

# 1

## Recording the Sounds of a Language

### 1.1 Deciding What to Record

We'll start with the overall aim of describing the major phonetic structures of a language – where the consonants are made, what kinds of articulation are involved, how the vowels differ in quality and length, how the pitch varies in different phrases, and other straightforward phonetic properties. What do you need to know before you can describe any of these characteristics? The first point, I would suggest, is how the sounds are organized. Every language has a certain number of contrasting sounds that can change the meanings of words. Discerning what sounds contrast, and how they can be combined to form words, is a major part of phonology. Without knowing the phonology of a language you cannot describe the phonetics. You need to know what it is that you have to describe.

Of course, without some knowledge of the sounds, you cannot describe the phonology of a language. It is a chicken and egg problem. The phonology has to be clear before you can make a meaningful description of the phonetics; and without a description of the sounds, you cannot get very far with the phonology. The two kinds of investigation have to advance hand in hand. Usually when you start working on the phonetic structures of a language, you will be able to find some previous work that will be helpful. Ideally much of the phonology will have been worked out already. In this book we will assume that there has been some prior work on the languages that we are going to describe. We will take it that at least the sounds that contrast in words – the phonemes – have been described to some extent. But, as we will

## 2 Recording the Sounds of a Language

My first fieldwork trip was not a very happy experience. I was a graduate student at Edinburgh University and my supervisor, David Abercrombie, suggested that I should go to the Outer Hebrides (the chain of islands off the north-west coast of Scotland) and learn something about the sounds of Gaelic. I had no idea what to do, and came back after a few days without achieving anything. I had to hide in my room for a week, because I did not want Professor Abercrombie to know that I had come back early. I hope this book will help you do better.

note later, you should never fully trust anyone else's description of the sounds of the language you are investigating. They may have been describing a different dialect, or the language might have changed since their account of it. Or they might have been wrong.

We'll consider first how to investigate the consonants. If we had all the time in the world, we might make a list of words illustrating every consonant before and after every vowel. In most languages, however, consonants have similar characteristics before most vowels, differing only before high vowels when they may become affricated. Given the practical limitations within which we all have to work, a good starting point is a list that illustrates each consonant before two vowels that are very distinct from one another, such as *i* and *a* if they occur in the language. If the language allows syllable-final consonants, then they should also be illustrated after two different vowels. Syllable-final consonants are often very different from syllable-initial consonants. For example both the consonants in the English word *leek* are very different from their counterparts in *keel*. Initial **k** is aspirated, whereas final **k** may be unreleased and preceded by a glottal stop in my English. Initial **l** in this context is slightly palatalized, whereas final **l** is velarized. All such variations should be illustrated in a complete account of the sounds of a language. A minimal word list for English might include *pie, tie, kye, . . . seep, seat, seek, . . . type, tight, tyke, . . .*, and so on. But, given enough time, it should also include consonants with many other vowels.

How do you go about finding all these words? The obvious place to start is with a dictionary of the language, if there is one. You will need to begin by studying the spelling system and working out which sounds correspond to which letters. If you were going to be recording English, for example, you would need to note that words illustrating initial **k** could be found under both 'c' and 'k', and that words illustrating both

My second fieldwork trip was much more enjoyable. By then I was married to Jenny, who had none of my fears about walking up to a stranger's door and asking to record his vowels. We were working for the Linguistic Survey of Scotland, and had been thoroughly prepared by Ian Catford. He told us how to find speakers, and provided us with word lists that he wanted recorded. The weather was sunny, and the rivers teeming with salmon swimming up to spawn. Fieldwork provides opportunities for seeing many wonderful sights. On later occasions we've seen thousands of pink flamingos rising from a lake, and Everest wreathed in clouds.

**θ** and **ð** could be found by looking for the spelling 'th'. Go through the whole dictionary carefully, noting how many pages are given to each consonant in initial position, and then, while looking more closely at a couple of the more common consonants, find out how many columns are devoted to each vowel. By using the most common consonants and vowels you can usually find good sets of minimally contrasting words.

