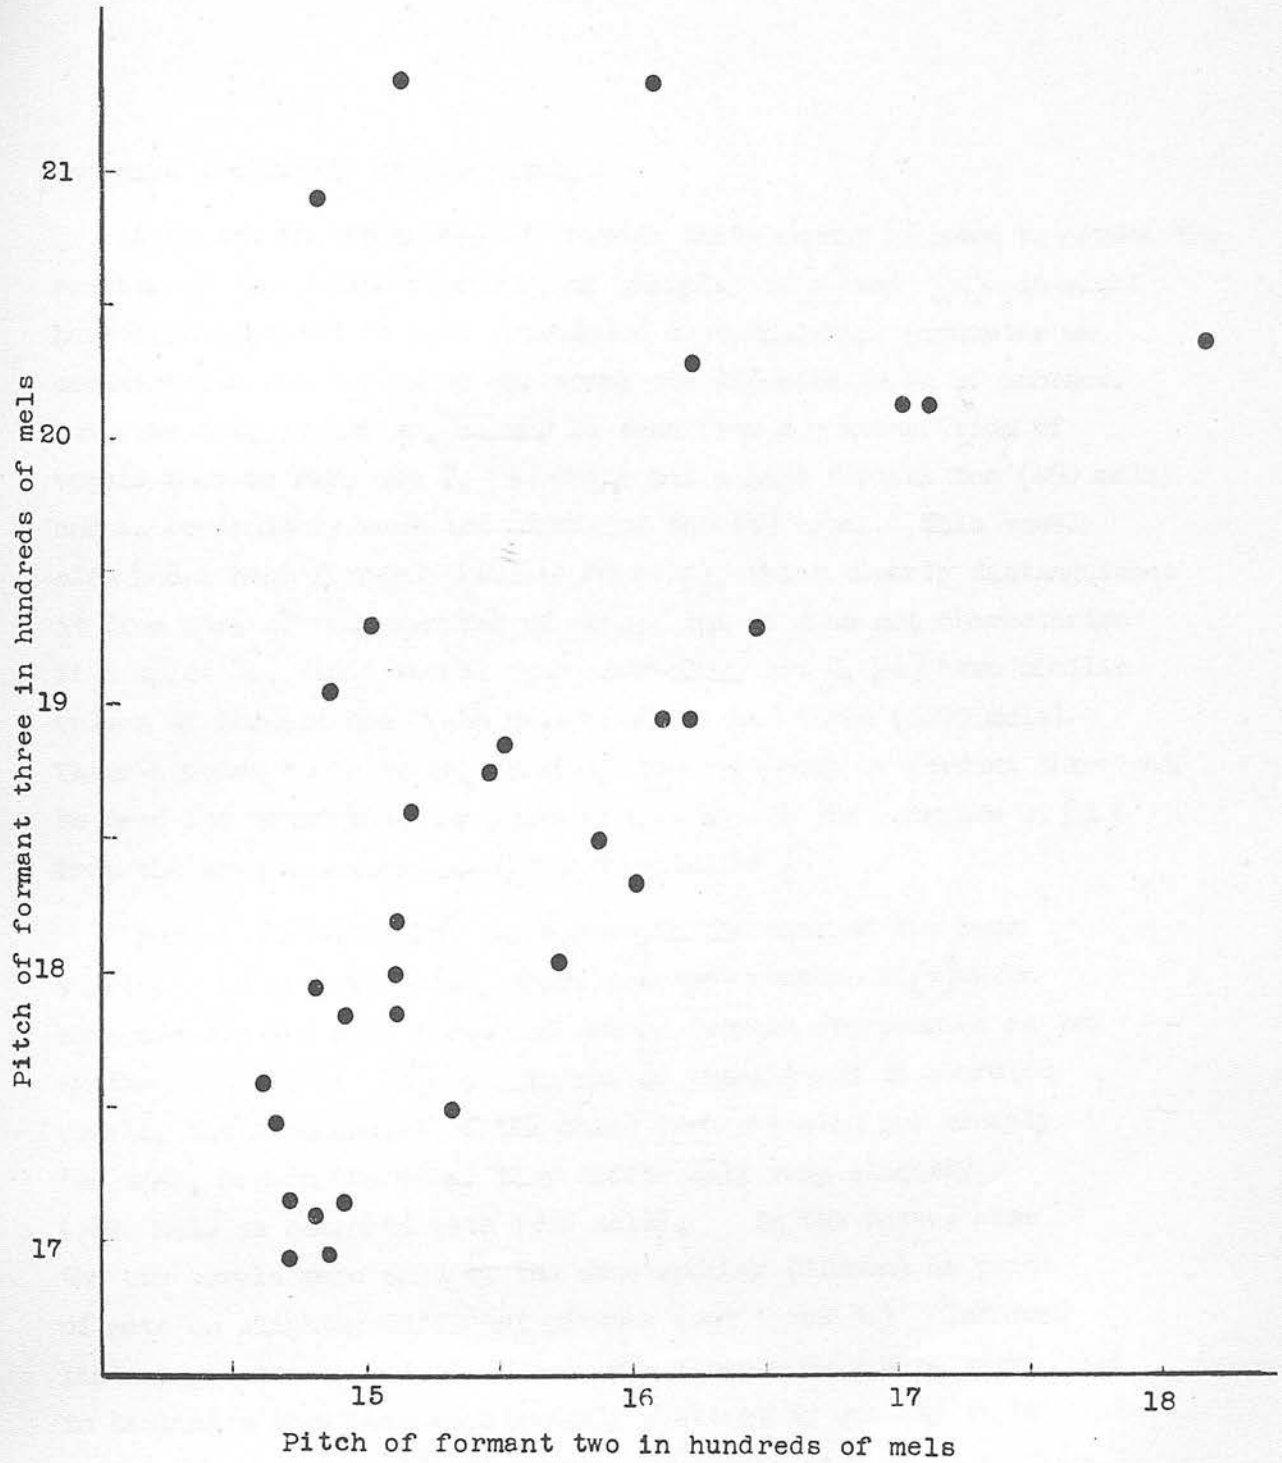


Figure 3.13 Formants two and three for the vowel [ɛ].

two with a slightly higher pitch.

Although the frequency of formant three cannot be used to reduce the scatter of the points representing examples of [ɛ] and [a] it might have been expected to have provided a distinguishing parameter on occasions when examples of one vowel overlap with those of another. But even this is not so, as may be seen from a consideration of vowels such as Fry, set 7, [i] which has a high formant one (480 mels) and is accordingly near the middle of the [e] area. This vowel also has a high formant three (1960 mels), which clearly distinguishes it from many of the examples of [e] ; but it does not characterise it completely, since vowels such as Tooley, set 4, [e] have similar values of formant one (485 mels) and formant three (2000 mels). There appears to be no way in which the frequency of formant three can be used for separating the area containing all the examples of [i] from the area containing all the examples of [e] .

Nor is the situation any better in the case of the back vowels. As we have noted, there are two examples of [ɔ] which have exactly the same first and second formant frequencies as two of the examples of [o] . In one of these pairs of confused vowels, the frequencies of the third formants also are exactly the same, and in the other they differ only very slightly (1890 mels as compared with 1900 mels). In the latter case the two vowels were said by the same speaker (Gimson) as parts of sets on slightly different pitches (set 1 and 3.) Informal listening tests showed that even when these vowels were reproduced in isolation they were sufficiently distinct in quality to be recognisable as (set 1) an example of [ɔ] and (set 3) an example of [o] . We therefore have a situation in which the specification of even three formants is quite insufficient as a specification

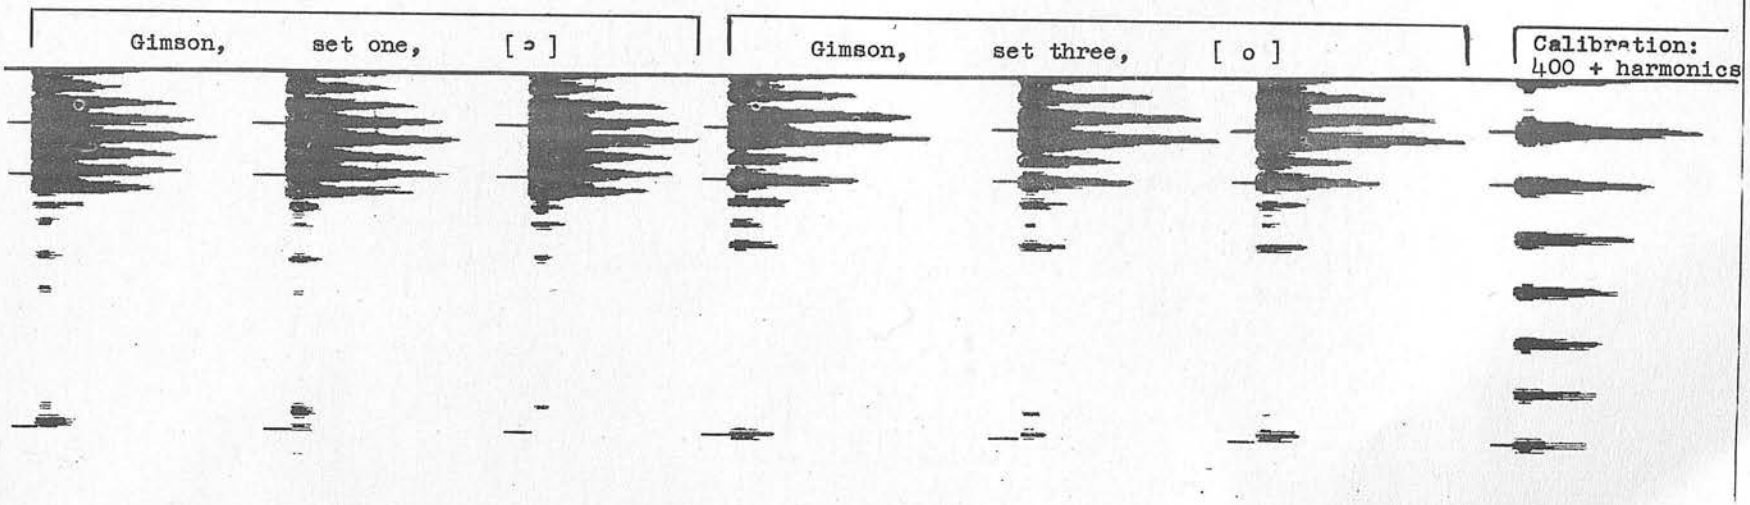


Figure 3.14. Gimson, set 1, [ ] and set 3, [o]; expanded scale, narrow band sections.

of vowel quality.

The reason for this is obvious as soon as we look at the acoustic analyses of these sounds. Expanded scale narrow band sections of both are reproduced as figure 3.14. In order that the reader may himself assess the formant frequencies, lines corresponding to the centres of the 400 and 800 cps calibration marks have been added on the left of each section. It is apparent that although the formant frequencies may be similar, the spectral envelopes are very different; the bandwidth of formant one is much greater in the example of [ɔ]. In the case of these vowels, therefore, it is not true that a formant specification is a complete account of the spectral envelope.

At this stage it is profitable to ask whether it is in any case possible to specify a vowel sound in terms of its own acoustic properties however precisely these may be known: in other words is it true that the fundamental acoustic (as opposed to phonetic) parameters should have the same value when the vowel value is the same, regardless of the type of speaker?

When we hear a vowel in isolation we can, of course, assess its quality to some extent. But our concept of the phonetic quality of a sound may become more precise or may be modified when we hear it in a context of other sounds. It may well be that the exact phonetic quality of a vowel sound does not depend on the absolute values of its formant frequencies, but on the relationship between the formant frequencies for that vowel and the formant frequencies of other vowels pronounced by that speaker. If this is so, then, as Eli Fischer-Jørgensen (1958) has said, it is "somewhat dubious to plot the vowels of different persons indiscriminately on the same