What do you do if no dictionary is available? The prior work on the language may not include a dictionary. It may be just a few notes by a linguist and perhaps a folk tale or short illustrative sample. Even this can provide a useful start. Type the story and any other words available into a computer and then sort them alphabetically. (Replace all spaces between words by paragraph marks, and then get the computer to sort the list.) Even a short word list like this can be useful in finding words. It will give you an indication of which are the most frequent sounds, and where you will find it easiest to find sets of minimal contrasts.

When you are looking around for sets of words, another major source is the speakers of the language themselves. Your short word list may have led you to find a set that almost illustrates the sounds you are interested in, but it lacks a particular example. Make up a word that might be appropriate, and ask the speakers if it exists. You may find that is the name of one of their clans, or a particular kind of tree, or some other perfectly acceptable but less frequently used word. Find out if your speakers can rhyme words or use alliteration in any way. This may help them produce good contrasting sets. If you are working on a well-known language, you might be able to find a rhyming dictionary (there are several for English), giving you a direct lead on minimal sets such as *pea, tea, key, . . .*

#### 4 Recording the Sounds of a Language

Generally speaking, it is best to avoid nonsense words, particularly if you are going to be dealing with naive speakers. Many people find it difficult to pronounce a nonsense word in a natural way. So should my list for English have included *kye*, a perfectly good word that has dropped out of use and is not in the vocabulary of most speakers of English? Usually literate speakers can pronounce simple forms such as this without any difficulty, particularly if you explain that *kye* is an older variant of the plural of *cow*. In any circumstances, you should always stick to words that speakers feel comfortable with. You might have to remind them that the word in question is what hunters call a particular type of arrow, or a bird that is rarely seen nowadays. But in the end, all that matters is that you have a list of words that speakers can pronounce in a natural way, with the true sounds of their own language.

It is very important to have a set of words that are as near minimal contrasts as possible, even if it means including a few unusual words. When it comes to characterizing the difference between sounds, the surrounding sounds should be identical. For example, as we will see in chapter 5, when you describe the vowels of a language you will find differences in the qualities of a set of vowels after **p** in comparison with those after **t**. A single list with some vowels after the one consonant and others after the other would lead to a false description of the differences in vowel quality. In most languages vowels differ even more before different syllable-final consonants. To take the example we looked at before, consider the vowels in *keel* and *leek*. Because of the final **l**, the vowel in *keel* is usually much more diphthongal than that in *leek*.

It is usually impossible to get a complete list illustrating all the consonants before the same vowel. Some consonants may occur only at the ends of words, or only after certain vowels. There is, for example, no set of words illustrating all the fricatives of English in the same circumstances, largely because English originally had no contrasts between voiced and voiceless fricatives, and the sound **ʒ**, a very recent import, occurs at the beginnings of only a few foreign words such as *Zsa Zsa* and at the ends of loan words such as *rouge*. You just have to find the best set you can, using a number of different possibilities. You might start with *fie, vie, thigh, thy, sigh, Zion, shy*, which contains all the fricatives in initial position before **aɪ** except for **ʒ**. Another set such as *proof, prove, sooth, soothe, loose, lose, rouge* contains all the fricatives in final position after **u:** except **ʃ**. Then you could add one of the few

contrasts between **f** and **ʒ**, *Aleutian* vs. *allusion*. Not very satisfactory, but there is no perfect answer for this set of sounds.

Linguists refer to a pair of words that contrast in only one segment as a minimal pair. A set of words that is the same except for a single segment is a minimal set. In Hawaiian, for example, it is easy to find a minimal set contrasting all the consonants (there are only 8 of them, **p, k, ʔ, m, n, w, l, h**). It took me only a few minutes using a Hawaiian dictionary to find: **paka, kaka, ʔaka, maka, naka, waka, laka, haka**, meaning *to remove the dregs, to rinse or clean, to laugh, eye, to quiver or shake, sharp or protruding, tame or gentle, shelf or perch*. With a set of words like this, we can compare the properties of each consonant with those of the others, knowing that they are not being affected by the context. If Hawaiian had not had the word **paka** (or you and your consultant had not been able to find it), you might have used the word **pika**. But then, when you compared the duration of the aspiration in **p** and **k** (a topic we will discuss more fully in chapter 4), you would have found that **p** had more aspiration than **k**. But, other things being equal (i.e. when they are in minimally contrasting words), it is the other way around; Hawaiian **k** has slightly more aspiration than **p**, as it does in most languages of the world. It is just that before **i** all stops generally have more aspiration. The initial **k** in the Hawaiian word **kika** has even more than in **kaka**.

Dictionaries often illustrate the sounds of a language (meaning the phonemes) simply as they appear in common words, rather than in words that are minimally contrasting. Similarly linguists who are not concerned with phonetics sometimes list the phonemes in words that are far from minimal sets. If the primary aim is to describe the syntax of a language, then there is little need to know more than that the sounds all differ from one another. But from a phonetic point of view it is well worth taking an enormous amount of trouble to find the best possible minimally contrasting sets of words.

You should also check that you have all the possible sounds. As I mentioned earlier, the source that you are using as a basis for making a word list may not be fully accurate. Your speakers will tell you if they no longer have a certain sound, or consider it to belong to another dialect, but they may not think of pointing out additional sounds. Look at the consonant inventory and see if there are any obvious gaps. For example, the language may have **p, t, k** and the ejectives **tʼ, kʼ**. Ask if there are any words beginning with **pʼ**. There may not be, as many languages have **tʼ, kʼ** but lack **pʼ**, but you should

Always when doing fieldwork you must be sensitive to the culture around you. Some words may be taboo because they are associated with people who have recently died. Others may be parts or functions of the body that are not mentioned in polite company. On some occasions, when looking for a particular sound in a particular context, I've said 'Do you have a word that sounds like such-and-such?', and provoked roars of laughter or embarrassed looks. In a polite book like this, I can't tell you about some of the words I hit on. When an unexpected reaction occurs I just go on with something different.

check for missing possibilities. Similarly, if neighboring languages or languages of the same family have sounds that do not seem to be in the language you are investigating, ask about them, demonstrating them in a variety of simple syllables, and inquiring whether there are any words with those syllables. You may find that your language consultants come up with sounds that are not in your sources.

A language will have a certain set of phonemes that form the contrasts between words. Each of these phonemes will have a number of allophones – members of a phoneme that occur in specific contexts. The word list should illustrate the principal allophones of consonants. If it already contains syllable-initial and syllable-final consonants, many allophonic differences will have been included, but there may be other interesting allophones that should be noted. A word list for American English, for example, should illustrate what happens to *t* in *pity* (where it is a tap) and in *button* (where it is often accompanied or replaced by a glottal stop). A really detailed study of English would include all the consonant clusters that can occur. But as there are approximately 105 ways of beginning an English syllable, and about 143 ways of ending one (depending on what you count as a cluster), this would be excessive for most purposes.

Illustrating vowels follows along the same lines. Make a list that has all the vowels in as similar contexts as possible. Sometimes that's easy. Japanese has several sets such as **ki**, **ke**, **ka**, **ko**, **ku**, all of which are meaningful words. In other languages it may be more difficult. You can very nearly get a complete set of English vowels in monosyllables beginning with **b** and ending with **d** or **t**, but to be complete you have to include another set beginning with **h** and ending with **d**, so as to get the vowel in *hood*. In any case, include sets of vowels after different

consonants. If you take the Japanese set **ki, ke, ka, ko, ku** you will find that the vowels may be a little different (they are usually slightly higher) than they are in Japanese **mi, me, ma, mo, mu** (but in this set they may be slightly nasalized).

Vowels often have important allophones that should be illustrated. In English, for example, the set of vowels that can occur before **ŋ** is restricted. There are different vowels in *seen* and *sin*, but before **ŋ** in *sing* there is no such contrast. In Californian English the vowel in *sing* is closer to that in *seen*, but in most forms of British English it is closer to that in *sin*. Special lists are needed to illustrate these context-restricted systems. Consonants often have a noticeable effect on vowel quality. Vowels before final **l** and **r** are very different from those before final **d** in English, and (to use more distant fieldwork examples) vowels adjacent to the uvulars **q, ɣ** are lower in Aleut.

Many languages contrast oral and nasal vowels. In some languages vowel length is important. These contrasts have to be illustrated, along with any other contrasts such as those involving different voice qualities, a topic that we will discuss more fully in a later chapter.