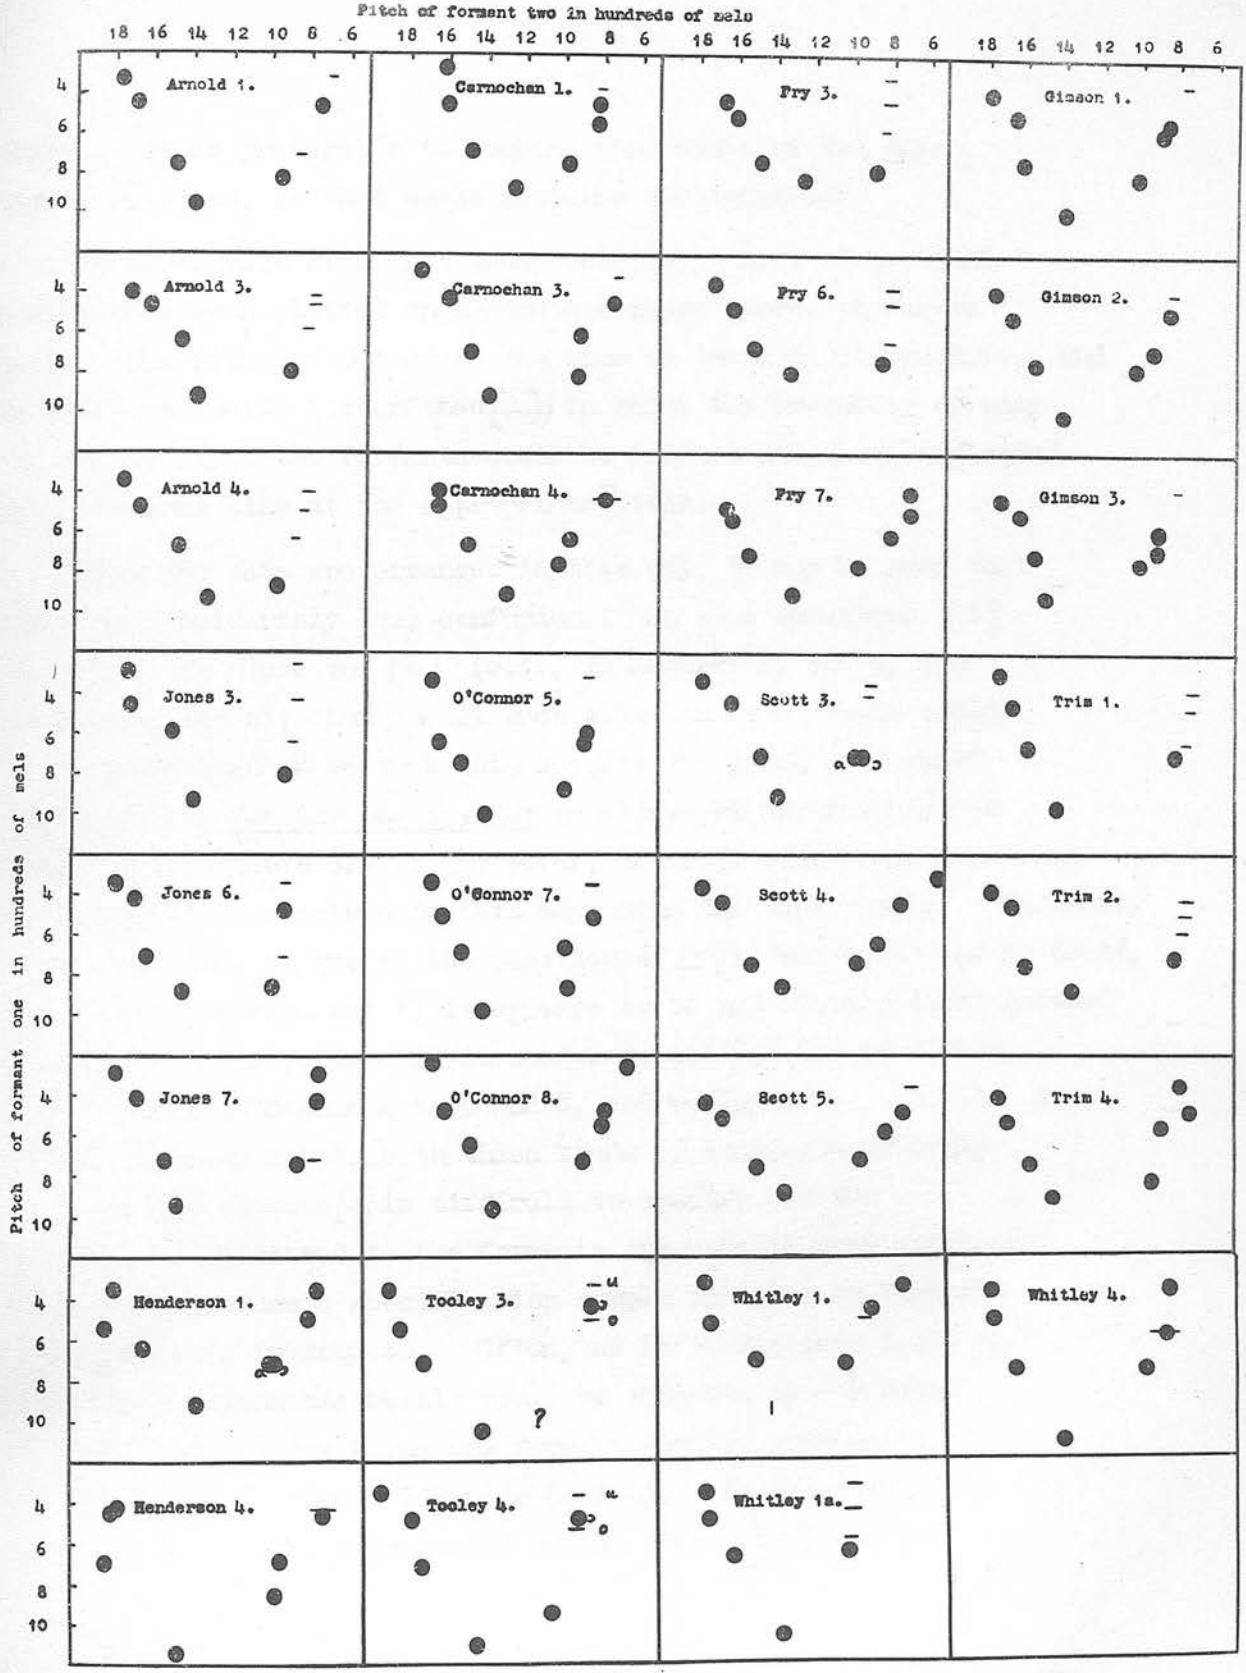


Figure 3.15. Formant charts of 31 sets of cardinal vowels.

chart. It is preferable to combine the vowels of the same person by lines, so that whole patterns are compared".

In accordance with this viewpoint the 31 sets of cardinal vowels have been plotted on separate formant charts in figure 3.15. The values plotted are the same as those in figure 3.11; and in addition vowels (other than [ɑ]) in which the frequency of only one of the first two formants could be located have been indicated by a straight line at the appropriate pitch.

When the data are arranged in this way, it may be seen that there is considerably less confusion. On some occasions [i] is still very close to [e] (e.g., in Henderson, set 4, and Carnochan, set 4); but in all such sets, the two vowels could be distinugished if we took into account the pitch of formant three, which, for any one speaker is always higher for [i] than for [e] (c.f. table 3.4.). However, there is still some confusion which cannot be resolved in this way among the back vowels. In nearly every case this is due to the position of [ɔ]; sometimes (as in Scott, set 3 and Henderson set 1) it appears to be not clearly distinguished from [ɑ]; and sometimes it is confused with [o] (as in Gimson, sets 1 and 3, O'Connor sets 5 and 8, and Tooley sets 3 and 4). As we have already noted, both these kinds of confusion probably arise partly because [ɔ] is difficult to specify and the frequencies allocated to the formants are usually very arbitrary, and partly because a specification simply in terms of formant frequencies is inadequate. Often, as in Tooley sets 3 and 4, a far more reasonable result could be achieved by regarding the examples of [ɔ] as being one formant vowels, with pitches of 640 mels (set 3) and 700 mels (set 4). In this way these two vowels could be represented on the charts by lines between

[ɒ] and [ɑ] in each case. But this solution is not open to us if we are forced into trying to specify all vowels in terms of two or three formants.

However, these difficulties should not lead us to overlook the considerable merits of a display of the form shown in figure 3.15; although there is still some confusion among some vowels, at least the vowels [e, ɛ, a, ɑ] form a reasonable pattern in the majority (about 27) of the sets. Of course, the points in these patterns have very different absolute values, so that the patterns for different speakers occupy different sections of the formant chart. In addition the patterns do not have exactly the same shape for every speaker. But these differences are irrelevant if we conclude that, when representing the acoustic characteristics of those vowels which can be specified in terms of two or three formant frequencies it is desirable to take into account the individual pattern shape which we learn to be characteristic of the speaker.

The data in figure 3.15 also enable us to consider the question of the "auditory equidistance" of cardinal vowels. We have already noted that although cardinal vowels are defined in terms of a number of properties which include auditory equidistance we cannot take this phrase at its face value, since whatever their theoretical specification, when a speaker pronounces a set of cardinal vowels what he actually does is to imitate a series of sounds which he has learnt through aural instruction. "Auditory equidistance" therefore may be a property ascribed to cardinal vowels solely by their originator; and, in fact, most of the phoneticians with whom the subject was discussed considered that the interval between each of the first five

vowels is greater than that between each of the last four. The data in figure 3.15 confirm this point of view. Nearly all the patterns are roughly isosceles triangles, with [i, a, u] at the apices; and since [ɛ] and [ɑ] are usually about the same distances from [a], the vowels [ɔ, o, u] are closer together than the vowels [i, e, ə, a, ɑ].

In summary, therefore, we may say that the analysis and representation of some cardinal vowels in terms of two (or more) formants is reasonably satisfactory; but this form of analysis is not practicable for other cardinal vowels. We must, therefore, consider alternative forms of representation. One possibility is that instead of starting from the premise that a formant specification should be used because it is claimed to be a complete statement of the acoustic structure of a vowel sound, we should try to specify the important auditory features of the sound. Then it might not matter that our present techniques of instrumental analysis do not allow us to specify the features which characterise the spectral envelope in the simplest possible way (i.e., in terms of the formant frequencies) since we can be sure that these same instrumental techniques will reveal quite adequately everything that can be heard. We do not know which, out of all the details that are revealed by an acoustic analysis, are the most important auditory attributes. But it is nevertheless useful to speculate briefly on the ones that a listener might use in his judgements of vowel quality. For, as we shall show later, it is a fact that trained phoneticians are reasonably in accord in many of their judgments. What then, are the features of the sound which they are assessing?

A partial answer to this question for a vowel in the unshaded area in the formant chart in figure 3.16 and the cardinal vowel chart in figure 3.17 is that information about the

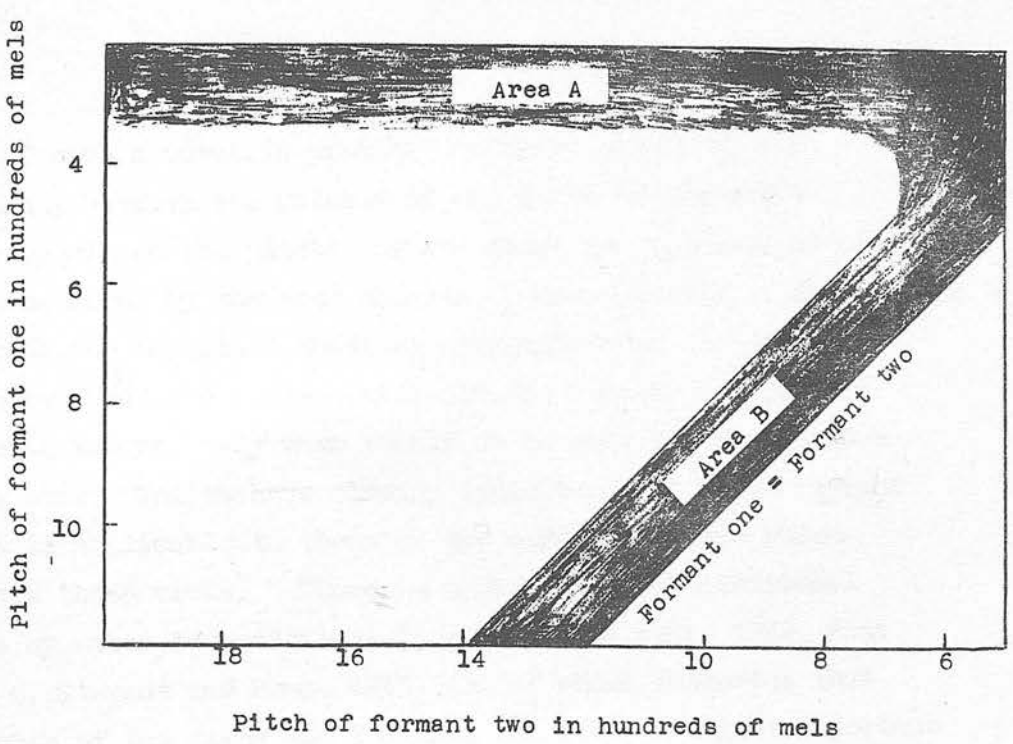


Figure 3.16 A formant chart, the shaded area of which shows the areas in which vowels cannot be specified simply in terms of the pitches (or frequencies) of their formants.

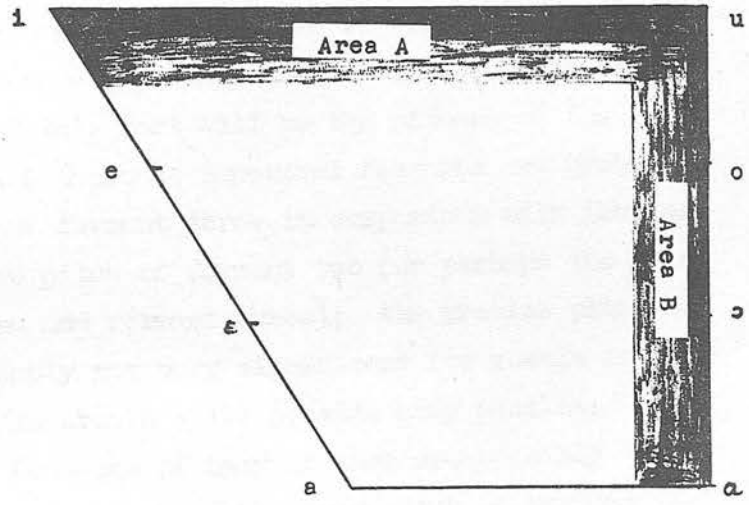


Figure 3.17 A cardinal vowel diagram, the shaded areas of which show the areas in which vowels cannot be specified simply in terms of the pitches (or frequencies) of their formants.

quality of such a vowel is probably conveyed mainly by the relationship between the pitches of the first two formants for that vowel and the pitches of the first two formants of other vowels pronounced by the same speaker. Consequently, for vowels in these areas the important auditory attributes and the principal acoustic features are equivalent. In the current analysis we have been concerned only with vowels on or near the peripheries of these areas; and we have already indicated that a two formant analysis is applicable to those of the cardinal vowels which are within these areas. There is a great deal of additional evidence by other investigators (e.g. Delattre 1951, 1952, Fant 1958b & c, Stevens and House 1955) all of which indicates that the pitches of the first two formants are among the most important auditory characteristics of all vowels.

However, this may not be true for vowels in the two shaded areas. These are the areas in which it is difficult to locate the centres of two formants either (area A) because the first formant has a very low frequency, or (area B) because the first and second formants are too close together. It seems unlikely that the cues utilized by the ear in assessing the quality of vowels of this sort will be the pitches of the formants. In area A the most important features are probably the high intensity of formant three in comparison with that of formant two, and the pitch of formant two (or perhaps the mean pitch of formant two and formant three); the precise pitch of formant one is probably not very significant for vowels in this area. The vowels in area B still provide many puzzles; but the characteristic features of many of them are possibly the pitch of formant one (or perhaps the mean pitch of formant one and formant two) and the comparative absence of energy in the

upper part of the spectrum. On these assumptions, when a vowel is being analysed and it is found to have a formant three with a greater intensity than formant two, it should be located in area A on a chart; and the greater the difference in intensity the deeper the vowel should be in area A. Of course such a vowel would also have a low formant one; but if this system of analysing and representing vowels were adopted, there would be no need to try to locate the centre frequency of formant one, although for some purposes (e.g., analysis- synthesis telephony) it might be convenient to allocate a conventional formant one frequency to such vowels. Similarly vowels which on analysis are found to have comparatively little energy above formant one could be said to belong in area B. The centre frequency of formant one of these vowels could be measured in the usual way and the amount of energy above formant one could be used (for analysis - synthesis telephony, or for plotting the vowel on a formant chart with the usual type of scale) as a measure of the frequency to be allocated to formant two, thus avoiding the difficult task of having to measure this frequency. In that part of the chart where area A overlaps area B will be found vowels which are characterised by a combination of the features described. These vowels have comparatively little energy above formant one; but what there is is associated with formant three, which has a far greater amplitude than formant two.

The usefulness of these speculations concerning the auditory cues which a listener uses in judging vowel quality can be assessed by further reference to figure 3.15. It may be seen that in nearly every set [i] is too close to [e] .

But if all the examples of [i] were regarded as vowels in area A in which the value of formant one is immaterial, and the relatively high intensity of formant three is the criterion for the value of the ordinate, then these sounds could be placed in more appropriate positions on the charts. Most of the vowels in area B also appear in more appropriate positions if we take into account the suggestion that the pitch of formant two be not used for the value of the abscissa. This may be said to have been done already in the case of those vowels which have been represented by a line at the pitch of formant one. On the present hypotheses we are entitled to consider these sets (which, if we were restricting our discussion to those sets of vowels which could be satisfactorily analysed in terms of two formants, we could not do). So we may now say that 29 sets can be represented in terms of their auditory attributes. This is a distinct improvement on the number (8 out of 31) which can be represented on acoustic charts which require the specification of 2 formant frequencies. But it must be noted that there are still several unexplained phenomena. Foremost among these is the peculiar position, on any kind of chart, of some examples of [ɔ] (e.g. in Tooley sets 3 and 4, and Gimson, set 3). The auditory characteristics of these vowels remain a mystery, and are likely to do so until there has been a great deal more research on the factors which can be correlated with perceived quality differences.

4. The relative nature of vowel quality.

The work described in the last chapter was concerned with what phoneticians call phonetic quality. This term means little or nothing to the man in the street; accordingly before we can discuss experiments involving naive subjects, we must extend the discussion which we began in chapter one on the nature of quality differences, and consider the kinds of information that are conveyed by speech. For the sake of convenience in exposition we may consider this information to be of three kinds. Firstly, when we listen to a person talking we can receive information about what he is saying; in other words, we can appreciate the linguistic significance of the utterance. Secondly, in addition to the information we receive as a result of considering an utterance in terms of a linguistic system conveying lexical and grammatical information, we also receive information of a different kind about the general background of the speaker; thus we can usually infer something about a speaker's place of origin and his social status from his accent. This kind of information may be termed accentual; it is conveyed by the features of a person's speech which he acquires through the influence of the particular groups of which he is (or was) a member. Lastly there is the kind of information conveyed by the idiosyncratic features of a person's speech. These, like the group and linguistic features, may be part of an individual's learned speech behavior; but unlike the other features, idiosyncratic features may also be due to anatomical and physiological considerations, such as the particular shape of the vocal cavities. The information which these features convey may be termed personal information. The relations between these three kinds of information

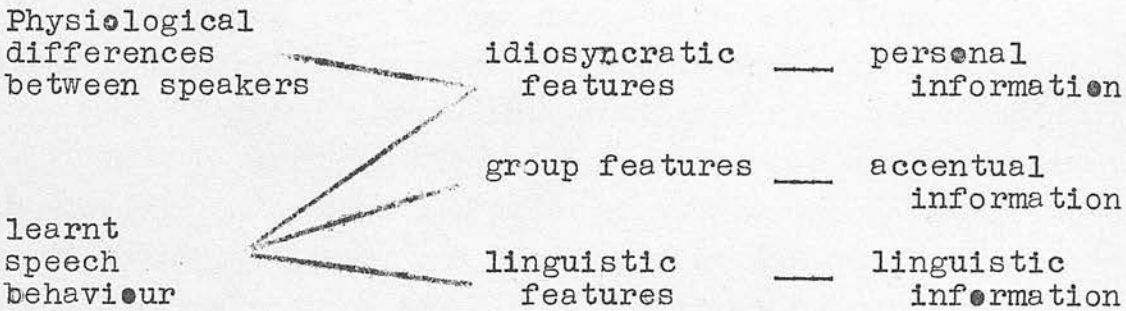


Figure 4.1 Differences in utterances and the information that they convey.

are summarised in figure 4.1 which is a more elaborate version of figure 1.1.

It is possible to arrange experimental situations which will elicit responses, from naive subjects, with respect to each of these three kinds of information. Thus one can ask a subject: Were these two sounds pronounced by the same speaker? (personal information); or: Is there any difference of accent between these two speakers? (accentual information); or: Do these two utterances consist of the same words used in the same way? (linguistic information). It is more difficult to arrange a situation where the socio-linguistic information and the linguistic information will be assessed concurrently by naive subjects. These two kinds of information taken together constitute, from the point of view developed in chapter one, the phonetic quality of a sound. As we saw earlier, they are normally assessed concurrently in situations such as those involving pronunciation teaching. But these situations do not lead to simple experiments in which naive subjects can make valid judgments of vowel quality.

In the current chapter we will consider experiments which are concerned mainly with those features of vowel quality which convey linguistic information, since these features are the aspect of phonetic quality which it is easiest to examine in an experimental situation involving naive subjects. But the same experiments will lead to a consideration of both the group and idiosyncratic features of vowel quality, so that tentative conclusions may be reached concerning the ways in which all

three kinds of information are conveyed.

We may begin from the principal conclusions reached in the last chapter, where it was suggested that the phonetic quality of any vowel that could be placed in the unshaded area of the formant chart in figure 3.16 or the cardinal vowel diagram in figure 3.17 (i.e., most vowels in normal speech) is partly dependent on the relationship between the frequencies of the first two formants for that vowel, and the frequencies of the first two formants of other vowels pronounced by the same speaker. On this hypothesis whenever a listener to speech has to assess the phonetic quality of a vowel, he utilizes whatever knowledge he has of the speaker's formant frequencies. Even when the vowel which the listener is considering is quite unlike any that he has heard the speaker produce before, he focuses his attention not simply on the absolute values of the frequencies of the formants, but on the relations between those frequencies and the general ranges of frequencies which he guesses to be characteristic of the speaker. Thus the hypothesis is that vowels are assessed at least partly in terms of the way in which their acoustic structure fits into the pattern of sounds which the listener has been able to observe or considers (on the basis of his knowledge of similar speakers) to be probable.

This hypothesis would be verified if it could be shown that naive subjects consider that the linguistic information conveyed by a vowel differs when the same physical sounds is heard in different contexts; two sounds can convey different linguistic information only if they have different phonetic qualities. Accordingly, an experiment was devised which was

designed to show whether subjects were influenced in their identification of a test word by variations in the introductory sentence preceding it.

In order to carry out this experiment it was desirable to obtain introductory sentences which were identical except in the ranges of their formants. This cannot of course be done by recording different people saying the same sentence, because the utterances are bound to differ in many ways. Accordingly it was decided to use synthesized speech, which can be precisely controlled in all respects. The particular instrument used for the purpose was the Edinburgh University Phonetics Department's copy of the Parametric Artificial Talking Device developed at the Ministry of Supply Signals Research and Development Establishment, (Lawrence 1953). The essential parts of the device are a generator producing a pulse corresponding to the larynx pulse which serves to excite the vocal tract; four formant generators which respond to the pulse excitation; and a generator which will produce noise corresponding to the excitation in fricative sounds. This instrument will synthesise speech which can be specified in terms of six variables, but which nevertheless sounds so natural that recordings of some sentences are always confused with recordings of normal speech. The six variables which are normally specified are the intensity and frequency of the pulse excitation, the frequencies of the lowest three formants and the intensity of the fricative noise. In order to control the synthesizer so that it will produce an utterance, information depicting the 6 variables as functions of time is painted on to a glass slide. The slide is then scanned by a mechanism which produces six controlling voltages which vary with time. The voltages control the appropriate generators of the synthesizer so that

Table 4.1. Differences in the six versions of the introductory sentence: Please say what this word is.

Sentence Version	Differences from sentence 1	Frequency range in cps	
		Formant 1	Formant 2
1	...	275-500	600-2500
2	F.1. down	200-380	600-2500
3	F.1. up	380-660	600-2500
4	F.2. down	275-500	400-2100
5	F.2. up	275-500	800-2900
6	F.1 down F.2. up	200-380	800-2900

a sequence of speech-like sounds is produced.

As well as the factors which are specified by the information painted on the glass slide, it is also possible to vary other factors, such as the frequency of the fourth formant, and the amplitudes and damping constants of all four formants; but no provision is made for controlling these factors as functions of time, and they were not in fact varied in the course of the experiment. In addition, it is possible to alter the frequency range over which each of the formant generators is operating. It was this facility that was used to produce the necessary variations in the introductory sentences.

Six versions of the sentence Please say what this word is were synthesized with the PAT device. This sentence was chosen as a suitable introductory context because the formant frequencies of the sounds vary over a wide range. Formant one varies between the low value necessary to produce the /i/ in Please to the high value required for the /ɒ/ in what; and formant two varies between the high value in the /i/ of please and the low value at the beginning of the /w/ in word.

In making all six versions of the introductory sentence the synthesizer was controlled by a single slide. Consequently the versions were identical with one another except for the variations which were introduced in the ranges over which the formant generators operated. The variations are summarized in table 4.1, which shows the highest and lowest values both of formant one and of formant two that occurred during a vowel in each version of this sentence.

It is interesting to note at this point that despite the

Table 4.2 The frequencies of the first two formants in the four test words.

Test word	Frequency in cps	
	Formant one	Formant two
A	375	1700
B	450	1700
C	575	1700
D	600	1300

great acoustic differences between the versions they were all readily identifiable as the same sentence. Moreover, all the trained phoneticians who listened to the different versions agreed that the variations which had been introduced did not appear to make any significant difference in either the linguistic or the accentual information which was being conveyed. With the exception of version six, which did sound rather unnatural and could not be judged as a sample of normal speech, all the different versions sounded like the same sentence pronounced by people who had the same accent but differed in their personal characteristics.

In addition to these introductory sentences, four test words were synthesized. Each of these was of the form b-(vowel)-t. The formant frequencies for the middle of the vowel in each of these words are shown in table. 4.2. The vowels in each of these test words were of comparatively short duration.

Several listening tests were devised with the aid of recordings of the material which has been described above. The first test was taken by sixty subjects. The first part of this test consisted of recordings of the test words, A, B, C and D, arranged in a random order. There were ten items in this part of the test. Subjects were told that they would hear ten words, each of which might be either bit, bet, bat or but. They were instructed to tick the appropriate word on the answer sheets with which they had been provided. The means of the responses in respect of each test word are shown in table. 4.3.

Between each of the first five words in the listening test there was a short pause during which subjects were requested to count aloud from one to ten. This was done in an attempt

Table 4.3 Means of the responses of 60 subjects for the ten words in the first part of the listening test.

Test word	Number of subjects identified as:			
	bit	bet	bat	but
A	52	8	...	...
B	14	46	...	...
C	...	27	33	...
D	...	1	14	45

to prevent the identification of a test word being unduly influenced by the auditory memory of the preceding word.

In the second part of the recording test words occurred immediately after the various versions of the introductory sentence. Subjects were given the following written instructions:

> You will now hear a voice saying Please say what this word is. This will be followed immediately by one of the words: bit, bet, bat, but. Please tick the appropriate word on the answer sheet below. There are twelve test sentences in this part of the recording; after answering in respect of each, there will be a short pause, during which you will be requested to count aloud, slowly, from one to ten.