It is often a good idea to record words within a carrier sentence, a frame that surrounds the contrast being illustrated. There are two reasons for this. Firstly, whenever anyone says a list of words, there is a tendency to produce them with the special intonation pattern used for lists. If you ask a speaker of English to say a list of words such as *heed, hid, head, had*, the last word will almost invariably be produced with a lower pitch and a longer vowel than it would have had if it had been earlier in the sequence. One can avoid this by adding some unneeded words at the end, making the list, for instance, *heed, hid, head, had, hid, heed*. But a better technique for producing stability in the pronunciation of each word is to put it into a frame such as *Say \_\_\_\_ again*. Even when eliciting lists in this way, it is a good idea to have the first item repeated again; otherwise the whole frame sentence may occur on a lower pitch, as sometimes happens for the last sentence in a paragraph when reading English.

The second reason for using a frame around the illustrative words is that it makes it easier to measure the lengths of the items that contrast. With a frame such as *Say \_\_\_\_ again* one can determine the beginnings and ends of stop consonants, which would be impossible if there were other stops or silence before and after them. It may be necessary to use more than one frame. If you want to measure the length of a word containing only vowels (such as, the words *eye, owe, awe* in English)

## 8 Recording the Sounds of a Language

you couldn't do it if it were in the frame *Say \_\_\_\_ again*, as there would be no separation between the words. The vowels of the words in the frame would run into the vowels of the test words. You would do better with a frame like *Repeat \_\_\_\_ twice*, which lets you see words beginning or ending with vowels.

Whatever language you are working on, while doing the preliminary work of determining a word list you also need to find suitable frames. In French I have used *Dis \_\_\_\_ encore*, the equivalent of *Say \_\_\_\_ again* both in meaning and in enclosing the test word between vowels. In Korean I used **ikəsi \_\_\_\_**, 'This is \_\_\_\_'. Even a simple frame helps speakers get into a set way of talking when recording, so that they say each word in the same manner. A frame ensures that each word occurs in the same rhythmic position, which is important as the position in an utterance can affect the stress pattern and length. It is usually best to begin by recording words in a frame that ensures that they have the equivalent of the nuclear stress in a sentence.

Finally, there must be material that illustrates the suprasegmental aspects of the language – variations in stress, tone and intonation. (Length distinctions have been mentioned already.) In a tone language, the word list must include words that illustrate the contrasting tones, each on at least two different vowels. If stress is significant, as it is in English, the list must include forms such as *an insult*, *to insult*, *to differ*, *to defer*. Speakers sometimes find such contrasts difficult to make when they are not in meaningful sentences. It may be necessary to use special frames that are not absolutely identical, using comparable sentences such as *Lance insults our Dad* vs. *Len's insults are bad*.

Intonation patterns are often hard to illustrate, as the phonological contrasts in many languages have not been systematized. But it is worth recording sentences illustrating commands, statements and

Once, when I was recording with a colleague in Africa, we had a speaker who wanted to tell us traditional stories. He was a good story teller, with lots of rhetorical flourishes that gave rise to the formulaic responses of the people around ('Say it again, father', 'So be it, so be it'). The first story he told us was a great performance, and my colleague wanted to record him telling some more. But we had very little knowledge of the language and would not be able to make a translation. We were never going to be able to use any of it for any scientific analysis. It was a great performance, but we asked him to stop so we could get on with our work.



different types of questions, and other syntactic devices that are conveyed by intonation changes, as well as differences in focus, such as *I want a red pen not a black one*, as compared with *I want a red pen not a pencil*.

So far I have been suggesting a very structured way of getting data that will illustrate the sounds of a language. The principal data will be word lists for the phonemic contrasts and specific sentences for describing the basic intonational patterns. It is also a good practice to record some more conversational utterances. I usually ask speakers to tell me about something simple, such as what they did yesterday. I try to elicit three or four sentences, not more, and then get the speaker, or someone working with me, to translate each sentence into English. There's not much that can be done with recordings that lack written translations and good phonetic transcriptions (which I try to do the same evening or as soon as I can). From a phonetician's point of view there is no point in making lengthy recordings of folk tales, or songs that people want to sing. Such recordings can seldom be used for an analysis of the major phonetic characteristics of a language, except in a qualitative way. You need sounds that have all been produced in the same way so that their features can be compared. From the point of view of a fieldworker wanting long-term ties to the community, it is worthwhile spending time establishing rapport with speakers. But at times you may have to tactfully remind your consultants that, as we will discuss in the next section, work on their language is work and not play.