The twelve items were arranged so that the predicted responses occurred in a random order.

The results of this part of this test are shown in figure 4.2. It may be seen that subjects are undoubtedly influenced in their identification of the test word by the auditory context in which it occurs. Thus word A is identified as bit by 87% of the subjects when it is preceded by version one of the introductory sentence; but as bet by 90% of the subjects when it is preceded by version two in which the first formant varies over a lower range. All that remains to be shown is that the influence of the introductory sentence is in accordance with the hypothesis which was discussed earlier concerning the relative nature of this aspect of vowel quality.

The relations between the formant st ructures of the

Test word	Version	Relation to normal	Means of the responses of 60 subjects *bit*	*bet*	*bat*	*but*
A	1	=	53	7		
	2	F1. down	4	54	2	
B	1	=	5	55		
	2	F1. down	1	57	2	
	3	F1. up	58	2		
	6	F1. down F2 up	7	46	6	1
C	1			25	35	
	3	F1. up		48	12	
	5	F2. up		23	37	
D	1	=			11	49
	4	F2. down		1	36	23

Figure 4.2 Identifications of the test words.

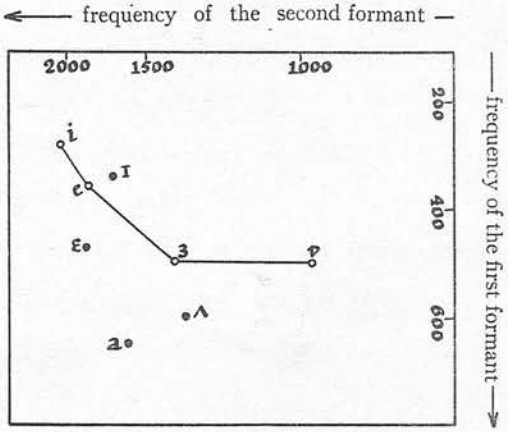


Figure 4.3 The formant structure of some of the author's vowels.

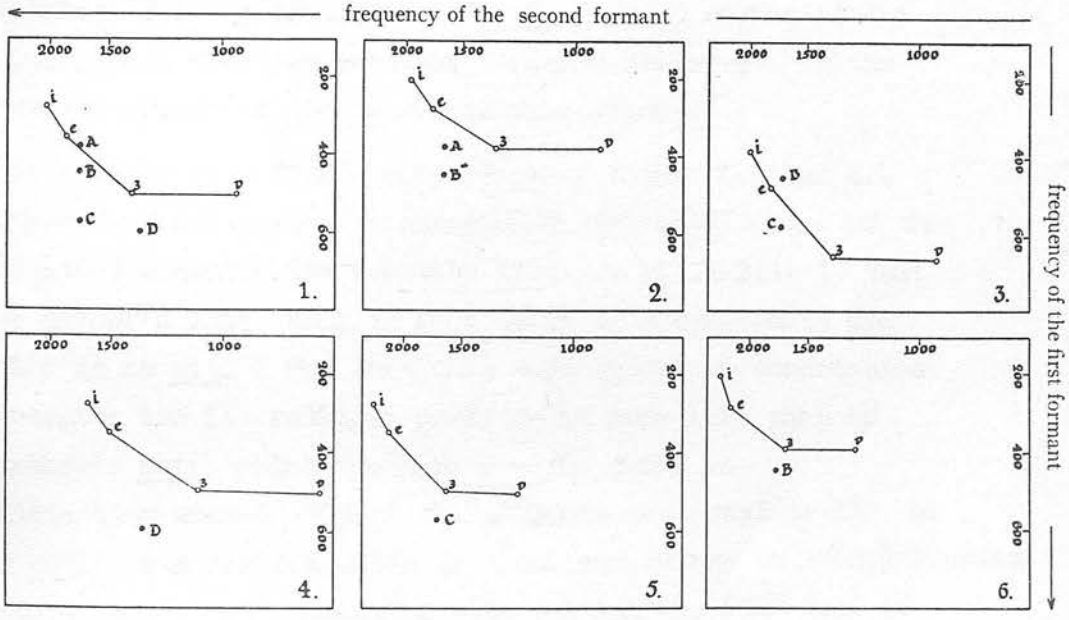


Figure 4.4 The formant structure of the six versions of the introductory sentence and the test words that they were each designed to be associated with.

vowels in a number of words can be conveniently represented by means of formant charts which show the frequency of the first formant at a time in the word when the formant structure is changing at a minimum rate plotted against the frequency of the second formant at the same time. In order to provide a basis for discussion, some of the author's vowels are shown in this form in figure 4.3. The symbols used are /I/ as in bit, /ɛ/ as in bet, /a/ as in bat, /ʌ/ as in but, /i/ as in please, /e/ as in say, /ɒ/as in what, and /ɜ/ as in word.

The pattern formed by the vowels shown may be taken as a representation of one of the kinds of relationships which can occur. Bearing this in mind, we may now consider the relationships between the vowels in each of the six versions of the introductory sentence and the test words with which they were designed to be associated. Figure 4.4 presents these data; solid points lettered A, B, C and D represent the test words, and the open circles indicate the vowels in the different versions of the introductory sentence.

It will be seen from a comparison of figure 4.3 and 4.4 that when the test word A is associated with version one of the introductory sentence its relative position is similar to that of the author's bit; and, in fact 87.5% of the subjects did identify it as bit. But when this word occurs in association with version two its relative position is more like that of the author's bet; which accounts for the shift in identification whereby 90% of the subjects now consider it to be bet. Similar reasoning explains the change in identification of word B when it is associated with version one (92% bet) as opposed to version three (97% bit). But we must also note in connection with word B that when it was associated

with versions two and six of the introductory sentence by far the majority of the subjects still identified it in the same way, (i.e., as bet) as when it was associated with version one. The probable reason for this is that the relative position of the vowel /ɛ/ as in bet can be anywhere in a comparatively large area. As Daniel Jones (1956) has noted: "The vowel (sc/ɛ/) varies a good deal with different speakers". Presumably, therefore the shifts in its relative position due to its being associated with versions two and six were not great enough to move it out of the part of the vowel pattern in which it is reasonable to expect to find a vowel of the /ɛ/ type.

The results shown in figure 4.2 indicate that there is a considerable amount of disagreement concerning the identification of word C. Some of the reasons for this can be appreciated from a comparison of the data presented in figures 4.3 and 4.4. Only when it is associated with version three of the introductory sentence does the vowel in this word have a relative position which is comparable with any of the relative positions of the author's vowels. In these circumstances 80% of the subjects did identify it as the same word, bet. But when it occurs in association with version one, where it occupies a relative position only slightly nearer the point in the pattern occupied by the author's bat than his bet, it is not surprising that 58% of the subjects identify it as bat and 42% as bet. The results obtained from the association of this word with version five however, are not so readily understandable. It might be expected that at least a small proportion of the subjects would identify this word as but in these circumstances But in fact this did

not happen, probably because the shift in its relative position was insufficient. On the other hand, tests with word D show that it is possible for the auditory context to influence the identification of a given test word so that it can be taken to be either bat or but. When this word was associated with version one, the majority of the subjects identified it as but (82%) as opposed to bat (18%); but in association with version four, in which the second formant was comparatively lower, then the results were bet (2%), bat (60%) and but (38%). Thus word D illustrates the fact that shifts in the range of the second formant in the introductory sentence can produce alterations in the identification of the test word which are of the same order as those produced by variations in the range of the first formant.

Taken all together, the results of this test show quite conclusively that the linguistic information conveyed by a given vowel is partly dependent on the relations between the frequencies of its formants and the frequencies of the formants of other vowels occurring in the same auditory context. The hypothesis propounded earlier must therefore be considered to have been verified in so far as one of the aspects of phonetic quality is concerned.

All the responses demanded by the listening test which has been described above are specifically related to linguistic information. But, on the basis of this test, two points may be noted concerning the socio-linguistic and personal information conveyed by vowels. Firstly as we have mentioned, there do not appear to be any differences in the socio-linguistic information conveyed by the different versions of the introductory sentence.

Table 4.4. Identifications of test word D in association with version one of the introductory sentence by different groups of subjects.

Number in group	Character of Group	Identified as: bat	but
7	Scots	3	4
19	English influenced Scots	4	15
19	English (R.P.)	1	18

It therefore seems that in this material the accentual information does not depend on the absolute values of the formant frequencies, but is, like linguistic information, partly a matter of the relative formant structure of vowels. In other words it would seem that both aspects of the phonetic quality of these vowels depend partly on the relative formant structure. This is further indicated by the second point: there is tentative evidence that subjects belonging to different socio-linguistic groups gave different responses to some of the test material. Consideration of the precise criteria that were used in dividing subjects into groups in accordance with their accents is outside the scope of this thesis. It must suffice to state that there were three main groups: in one there were seven subjects who had what is known as a Basic Scots vowel system (Abercrombie and Aitken, forthcoming); in the second there were nineteen Scottish speakers who had vowel systems that had been slightly modified due to the influence of the English of England: and in the third there were nineteen subjects who were speakers of the form of English of England known as R.P. Table 4.4 shows the responses of each of these three groups in respect of test word D preceded by version one of the introductory sentence. It can be seen that there is a greater tendency among the Scottish speakers to favor the identification of this word, as bat presumably because in their speech the relative position of the vowel in bat is similar to that of the vowel in the test word. When we consider that a speaker has vowel sounds which are typical of a Scottish speaker (i.e., when we interpret the accentual information conveyed by his vowels), we probably do so by appreciating the relative

formant structure of the vowels.

On the other hand, the personal information conveyed by vowels does seem to depend partly on the absolute values of the formant frequencies. Thus all the versions of the introductory sentence sounded as if they had been spoken by different voices. As we have already indicated, the reasons for this are best understood by reference to the articulatory processes involved in speech. The formants of a sound are essentially properties of the shape of the vocal tract. Consequently the ranges over which a speaker's formants can vary depend to a great extent on the size of his head. Because the ranges cannot be altered at will, they are not part of a speaker's learned speech behaviour, and can therefore convey only personal information. Additional personal information is, of course, conveyed by the relative positions of some of a speaker's vowels, insofar as these are idiosyncratic features of his speech and not aspects which identify him as belonging to a particular group.

But we may conclude that it seems reasonably certain that in assessing the phonetic quality of vowels which can be specified in terms of two formants the listener compensates for some of the individual characteristics of the voice he is hearing.

5. Adaptation to different personal characteristics

In the previous chapter data concerning a new auditory phenomenon were presented. The current chapter will discuss the psychological mechanism responsible for the adjustments made when listening to voices which differed only in their personal characteristics. The results which were described are suggestively similar to findings in other fields of sensory judgment, and there is a theory of "adaptation level" which has been put forward to explain such findings (Helson 1948). This theory may be briefly described as follows.

Responses along any dimension of judgment form an ordered series which corresponds to the physical ordering of stimuli; and at the middle of the series of responses there is a neutral point. Thus on the dimension of weight there is the series of responses very heavy, heavy, medium, light, very light. The neutral response ('medium' in the example) is not attached permanently to some particular physical stimulus, but rather tends to be given in response to the weighted mean of the stimuli which have been experienced up to the present. Responses on one side of the neutral point are given to stimuli which are physically greater than the weighted mean, and responses on the other side to those which are less. Thus a professional weight-lifter might regard a weight of 100 lbs as 'medium' and so call a weight of 50 lbs 'light'. A clerk on the other hand might call only 20lbs 'medium' and therefore would call 50 lbs 'heavy.' The subjective scale of responses may be shifted in the laboratory by presenting subjects with a series of weights; when a 50 lb weight follows a number of weights which exceed 50 lbs, it is likely to be judged 'light'. The stimulus to which at any

time the neutral response would be given is called the 'adaptation level' for that time.

It is important to note two points. In the first place this theory has been shown to be applicable not only to judgments of qualities such as heaviness, which is in ordinary speech a relative concept, but also to judgments which are usually regarded as absolute in the sense that there is a "correct" response, i.e., not one prescribed by the experiment, but one which is normally attached to the stimulus. Thus Campbell, Lewis and Hunt (1958) have shown that the key on a piano key-board named as that producing a certain note will be higher if a series of low notes has recently been presented for naming than if the series has been one of high notes.

In the second place, the formation of an adaptation level is usually conceived by psychologists as being directly analogous to the process of sensory adaptation as described by physiologists. Thus the adaptation level is thought to be dependent only on the stimuli presented and not upon the responses; just as the state of dark adaptation of an eye varies with the intensity of light which has been experienced, and does not vary with the opinion of the eye's owner regarding the brightness of the light. There are, however, certain experiments in the literature which are hard to explain on this basis. For example, Brown (1953) found that in judgments of lifted weights there was no effect on judgment if a heavy 'anchor' weight was inserted repeatedly into the series of weights for judgment - provided that the heavy weight took the form of a tray on which the test weights were presented, so that the subject regarded his lifting the tray as a help to the experimenter and not as part of the test sequence. If the interpolated weight

looked like the test weights, it produced the usual effect of increasing the number of 'light' judgments. In the latter case Brown also found that instructions to lift the interpolated weights without judging them produced less effect on the scale of judgments than did instructions to judge all weights.

Another instance is provided by Sherif, Taub, and Hovland (1958), who used anchor weights of various values. In one series of trials the anchor weight was equal to the heaviest weight of the test series; in other series the anchor weight was greater (sometimes much greater) than any weight in the test series. In the latter case the usual adaptation level effect appeared; the judgments were shifted in the direction of 'light'. But when the anchor weight was the same as the heaviest of the test series, there were actually more judgments of 'heavy' than there were in a control series, without an anchor. The authors consider that this is not easily explained in terms of adaptation level theory. However, it is stated in their experimental procedure that the weights were presented in pairs; and subjects were told that they should always regard the first (i.e., the anchor weight) as corresponding to the heaviest point on their response scale. A little armchair introspection would suggest that subjects might, in an anchor-less series, avoid the extreme end of the response scale in case a still heavier stimulus should appear in the future. But in a series with an anchor, whenever they were presented with two weights which were the same, they would be emboldened by their instructions to use the heaviest possible category of response, knowing that the anchor represented that category. However, in the series when the anchors were always far heavier than the following weight which had to be judged they would obviously

according to their instructions, never use the heaviest category of response. Thus the magnitude of the anchor weight is affecting, not the stimulus to which the neutral response is given, but that to which the response 'heaviest possible' is given. If this is taken into account, the results are perfectly consistent with those using different instructions to the subject: and most experiments supporting adaptation level do not specify the response to be given to anchor weights.

Both these experiments suggest that, in considering the effect of preceding stimuli upon the response given to a test stimulus, we must take into account the type of response made to the preceding stimuli. As we shall see, experiments on vowel judgements reinforce this suggestion. The application of adaptation level theory to vowel judgements is that a series of speech sounds in which the formant one position was high would raise the adaptation level, and so would cause a given test word to be judged as having a value of formant one lower than the adaptation level; and thus bit rather than bet. The theory of adaptation level has been put forward by Helson as explaining such phenomena as the constancy of apparent colour under changing illumination. The extraordinary efficiency of the compensatory mechanisms which perform this adjustment, and the efficiency of the related adjustments in perception which allow us to compare accurately the sizes of objects at different distances or of shapes in different orientations are all long-standing problems. Since the vowel judgment case is also at least superficially similar to that of constancy of hue under different illuminations, considerable general interest attaches to the degree of success achieved by adaptation level theory in explanations of vowel judgements. Three series of experiments were therefore carried out to evaluate the theory

as applied to these judgments.

The first series of experiments was designed to determine the effect of different instructions and conditions of presentation. As in the experiment described in the previous chapter, each group of subjects heard a recording of the sentence Please say what this word is followed by a test word which they had to identify as bit, bet, bat or but. They indicated their choice by putting a pencil stroke on a printed answer sheet. After making only one such judgment the subjects carried out an intelligence test for an hour, and then made another judgment on a second recording. In the first recording version one of the introductory sentence (see table 4.1) was used, and in the second recording version two (in which the range of the second formant was much lower) was used. The test word was in both cases test word A.

All the subjects were Naval Ratings, and none served in more than one group. The various groups were treated in the following ways:

Group I (15 subjects) were given no special instructions.

Group II(12 subjects) were instructed that the test word was spoken by a different person from the introductory sentence, and that therefore they were to ignore the introduction.

Group III (17 subjects) were presented with the test word from a loudspeaker separated from that presenting the introductory sentence. Since subjects were not tested individually the angular displacement of the speakers was not the same for all subjects, but was at least 45° for all of them. In addition the instructions emphasised which loudspeaker was operating for each part of the

Table 5.1. Number of subjects showing a change in vowel judgment when the introductory sentence has a different Formant I range.

Group	Subjects showing change in expected direction	Subjects showing no change	Subjects showing change in opposite direction
I (No special instructions)	14	1	0
II (Told to ignore voice)	11	1	0
III (Sentence and word from different loud-speakers)	12	4	1
IV Interval of 5 sec between sentence and word	9	1	0
V $7\frac{1}{2}$ sec interval	5	6	1
VI 10 sec interval	6	7	0
VII 10 sec interval filled by counting aloud	2	14	0
VIII Test word before sentence	1	10	0

recording.

Group IV (10 subjects) received the normal test procedure except that the time interval between the end of the introductory sentence and the test word was 5 sec instead of the usual natural interval of about $\frac{1}{4}$ sec.

Group V (12 subjects) resembled Group IV except that the interval was $7\frac{1}{2}$ sec.

Group VI (13 subjects) resembled Group IV except that the interval was 10 sec.

Group VII (16 subjects) resembled Group VI except that the subjects were instructed to count aloud from one to ten, in chorus with the experimenter, during the interval between the sentence and the test word.

Group VIII (11 subjects) received the normal test procedure except that the test word occurred before the introductory sentence.

On the basis of the results discussed in the last chapter, we would expect that if subjects heard a difference between the two presentations of the test word, they would consider that the first presentation differed from the second in a way which would correspond to its having been judged as having a lower formant one (i.e., as being bit as compared with bet, or bet as compared with bat).

Table 5.1 gives the number of subjects in each group who considered the two presentations of the test word to be different. It will be seen that the nature of the introductory sentence affects perception of the test word even when the instructions

are to ignore it. The effect is also present even when the word and the sentence are separated in space; and when they are separated by a time interval of 5 sec. Longer time intervals produce some reduction of the effect, although it is still present. It disappears virtually completely when the interval between sentence and test word is filled by the listener's own counting aloud, or when the test word comes before the sentence.

These results are similar to those long known in visual constancy phenomena in that they minimize the importance of conscious knowledge or inference. Telling a man that a colour-wheel is brightly lit does not alter his perception of it, and it is necessary for his field of view to contain the source of illumination or at least other objects lit by it. So also telling him to ignore a spoken voice does not stop him judging the test word with reference to that voice, and acoustic separation of sentence and word is needed.

In addition, the results are all consistent with the simple adaptation level theory stated earlier. Since the response to the word depends upon the adaptation level at the time when the word arrives, the effect is naturally abolished by placing the sentence after the word. Since the counting between sentence and word shifts the adaptation level away from that produced by the sentence, it abolishes the effect. The lesser but still real shift produced by an empty time interval (Group VI) may require further explanation: it can be compared to the error introduced in psychological experiments by presenting a stimulus some time after the standard with which it is to be compared. As is well-known, under these circumstances, with stimuli of fair

intensity the second will be judged subjectively as being more intense than the first: but with very faint stimuli the reverse is true. On adaptation level theory these phenomena are explained by supposing that the level drifts gradually from the positions in which each stimulus leaves it, back towards its normal position based on more remotely past experience. Thus ten seconds after a sentence in which a particular voice has been heard, the level of adaptation is a compromise between that produced by the voice and the normal resting level. This effect also is therefore consistent with adaptation level theory.

However when we come to consider the second series of experiments, in which the effect of changing the words in the introductory sentence was studied, we find that the theory of adaptation level is not entirely adequate. As we have noted, the theory takes into account only the stimuli which have been experienced, and not the responses made to them; in this the theory resembles ordinary accounts of sensory adaptation. It follows from this theory that it would be possible to alter the adaptation level (and hence vary the interpretation of the test word) simply by altering the words chosen for the introductory sentence, since (for any particular speaker) the vowels in some words have lower formant one values than the vowels in others. It is possible that if all the vowels in the introduction happened to have a low formant one position like the vowels in please and this, the adaptation level might be low; in this case a test word would have a greater likelihood of being judged as bet, or bat rather than bit. If on the other hand the vowels in the introductory sentence had rather high formant one positions the test word might be more likely to elicit bit responses. Thus it is interesting to compare the effects

of differently worded introductions upon judgment.

In addition, it is of value to compare long and short introductions; according to the theory of adaptation level, the differences in the interpretation of a test word after two different but long versions of the introductory sentence should be greater than the differences in interpretation of a test word after two different but short sentences, since the longer introductions would contain more stimuli to alter the adaptation level. An alternative theory might be that this was not so because the test word was compared with vowels of which there happened to be samples in the introductory sentence, and identified as that vowel which was physically most nearly identical to it. Thus in the recordings used in Series I, the test vowel had approximately the same first formant position as the vowels in this and is in the first version of the introductory sentence. Therefore, it might be argued, bit responses were given with that sentence; whereas with the second version the vowel of the test word was not similar to those in this and is and therefore was not judged as bit. On this view the simple sentence This is should produce as much effect as Please say what this word is.

In order to obtain more data on these points the following experiments were carried out. A group of 44 students were presented with a series of seven recordings containing the synthesised words What's this followed by a test word. At the end of the seven recordings another seven were played each of which consisted of the synthesised words This is followed by a test word. Subjects were instructed to indicate in each one whether the word was bit, bet or bat, using a printed answer

Table 5.2. Responses to different test words with introductory sentences having various wordings and shifts of first formant.

Test word	Sentence Version	Sentence Wording: Please say what this word is (60 subjects) response			Sentence Wording: What's this (44 subjects) response			Sentence Wording: This is (44 subjects) response		
		bit	bet	bat	bit	bet	bat	bit	bet	bet
A	1	53	7	0	10	33	1	40	4	0
	2 (formant one lowered)	4	54	2	1	40	3	8	34	2
B	1	5	55	0	2	40	2	17	27	0
	2 (formant one lowered)	1	57	2	0	36	8	2	39	3
	3 (formant one raised)	58	2	0	21	23	0	32	12	0
C	1	0	25	35	20	14	30	1	18	25
	3 (formant one raised)	0	48	12	2	28	14	4	34	6

Data for the longest sentence are from the experiment reported in the previous chapter

Table 5.3. Numbers of subjects giving different responses to the same word when the introduction is altered, for each wording of the introductory sentence.

Word	Versions Compared	Sentence	Predicted Effect	Number of Subjects showing No Change	Opposite Effect
	1 & 2	Please say what this word is?	58 = 87%	8 = 13%	0 = 0%
		What's this?	13 = 30%	29 = 65%	2 = 5%
		This is?	34 = 76%	10 = 23%	0 = 0%
	2 & 3	Please say what this word is?	56 = 93%	4 = 1%	0 = 0%
		What's this?	18 = 40%	26 = 59%	0 = 0%
		This is?	29 = 64%	15 = 34%	1 = 2%
	1 & 3	What's this?	18 = 40%	25 = 57%	1 = 2%
		This is?	23 = 52%	19 = 43%	2 = 5%

sheet. Between each recording there was a break in which the group counted aloud from one to ten.

In each group of seven recordings each of the three test words A, B and C ~~were~~ was used. There were also three versions of each of the introductory sentences. These three versions were produced with settings of the controls for the formant ranges which were identical with those used in producing versions one, two and three of the introductory sentence Please say what this word is described in the previous chapter. Test word A was presented with versions one and two of each of the sentences, test word B with all three versions of each sentence, and test word C with versions one and three of each sentence.

Tables 5.2 and 5.3 show the results of this test. Data from the previous chapter for the similar tests with the introductory sentence Please say what this word is are also included for comparison.

There are two main features of these results. Firstly, there is some evidence that the longest introductory sentence produces more shifts in judgement than either of the two short ones. The significance of this difference was assessed for each test word by examining the number of subjects showing the effect with each length of sentence. Thus for Test Word A, the longest sentence produced different responses to Versions 1 and 2 in more subjects than What's this did. The corresponding difference between This is and the longest sentence is not significant. For Test Word B, the longest sentence produces different responses to Versions 2 and 3 in more subjects than either What's this or This is. The differences between sentences are

not significant on Test Word C (although each sentence produces differences between Versions 1 and 3). As was noted in the previous chapter the values of the formant frequencies in Test Word C are such that this particular sound is less likely to be affected by different versions of the introductory sentence.

This greater effect of the longer sentence cannot be explained as due to a difference in the success of synthesis, producing a greater effect for the most natural sounding sentence. In a subsidiary experiment twelve Naval Ratings were asked to assess the relative 'naturalness' of the three synthesised sentences. Please say what this word is, What's this, This is and an additional synthetic utterance What did you say before that, using the method of paired comparisons. Each of the six possible pairs of sentences was presented, and the subjects indicated for each pair which sentence sounded more natural. Half the subjects heard the members of each pair in one order and half in the opposite order. All the subjects gave This is the highest place for naturalness, a result which is of course significant well beyond the .01 level. There were no significant differences between the other sentences.

The smaller effect produced by This is and What's this is consistent with adaptation level theory. Indeed it should be noted that the difference between This is and the longest sentence disproves the alternative view, mentioned above, that the test vowel is identified as the most similar of the vowels in the introduction. This is contains as many samples of the vowel in bit as the longer sentence does, yet its effect is smaller. The other vowels in quite different parts of the spectrum must be affecting judgment.

The second main feature of table 2 is not consistent with simple

adaptation level theory. It has already been said that on that theory an introduction containing a number of vowels with low formant one frequencies should give fewer bit responses than occur after an introductory sentence containing a number of words with vowels of a kind such that they have a comparatively high formant one. Now each version of the sentence Please say what this word is has a formant one with a mean value lower than the mean value of formant one in the corresponding version of What's this (since the added words in the longer sentence contain vowels with a low mean value of formant one). Therefore the adaptation level produced by the longer sentence should be lower than that produced by the shorter sentence; and it should give fewer bit responses if the choice is between bit and bet, or fewer bet responses if the choice is between bet and bat. Table 5.2, shows that this is not the case for any combination of Test Word and Introductory Version. The difference is in fact significantly in the opposite direction in three cases, viz, Test Word B with Versions 2 and 3, and Test Word A with Version 1.