When you have a tentative list of words and phrases that you want to record, it is time to start working with one or two speakers of the language. Even if you know almost nothing about the language, and there are no books or articles on it or any closely related language, it is

One of the best lists of words that I was able to record illustrates the 20 clicks of Nama, a Khoisan language spoken in Namibia. I was working with an excellent consultant, Mr. Johannes Boois of the Namibian Literacy Bureau. After a little while he said, 'Oh I see what you are trying to do, you want a set of words, each with one of the Nama clicks at the beginning, and each before the same vowel. Let me think about it.' Next morning he came back with a list of words of just that kind. He added, 'You didn't say whether you wanted them all on the same tone, but I found words that all had high tone vowels.'

advisable to make some sort of list before contacting anyone. Your list will give the speakers some indication of what you are trying to do, and working together you will be able to develop a better list.

You will need to go through your list very carefully with your selected speakers. (We'll see in the next section how to find appropriate language consultants.) Usually a number of problems arise the first time through. Some of the words will turn out to be used only by old people, or women, or speakers of another dialect. If you are lucky your language consultants will grasp what you are trying to do, and suggest alternative words. Get them talking among themselves, and they may even be able to suggest contrasts that you had not anticipated. Pay attention to what they have to say, and try to imitate their pronunciation. When you have learned to say a word yourself in a way that a native speaker will accept, you will have found out a great deal about how it is articulated. Moreover it is very satisfying to your language consultants to hear your attempts at speaking their language. One of the first things you should do when working with someone is to learn how to say 'hello' and 'goodbye'. If you greet people properly, showing respect for their language, they will be much more relaxed and willing to help you find a set of illustrative words. You should also be able to say 'please' and 'thank you' to people in their own language (make your mother proud of how you behave).

There is no special trick involved in imitating the pronunciation of a word. It's not like being able to imitate someone's voice or impersonate a particular character, skills I've never had. Repeating a single word that you have just heard is simply a matter of learning to listen for subtle shades of sounds, and learning to produce them. In my experience most students can produce nearly all the sounds on the IPA chart fairly well after a ten-week course in phonetics, provided that they have an instructor leading them through a set of practical exercises such as those in my textbook *A Course in Phonetics* (Heinle & Heinle, 2001). Becoming fully proficient in producing and hearing some complex sounds may take a little longer. I'm still unable to produce a voiced palatal click in a way that fully satisfies a speaker of !Xóõ. But given a proper basic training in phonetics, with a little practice anyone can achieve a near-native pronunciation of most short words. Phonetic ability is not an esoteric skill possessed by only a few.

When you are checking that you have pronounced a word correctly, be careful how you phrase your question. If you just ask 'Am I saying this correctly?', some language consultants will approve your

mispronunciation simply because they want to make you feel good. Tell them that you want them to be very strict teachers, and that you want to sound exactly like them. You should also be aware that when a consultant doesn't like your pronunciation of a word it may not be for the reason you expect. You may be trying to get the correct pronunciation of the consonant cluster at the beginning of a word and find that you are continually told that you are saying the word wrong. But it may not be the consonant cluster that is wrong; you may have got the wrong tone on the vowel, or something else quite different.

A technique that I find better than simply saying a word and asking 'Am I saying this correctly?', is to produce two slightly different versions, and ask 'Which is better, one or two?' I usually repeat a pair of possibilities several times, holding up one finger when I say the first possibility and two fingers when I say the second. If neither is approved it may well be that I am working on altering the wrong part of the word. When approval has been given to one of the two possibilities, I go on with something else and later return to the same pair, but saying them in reverse order, to see if they still prefer the same one. With good consultants you can even take a word one syllable at a time, checking several possibilities for each syllable.

I cannot emphasize too much the importance of working with speakers and thinking out what it is that has to be illustrated before making formal sets of recordings. The word list will no doubt be revised several times during later work with additional speakers, but almost nothing in the study of the phonetic characteristics of a language is more important than this initial work on a word list. You should be prepared to spend many hours checking everything out. You have to be sure that you record all that you will want to analyze. When you have left the field you will no longer have access to a large group of speakers.