Since this is the first result we have found which is flatly inconsistent with a simple adaptation level theory, it is worth examining it carefully. One possible explanation which might have been advanced is that the listeners began the experiment with a very low adaptation level, so that all the introductory sentences tend to raise the level by varying amounts. Although the mean of What's this is higher than that of Please say what this word is, the shorter the sentence might be less successful in pulling the adaptation level away from its initial value, and therefore given responses more toward the right-hand side of table 5.2. This explanation is unacceptable however, since Version 2 of the introductory sentence must have an adaptation level below that

brought by the listener to the experiment. We know this because Version 2 causes Test Word A to be heard as bet, although in isolation that test word is heard as bit. Yet Version 2 of What's this still gives responses more to the right-hand side of table 2 than does the longer introductory sentence.

The further implications of the result may usefully be discussed following the results of the third series of experiments which are also inconsistent with simple adaptation level theory. For the moment it may be noted that the vowels in What's this are not merely different stimuli from those in Please say what this word is but also produce different responses.

The third series of experiments was concerned with prolonged experience of the situation. If we think of the listener shifting his scale of vowels so that the mid-point corresponds to the weighted mean of the physical stimuli he has experienced, it would seem logical to expect larger effects if the exposure to a particular set of vowels was repeated again and again. The experiments in this series were intended to produce such exposures and to compare the results with those expected.

Three groups of Naval Ratings, each containing twelve individuals were tested. Each subject was provided with an answer sheet containing the possibilities, beat, bit, bait, bet, bat, bought, but and asked to indicate the word heard in each of a number of recordings.

Group X received prolonged experinece of Version 2 of the introductory sentence Plese say what this word is, and heard no other version. The details of procedure were as follows. They heard the four test words, A, B, C and D, each in isolation.

The order of presentation was ADBC. After these four isolated words came four recordings of the same words each preceded by Version 2 of the introductory sentence Please say what this word is but with a 10 second gap between sentence and test word. The order of presentation was BCDA. There then followed 60 presentations of the same introductory sentence, each followed at a normal short interval by a test word. The different test words were given in random order subject to the restriction that each successive block of four words contained one example of each word: no indication was given to the subject of the end of blocks. This restriction on order and the corresponding restrictions in other groups were intended to avoid temporary shifts in adaptation level within the sequence due to a block of similar test words. Following these sixty recordings the four recordings with a ten-second gap were repeated, and finally the four test words were repeated in isolation.

Group Y received prolonged exposure to Version 2 and Version 3 concurrently; that is, they heard sentences with high first formants interspersed with sentences with low ones. They started by hearing the four test words in isolation, in the order ADBC. They then received eight recordings with the ten second gap between sentence and word; four of the recordings, one with each test word, were with Version 2 of the introductory sentence and four with Version 3. The order of presentation was so arranged that all four test words, and two each of the introductory Versions occurred in the first four recordings. Following these eight recordings, there were fifty-six presentations with no extra gap between sentence and word. Half of these used Version 2 and half Version 3, and in each third of the test there

Table 5.4 Some of the responses of Group X to isolated words before and after prolonged experience of a sentence with a low mean formant one

		beat	bit	bait	bet	bat	bought
Word A	Initial	-	7	2	3	-	-
	Final	-	9	1	2	-	-
Word B	Initial	-	4	4	3	1	-
	Final	-	1	-	11	-	-
Word C	Initial	-	-	-	5	6	-
	Final	-	-	-	3	9	-
Word D	Initial	-	-	-	1	2	-
	Final	-	-	-	-	1	1

were equal numbers of each version: the order of test words was the same as in Group X. After these fifty-six recordings the eight with the ten second gap were repeated, and finally the four isolated test words.

Group Z had the same experience as Group Y, but were given knowledge of results after each sentence. They heard exactly the same recordings as Group Y. During the first fifty-two of the main block of recordings, however, the experimenter called out the "correct" answer after each combination of sentence and word. The correct answer given was the most common response accorded to this combination of introduction and test word in the experiment reported in the previous chapter. Thus throughout the fifty two items in the current experiment subjects were being told to judge differently if the introductory sentence was different. The answer was indicated by numbering the answers on the subjects' sheets and calling out a number, so as to avoid the possibility of mishearing the experimenter's vowel. The subjects entered this answer on their sheet to ensure their attention.

It was thought before the experiment that prolonged experience might shift judgment of the isolated words in one direction or the other, especially in Group X. In fact, no consistent change in either direction occurred, the normal result being that the final judgements showed less deviation from the modal response than the initial judgments did. Some of the results of Group X are given in table 5.4 as an example It will be seen that although there was some sign that the responses to B, C and D shifted towards the right of the

Table 5.5. Responses of group Y to test word B 10 secs after an introduction, before and after prolonged experience of two versions of the introduction.

		beat	bit	bait	but	bat	bought	but
initial	version 3	-	6	2	4	-	-	-
	version 2	-	1	1	9	1	-	-
final	version 3	-	3	3	6	-	-	-
	version 2	1	2	1	8	-	-	-

table, those to A shifted if anything to the left. None of these shifts are significant statistically. Thus it does not appear that prolonged experience of a synthetic voice with low first formant frequency produces a drop in adaptation level.

Furthermore in the situation in which the words followed the sentences after a ten second interval expectations were not borne out. It was thought that the effect of an introductory sentence might last longer if subjects had repeatedly experienced that sentence, either alone or in contrast with another. Thus the final group of tests might show more effect than the initial ones. This was not so: if anything there was less effect. The results of Test Word B for Group Y are shown in table 5.5 since they include the only comfortably significant result observed: Test Word B gave different responses in the appropriate direction with Versions 2 and 3 in the intial tests for seven subjects none showing an opposite effect. In the final tests, however, the distribution of responses to the two Versions was practically identical and quite insignificantly different.

In the case of words immediately following sentences the chief interest attaches to Test Word A following Version 2 and Test Word B following Version 3, since these are the two combinations which according to earlier results should give a substantial shift in response away from that given to the isolated word. The results for these combinations are given in table 5.6 for each group. The initial category gives the number of responses of each type for the first occasion on which the combination occurred: and the final category that for the last occasion. There was no intial category for Group Z, since they were being told what to answer by the experimenter

Table 5.6 Responses to words immediately following sentences before and after prolonged experience of the situation.

Version 2, test word A (heard as bit when isolated)

		beat	bit	bait	bet	bat	bought	but
Group X	initial	-	2	-	9	1	-	-
	final	-	9	3	-	-	-	-
Group Y	initial	-	8	2	2	-	-	-
	final	2	5	2	3	-	-	-

Version 3, Test word B (heard as bet when isolated)

Group Y	initial	-	10	2	-	-	-	-
	final	-	6	2	4	-	-	-
Group Z	final	-	10	-	2	-	-	-

at that stage of the recording and so all of them marked their sheets in the pattern of table 5.2.

It can be seen that although the now familiar effect appears in the initial judgment it weakens after prolonged experience. Group X, who heard sixty presentations of Version 2, gave a modal response of bet to Test Word A initially and of bit finally. Ten subjects showed a change in this direction and none in the opposite sense, so that change is significant well beyond the .01 level. There are signs of a similar shift in Group Y, whose experience was of both Versions, but the statistical significance is less satisfactory. The change of Group Y on Test Word A is quite insignificant, possibly because this group happens to have a pronounced tendency to give bit responses to any test word even initially. Group Y does show a just significant trend in its responses to Test Word B and Version 3. Five subjects modify their responses from bit or bait in the direction of bet, and none do the opposite. The final responses are not significantly different from the responses to isolated words obtained at the beginning of the run.

It only remains to note that Group Z, the knowledge of results group, is quite different from Group Y. They show no sign at all of wearing-off of the effect. Their final judgments on Test Word B with Version 3 are almost exactly the same as the initial ones of Group Y, and are significantly different from the responses of Group Z to Test Word B in isolation; nine subjects showed a difference in this direction and none in the opposite sense, so $p < .01$. This result is very much in accordance with everyday experience: we do not find that a prolonged oration from one voice causes the speaker to

become unintelligible, although he may be boring. In everyday life we are, of course, provided with a check upon the accuracy of our adjustment to a speaker's voice: if we are not identifying his vowels correctly he will not make sense. The real-life situation is therefore usually analogous to that of Group Z.

To summarise the results of the third series of experiments, repeated experience of this vowel judgment situation does not cause a permanent bias of response in one direction or another: if anything it produces a loss of the effect. The only exception to this rule is the case in which knowledge of results is given to the subjects. This result, like some of those in the second series of experiments, is inconsistent with the simple type of adaptation level theory given at the beginning of the paper.

There are therefore two cases in which responses to vowels behave in a fashion requiring further explanation. One of these is that an alteration of the Formant I range in a sentence by altering the words used does not appear to shift the judgment of subsequent words: it is only when the Formant I range is shifted by altering the physical sounds used for particular words, that the listener is led to alter his judgments. The other case requiring explanation is that of repeated presentation of test words and sentences, which causes the effect to disappear.

These difficulties are in some ways similar to those mentioned earlier as having arisen in other psycho-physical situations: for example, Brown's demonstration that a heavy weight has no effect on judgements of weight if it is regarded as irrelevant to the other weights lifted. It was said previously that these earlier experiments suggest that past responses as well as past stimuli should be taken into account

in determining the adaptation level. For instance, we might suggest that the response which will be given to a stimulus corresponding to the weighted mean of all past stimuli is not necessarily the neutral or middle response; instead the response given will be the mean of all past responses. In many cases this modification will make no difference, since after a long series of trials the mean of all past responses will in fact be the middle of the response scale. The modified theory does, however, lead one to expect that extreme stimuli will not have a marked effect if they are recognised as such and given extreme responses. Thus in our experiments a vowel with a low formant one position will not bias subsequent perception if the situation requires it to be recognised as the vowel in please rather than that in what - unless the formant is placed unduly low even for a vowel in please.

On the other hand, we might modify the theory of adaptation level more thoroughly, and suggest that each category of response is separate and forms its own adaptation level. Either type of modification will explain the departure of the second and third series of experiments from the simple theory. In the second series, changing the words of the introductory sentences changes the response to them as well as the mean level of Formant I; whereas changing the Formant I position in a sentence whose wording remains constant alters the stimulus without altering the response. The latter case will therefore produce an effect on subsequent stimulus-response relationships while the former need not.

In the third series, the response is often made to the test words alone. In single trials they may be influenced by the introductory sentence to which no overt response is required.

When trials are repeated over and over, there is a context of other test words which is more important than that of introductory sentences, which probably come to be regarded simply as a cueing signal for the test word. This is particularly true becase the speech is in fact artificial and thus may become increasingly unconvincing on repeated hearing, since no human voice can repeat a phrase identically in this way. So the effect of the introductory sentence is negligible in that the words in it are not considered as sounds containing particular vowels. On the other hand, the knowledge of results experiment, on this interpretation, continues to show the effect of the variations of the introduction since the experimental situation ensures that the sentence receives a response.

With such modification, then, the theory of adaptation level is adequate for the case of vowel judgments. But it is certainly not adequate if it is conceived as a close analogue of sensory adaptation, so that only the stimulation received need to be taken into account in prediction.

6. Judgments of vowel quality.

We may profitably begin this final chapter by reviewing the conclusions which we have reached in the last three chapters. It seems that:

1. The acoustic quality of most vowel sounds can be conveniently specified by stating the frequencies of their first two or three formants.

2. This is not true of vowels which are called in traditional terms close vowels, nor of so-called back vowels. It is not at all easy to analyse these vowels in terms of their formants; nor is such an analysis an adequate account of their spectral characteristics.

3. The perceptual quality of a vowel usually depends on the relationship between the pitches of the formants of that vowel and the pitches of the formants of other vowels pronounced by the same speaker.

4. The listener to speech uses his past experience to form an adaptation level, the immediate past experience of a particular voice being the most important factor in this process.

5. Neither of points 3 and 4 above has been shown to be true for the vowels mentioned in 2 above.

This is as far as we can go in the present state of our knowledge. The next major steps will probably involve a large number of psychophysical experiments on the perception of quality differences.

At the moment we do not know which are the important auditory cues for a listener assessing the qualities of close or of back vowels. We can therefore make neither a satisfactory acoustic analysis nor an adequate representation of the correlated auditory characteristics of these vowels.

Accordingly, we must now discuss the hypothesis that the most useful form of description for a phonetician who wishes to describe a vowel to other phoneticians in such a way that they can produce similar scunds without ever having heard the original sound (or a recording of it) may be in terms of the cardinal vowel system. This hypothesis was evaluated by means of the following experiment.

A recording was made of the following ten Gaelic words as pronounced by a native Gael: (a) Bèid (b) sgó (c) cùl (d) reub (e) lon (f) bìg (g) fàl (h) laochan (i) stagh (j) gaoth. These words contained vowels which were as monophthongal as possible, but which differed greatly in their phonetic quality. Eighteen phoneticians were asked to plot the vowels in these words on cardinal vowel diagrams which were provided. Fifteen of these subjects were trained in the British tradition of Phonetics, and consequently had had extensive practice in performing cardinal vowels, and assessing the quality of other vowels in terms of them. The other three subjects were experienced phoneticians well acquainted with the theory of cardinal vowels, but they had not undergone the rigorous formal training in the use of the system which had served to provide the first group of subjects with fixed reference points. Details of the background of each subject are given in table 6.1. Each subject listened to the recording by himself, playing it back as often as he wished

Table 6.1 The eighteen phoneticians who took part in the experiment of judging the quality of the vowels in ten Gaelic words.

No.	Subject	Status at time of expt.	Cardinal vowels learnt from:
1	D. Abercrombie	Staff, University of Edinburgh	Daniel Jones
2	E.T. Uldall		
3	J.C. Catford		
4	P. Ladefoged		
5	J. Woolley	Post-graduate students, University of Edinburgh.	D. Abercrombie, E.T. Uldall & J.C. Catford.
6	L. Criper		
7	A. Vos		
8	A Rodger		
9	M. Fraser		
10	A.C. Gimson	Staff, University College, London	Daniel Jones
11	J.D. O'Connor		
12	G.F. Arnold		
13	O. Tooley		
14	J.L. M. Trim		
15	P.D. Strevens	Staff, University College of the Gold Coast.	H. Coustenoble
16	G. Hammarström	Staff, Uppsala University	-----
17	S. Sapon	Staff, Ohio State University	-----
18	T. Hill	Staff, University of Edinburgh	-----

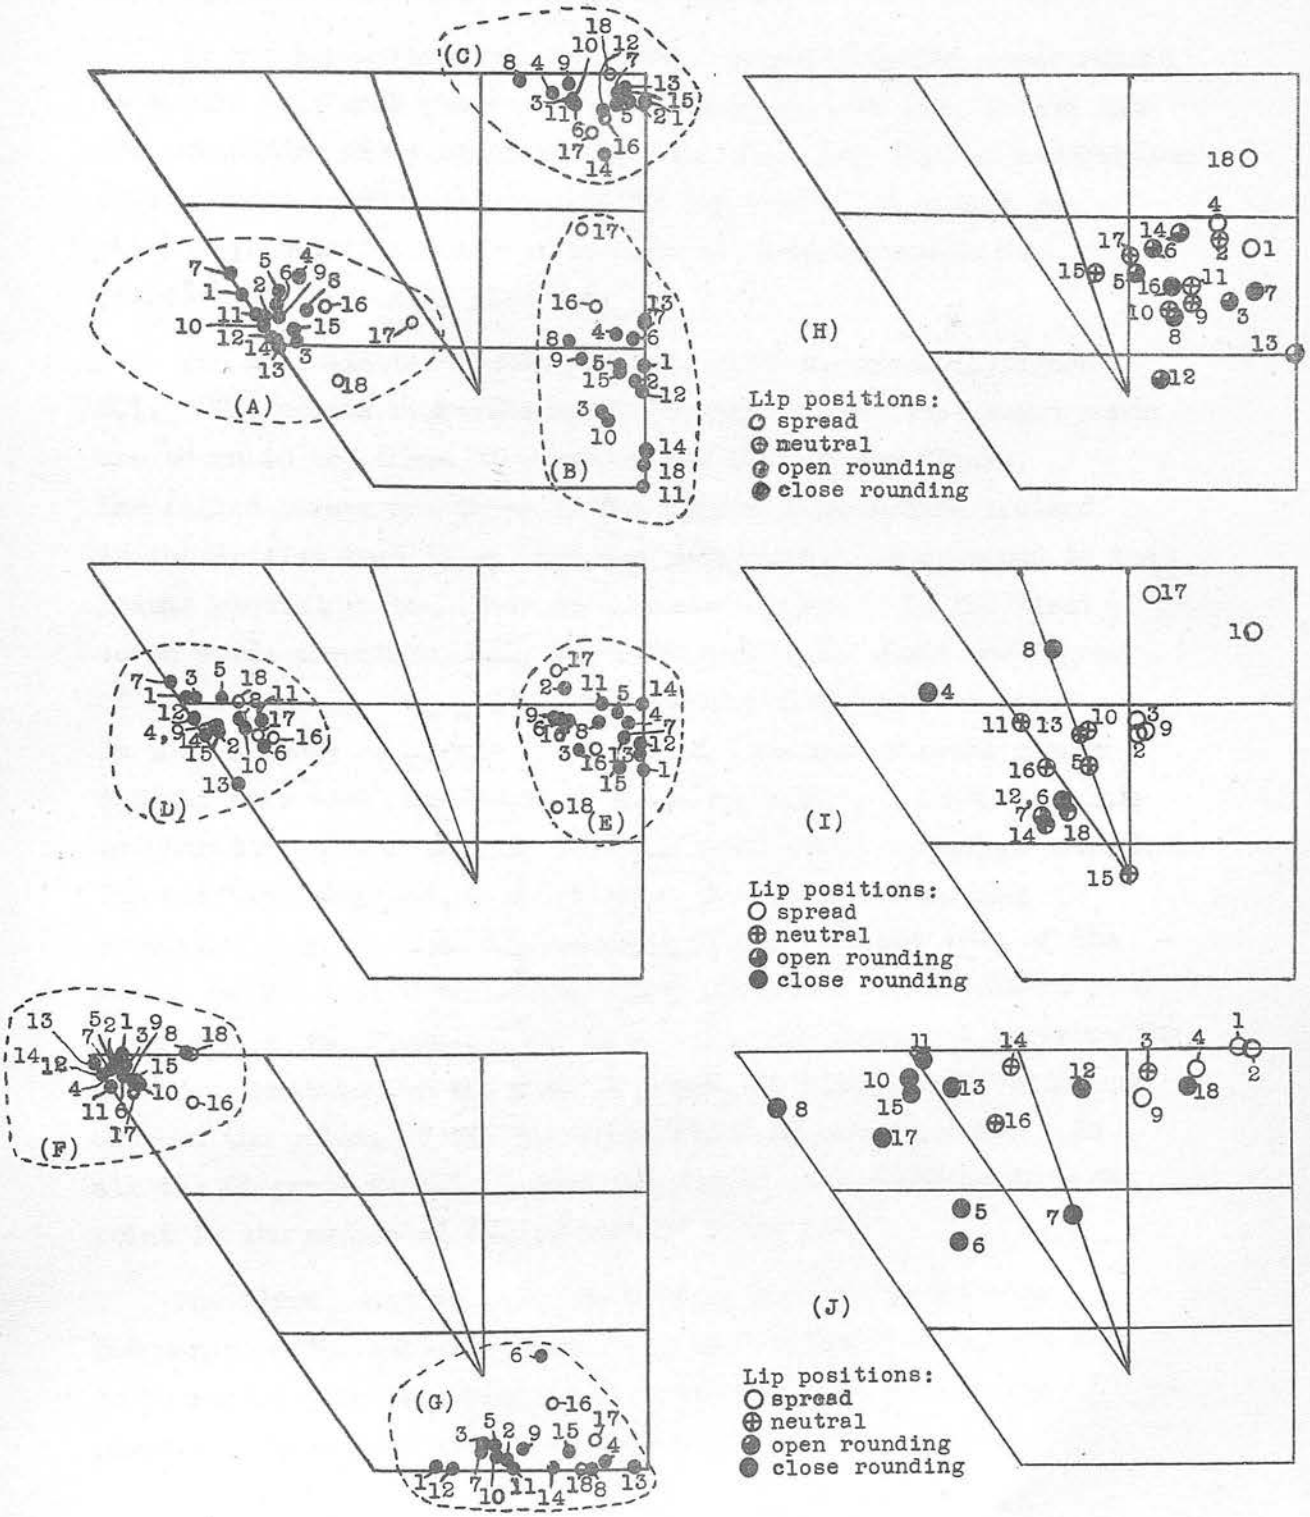


Figure 6.1 The judgments of 18 phoneticians of the vowels in ten Gaelic words.

and in any way that he found convenient. None of the subjects was acquainted with this variety of Gaelic.

It was hoped that this procedure was sufficiently standardised to ensure that each phonetician was assessing the same phonic data and presenting his results in the same way; but that it nevertheless corresponded as much as possible to the typical situation in which a phonetician needs to be able to describe vowels for purposes of linguistic research.

The consolidated results of this test are shown in figure 6.1. The points representing the vowels in the first seven words are shown in the three diagrams on the left of the figure. The filled points are those of the fifteen phoneticians trained in the British tradition; and the open circles correspond to the points plotted by the other three phoneticians. In the first seven words there were only minor disagreements about the degree of lip-rounding; e.g., some phoneticians described the vowel in word B (sgò) as having "rather open rounding" whereas others thought this vowel had "open to close rounding". In these words no phonetician ever thought that any vowel had a spread or a neutral lip position when other phoneticians thought that the same vowel had open or close lip-rounding. This was not true of the vowels in the last three words, which are shown in the diagrams on the right of the figure. In these words the degree of lip-rounding has been indicated by the kind of point, no distinction being made between the points of the different kinds of phoneticians. In all the diagrams in this figure the number corresponding to each point is the number of the subject in table 6.1.

The first point to note about these results is that the judgments of the subjects trained in the British tradition seem to be usually more in agreement than those of the three other phoneticians and any twelve of the first fifteen subjects.

Thus the points for the vowels in words, A, B, E and F as represented by subjects 16, 17 and 18 are comparatively widely scattered, showing a lack of agreement both among themselves and with the other subjects.

All the last three subjects were good phoneticians with a knowledge of many different languages and experience of dialectology. They thus had a higher professional standing and a great deal more experience than subjects such as the post-graduate students or the junior members of the staff of the University of Edinburgh. But they are nevertheless relatively unable to communicate in writing with one another in an unambiguous way about the quality of a vowel sound.

The superiority of the agreement among phoneticians trained in the British tradition cannot be explained as being due to aspects of their linguistic background which had nothing to do with their phonetic training. Admittedly eleven of them were speakers of that variety of English known as RP (Jones 1957); but the other four were an American (E.T. Uldall), a Dutchman (A Vos), a Scot (A. Rodger) and a Yorkshire girl (M. Fraser). However, despite these diverse backgrounds these subjects were not in disagreement with the other eleven as were the three phoneticians, a Swede (G. Hammarström), an American (S. Sapon) and an Englishman (T. Hill), who had had no formal training in the cardinal vowel system. We may conclude, therefore, that in so far as there is agreement among the phoneticians who have graduated in the British tradition, it is to some extent due to the rigorous training which they have all undergone.

Table 6.2

	Word	Minimum area		Rank order
		sq.mm.	% of total area	
A	Beid	84	1.5	3
B	sgò	206	3.8	7
C	cùl	60	1.1	2
D	reub	90	1.6	4
E	lon	135	2.5	6
F	big	30	0.6	1
G	fàl	114	2.1	5
H	laochan	396	7.3	8
I	stagh	604	11.1	9
J	gaoth	1449	26.7	10

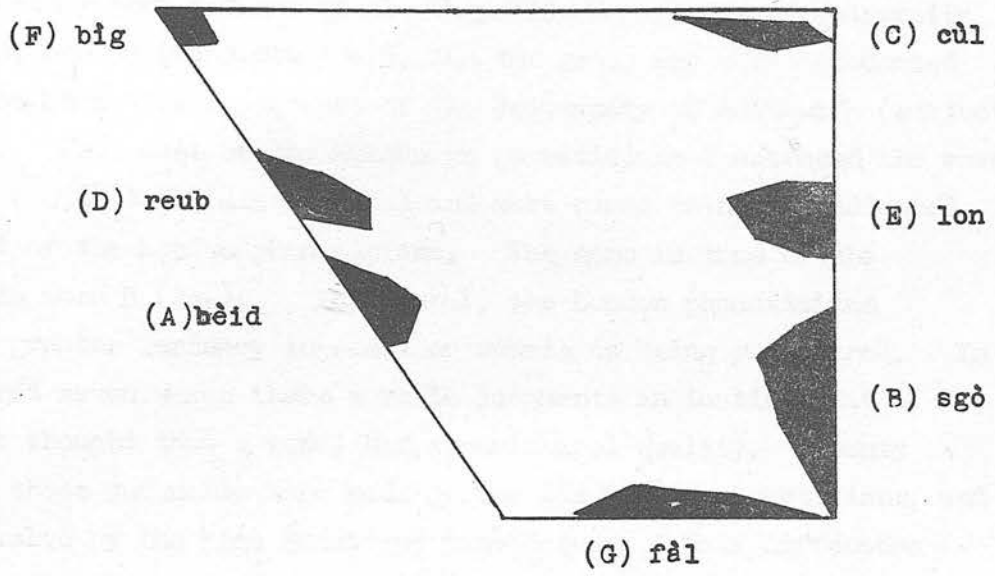


Figure 6.2 The minimum areas containing the points indicating the judgments of at least 14 out of the 15 phoneticians trained in the British tradition.

There are, however, disagreements even among these phoneticians. Some of these are between the group of phoneticians who were or had recently been members of the Phonetics Department of University College, London (subjects 10-15) and the group who were associated with the Phonetics Department of the University of Edinburgh (subjects 1 - 9). Thus most of the Edinburgh phoneticians considered the vowel in word A (bèid) to be more central and more close than was indicated by most of the London phoneticians. The same is true of the vowel in word B (sgɔ). In general, the London phoneticians have a greater tendency to consider vowels as being peripheral. In the first seven words there were 34 judgments indicating that a subject thought that a vowel had a peripheral quality. Twenty two of these judgments were made by the six London phoneticians, and only twelve by the nine Edinburgh phoneticians. This difference is very significant ($X^2 = 10.5$, $p < 0.01$).

It is also apparent that in some words the agreement is much greater than in others. Figure 6.2 shows for each of the first seven vowels the minimum area which will contain the points representing the judgments of more than 90% (i.e., fourteen out of fifteen) of the subjects trained in the use of the cardinal vowel system. Comparison with figure 6.1 will show how these areas have been drawn. Table 6.2 shows the size of the area for each of the ten vowels expressed both in square mm and as a percentage of the total area of the diagram. When the judgments of the vowels in the first seven words are compared in this way, it may be seen that the degree of agreement is often very great (far greater than might appear from figure 6.1, where the size of the points and the attached numbers tend to exaggerate the extent of the spread of the points).

Some interesting facts emerge from a study of the rank ordering

of these areas. Thus the vowels in words C (cùl) and F (bìg) are both of the kind in which it is difficult to locate the exact centre of the first formant because of its proximity to the fundamental. But nevertheless these are the vowels concerning which subjects are in the greatest agreement; and the major part of what disagreement there is in the case of word C (cùl) is in judgments of the front-back dimension, which, according to the usual formant theories for vowels of this type, can be correlated with the frequency of formant two. This agreement of judgments of the close-open dimension in the case of vowels in which it is difficult to locate the centre of the first formant supports the hypothesis that these vowels are assessed in terms of some quality other than the pitch of formant one.

After the judgments of the vowels in words F (bìg) and C (cùl) the next greatest agreement is in the judgments of the vowels in words A (beid) and D (reub). These vowels are of the type which it is comparatively easy to analyse and specify in terms of their formant frequencies. The vowels in words G (fàl), E (lon) and B (sgò) however are more difficult to specify in this way; and there is also more disagreement in the judgments of their qualities. In addition, as we found in chapter three, a vowel in the neighbourhood of cardinal six - i.e., that in word B (sgò)-produces more difficulties than one near cardinal seven - i.e., that in word D (lon). We may conclude therefore, that with the important exception of vowels of the of the [i] and [u] type, the degree of difficulty which phoneticians have in auditorily assessing vowel qualities is paralleled by the difficulty of specifying similar vowels in terms of instrumental data.

The vowels in the last three words cannot be considered in the same way, because there has been no instrumental analysis of a large number of vowels of these types. (This is obviously a task which must be undertaken as soon as possible). It is readily apparent that subjects have a great deal of difficulty in assessing the quality of these relatively unfamiliar vowels. As in the case of the vowels in the first seven words, they would no doubt have found the task easier if they had been listening to the informant himself instead of a recording. There might then have been less disagreement about the degree of lip-rounding; and this in its turn would probably have resulted in greater agreement in the judgments of the other aspects of vowel quality. Nevertheless it should be noted that in each of these three vowels there is a wide scatter of judgments, even among those who agree about the lip positions.

The great disagreements about the degree of lip-rounding indicate that this feature of vowel quality is not easy to assess in auditory terms alone; and, in fact, although eleven out of the eighteen subjects considered that the vowel in word J (gaoth) had close lip-rounding, all the subjects who met the informant after the experiment then considered that this vowel had a spread or a neutral lip-position. Furthermore the judgments of this vowel and of the vowel in word I (stagh) indicate that the degree of lip-rounding is not always considered as an independent variable. There is a very significant correlation ($X^2 = 17.3$, $p < 0.01$) between the judgments indicating open or close rounding which are associated with points in the left hand section of the diagram, as opposed to judgments indicating a spread lip-position which are associated with points in the right hand section of the diagram.

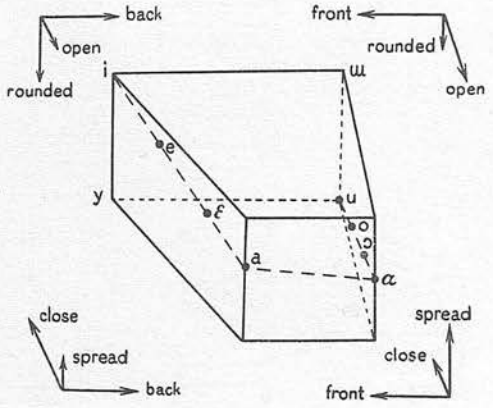


Figure 6.3 A three dimensional vowel diagram.

According to the traditional theory, there are three main parameters of vowel quality which are independently variable: the position of the highest point of the tongue in the close-open dimension, and in the front-back dimension; and the degree of lip rounding. Thus the quality of a vowel may be indicated in terms of a point in a three dimensional space as shown in figure 6.3. The cardinal vowels do not, of course, lie on a surface of this figure, but occur in the positions indicated. We may summarise the viewpoint which is being put forward here by suggesting that these dimensions are not the most appropriate parameters with which to specify vowel quality. It would be better if phoneticians compared vowels in terms of the two dimensions which specify the plane of the cardinal vowels (perhaps without using articulatory terms which are apt to be misleading if considered literally, but using instead purely auditory terms or descriptions indicating vectors from auditorily known reference points), and also in terms of a third auditory dimension for which there is as yet no convenient name nor adequate data concerning its nature.

These speculations, however, should not lead us to neglect the great value of the existing cardinal vowel system. Considering the judgments of the vowels in the ten Gaelic words, we may conclude that the traditional training in the use of the cardinal vowels allows phoneticians to make very adequate judgments of at least those vowels which are judged as having lip positions like those of similar primary cardinal vowels. It is difficult to produce a valid statistic based on the experimental data reported here, partly because the subjects cannot be considered as a sample of a population and partly because of the inter-dependence of all the judgments. But we can attempt to summarise part of the data in table 6.2 by saying that the mean minimum area for the phoneticians'

judgments of the vowels in the first seven words was under 1/2% of the total vowel area. There is at the moment no other way in which vowels can be specified with equal accuracy. Even the forty two symbols and additional modifiers suggested by Bloch and Trager (1942) will be of no avail unless they constitute reference points that are accurately known to all users of the system. Otherwise it is impossible to say precisely what sound is implied by, for instance, the symbol [E]. This sort of difficulty becomes especially obvious when we try to interpret remarks such as "by [æ] we mean here the cardinal higher low front unrounded lax vowel" (Trager & Smith 1951); and the subsequent statement "many will have a slightly higher and tenser quality say [æ^]". These authors seem to be implying that there is a precisely determined vowel quality, symbolised by [æ] and known to all their readers. But no phonetician relying solely on their descriptions could reproduce these vowel qualities with certainty; the articulatory statements (e.g., about tenseness and laxness) are far too vague and physiologically meaningless. No doubt Trager and Smith and their immediate associates know exactly what vowel sounds are intended by the symbols and the descriptions. But they could convey this information only by oral instruction in their reference points. Any phonetician who wishes to convey the precise difference between two vowel sounds can do so only in so far as there are equally precise reference points known to him and his reader.

At the moment the best means of providing these reference points is by oral instruction in the cardinal vowels; consequently in the present state of our knowledge, to abandon the cardinal vowel system is to abandon the only internationally known method of

specifying vowels at all accurately. This will remain the state of affairs until it is possible to make acoustic measurements which can be used for specifying vowels. We have shown in this thesis that this can be done to some extent; but it seems that we are a long way from a simple acoustic specification which is applicable to vowels of all types, and which allows for the personal features of a speaker's vowels, which a skilled phonetician automatically recognises and discounts. Consequently, the traditional rigorous training in the performance and use of known reference points remains essential for all who wish to give detailed specifications of vowel sounds.

Bibliography

Abercrombie, D., (1954) "The recording of dialect material" Orbis, 3, 231-235.

Abercrombie, D., and Aitken, A.J., (forthcoming) A Scots Phonetic Reader.

Bell, A.G., (1879) "Vowel theories" Amer. J. Otol., 1. (Reprinted in Bell, A.G., (1906) The Mechanism of Speech.)

Bell, A.M., (1849) Principles of Speech [not seen, quoted in Bell, (1867) q.v.]

Bell, A.M., (1867) Visible Speech.

Beranek, L.L., (1949) Acoustic Measurements.

Bloch, B., and Trager, G., (1942) Outline of Linguistic Analysis.

Brown, D. E., (1953) "Stimulus similarity and the anchoring of subjective scales" Amer. J. Psychol. 66, 199-214.

Brightland's (1711) Grammar of the English Tongue.

Campbell,D.R., Lewis,N.A. & Hunt,W.A. (1958) "Context effects with judgmental language that is extra-experimental", J.Exp. Psychol 55,220-8

Crandall, I.B., (1925) "Sounds of Speech" Bell Technical Journal, October. [not seen, quoted in Fletcher (1929) q.v.]

Delattre, P., (1951) "The physiological interpretation of spectrograms" Pub. Mod. Lang. Assoc., 66, 864-875.

Delattre, P., Liberman, A.M., Cooper, F.S. and Gerstman, L.J., (1952) " "An experimental study of the acoustic determinants of vowel colour" Word, 8, 195-210.

Dunn, H.K., (1950)"Calculation of vowel resonances" J. Acoust. Soc. Amer.22, 740-753.

Ellis, A.J., (1885) Translation of H. Helmholtz (1863, 1st Ed, & 1877, 4th Ed) Sensations of Tone, with numerous additional notes, etc.

Essner, C., (1947) "Recherches sur la structure des voyelles orales" Arch. Neerland. Phon. Exper.20, 40-77.

Fant, G., (1956) "On the predictability of formant levels and spectrum envelopes from formant frequencies" in For Roman Jakobson.

Fant, G. (1958a) Acoustic Analysis & Synthesis of Speech

Fant, G. (1958b) Acoustic Theory of Speech Production

Fant, G. (1958c) Modern Instruments & Methods for the Acoustic Study of Speech.

Firth, J.R., (1957) Papers in Linguistics, 1934-1951

Fischer-Jorgensen, E., (1958) "What can the new techniques of acoustic phonetics contribute to linguistics?" in Proc. of the VIII Internat. Congress of Linguists

Fletcher, H., (1929) Speech & Hearing

Greenwood, J., (1711) An Essay Towards a Practical English Grammar

Helmholtz, H., (1863) [not seen, translated by Ellis (1885) q.v.]

Helson, H. (1948) "Adaptation level as a basis for a quantitative theory of frames of reference" Psychol. Rev. 55, 297-313

Hockett, C.F., (1955) A Manual of Phonology
Holder, W. (1669) Elements of Speech
Jones, D., (1917) An English Pronouncing Dictionary (1st Ed.)
Jones, D., (1956) An Outline of English Phonetics
Jones, D., (1956a) English Pronouncing Dictionary (11th Ed)
Jones, S., (1929) "Radiography & pronunciation" Brit. J. Radiol. 2.15, 149-50
Joos, M., (1948) Acoustic Phonetics
Koenig, W. (1949) "A new frequency scale for acoustic measurements" Bell Labs. Record 27.299-301.
Ladefoged, P., & Broadbent, D.E., "Information conveyed by vowels" J. Acoust. Soc. Amer. 29.1, 98-104
Lawrence, W., (1953) "The synthesis of speech from signals which have a low information rate" in Communication Theory.
Lehnert, M. (1936) Die Grammatik des englischen Sprachmeisters John Wallis (1606-1703)
Lloyd, R., (1891) "On recent progress in the analysis of vowel sounds" in Report of the Brit. Assoc.
Munson, W.A., & Gardner, M.B., (1950) "Loudness patterns - a new approach" J. Acoust. Soc. Amer. 22, 177-190
Newton, I. (c1665) (Photostats of the original notebooks in the British Museum)
Paget, R. (1923) "The production of artificial vowel sounds" Proc. Royal Soc. A. 102, 752-65
Passy, P. (1887) "Kurze darstelling des Französischen Lautsystems" Phonetische Studien, 18-40
Peterson, G.E. (1951) "The phonetic value of vowels" Language 27, 541-553
Peterson, G.E. (1952) "Information bearing elements of speech" J. Acoust. Soc. Amer. 24, 629-37
Peterson, G.E. & Barney, H.L. (1952) "Control methods used in a study of vowels" J.Acoust. Soc. Amer. 24, 175-84
Pike, K. (1943) Phonetics
Potter, R.K. et al. (1946) "Technical aspects of Visible Speech" J. Acoust. Soc. Amer. 17, 1-89
Potter, R.K. & Peterson, G.E. (1948) "The representation of vowels and their movements" J. Acoust. Soc. Amer. 20, 528-35
Potter, R.K. & Steinberg, J.C. "Towards the specification of speech" J. Acoust. Soc. Amer. 22, 807-20
Robinson, R. (1617) The Art of Pronuntiation
Russell, G.O. (1928) The Vowel
Sherif, M., Taub, D., & Hovland, C.I. (1958) "Assimilation & contrast effects of anchoring stimuli on judgments", J.Exp.Psychol. 55, 150-5
Stevens, K. & House, A. "Development of a quantitative description of vowel articulation" J. Acoust. Soc. Amer. 27, 484-93

Steven, S.S. & Volkman, J. (1940) "The relation of pitch to frequency - a revised scale." Amer J. Psychol. 53, 329-38
Sweet, H. (1890) A Primer of Phonetics
Trager, G. & Smith, H.L. (1951) An Outline of English Structure
Wallis, J. (1653) Grammatica Linguae Anglicanae
Ward, I. (1945) The Phonetics of English
Wheatstone, C. (1837) [an untitled review over the initials "C.W." in The London & Westminster Gazette, October, 1937, pp 27-41]
Willis, R. (1829) "On the vowel sounds and on reed organ pipes" Trans. Camb. Phil. Soc. 3.