Finally in this section, I should note that the formal recordings that I have suggested making are only the first step in describing the phonetic events that occur in a language. In normal conversation we don't use the precise pronunciations that are used when repeating lists of words. A full description of a language will try to account for all the elided forms and the vagaries of casual speech. Making suitable recordings for this kind of analysis involves catching people when they are bound up with what they are saying and have forgotten that they are being recorded. Such recordings are valuable for many tasks,

such as building speech recognition systems. Make them, and, most importantly, annotate them while you can still remember what they are about. But they are not a good basis for a description of the phonetic structures of a language. You need to know the contrastive sounds that occur before you can describe conversational utterances.

## 1.2 Finding Speakers

Probably the question that I am asked most frequently is: 'How do you find appropriate speakers?' There is no general answer. Local conditions vary greatly. In many parts of the world official permission is required before doing any kind of research, and sometimes this works out to one's advantage. There may be a local official to whom one must report, who may be able to help. Often a local schoolteacher, who may not speak the language you are interested in, nevertheless knows people who do. The local postmaster is another good source of information, as are local clergy and police (though in many societies these authority figures may have been brought in from another region, and not be looked on with favor by the local population).

It may turn out that none of these suggestions is any use when you are trying to work with a very isolated group of people. In these circumstances the only way I have ever been able to achieve anything is to find a linguist or a missionary who has lived in the area and can point out the best consultants to me. They can also help me find an assistant who knows both English and the local trade language that the speakers of the language I am investigating will know. I've done well with the help of schoolboys who translated what I wanted to ask into Swahili, which is the lingua franca of many parts of East Africa. The language consultants I was working with knew Swahili (most speakers of small community languages are multilingual), and could produce the words I wanted in their own language.

Sometimes it is very difficult to work on a particular language. Native American communities may want you to get permission from their tribal council. They may be suspicious of your motives and feel that you may be stealing something from them. The language is sacred for many Native Americans in a way that is hard for outsiders to appreciate. It was given to them by the Gods, and is an essential part of their religion and their identity. In these circumstances one should be very careful not to give offence, and offer to show everything to the

One of our former UCLA linguistics students who is a Navajo tells how she was once giving a talk in a Navajo community. She was showing how words could be put together to create new words (just as *sweet* + *heart* creates a word with an entirely new meaning). While she was explaining this an elder called out: 'Stop this blasphemy! Only the gods can create words.' The Navajo language is holy in a way that is very foreign to most of us.

tribal council before publishing it in any way. Speakers may also feel they are being exploited and think that fieldworkers are making money out of them (as, indeed, we are, in that we are usually being paid to do fieldwork and are receiving the rewards of the publication of the knowledge we acquired from them). However, although a few Native American communities may be hard to work in, others are welcoming. They are very conscious of the fact that their language is spoken only by people who are middle-aged or older. When their children grow up the language will be gone. As a result they are eager to have linguists do whatever they can so that at least a record of the language exists.

There are various points to watch for when selecting speakers. The first is to make sure that the language is really their mother tongue. This may seem too obvious to mention, but there have been cases of well-known linguists publishing authoritative grammars based on the speech of non-native speakers. You want speakers who use the language in their daily life, have all their teeth, and are not hard of hearing. You should also consult to find out who are considered to be the 'best' speakers, particularly if you are working on an endangered language that people no longer speak fluently. Unfortunately you will often be told that the best speaker is a very elderly person, who may be missing their teeth, and who would rather reminisce than concentrate on going through a word list. In literate countries you may be directed to a local scholar whose help you should certainly seek, while emphasizing that you want to learn how ordinary people speak. You'll probably get a shrug of the head, and be told that they don't speak properly nowadays. But be insistent and you'll avoid all the spelling pronunciations and hypercorrections you might otherwise get from someone with a vested interest in the language.

It is worth spending time trying to find somebody who is quick and eager to help. If after a short session you find any problems, move on. When you are working on the sounds, people may not realize that you

are not interested in the precise meanings of each word. Once you have a meaning that is sufficient to identify the word for other people, you do not need to know that the word also means a particular type of tree, or the left arm of a new-born babe. Old people can be hard to control. Given a choice I like to work with high-school students. They quickly understand what I am trying to do, and like to act as teachers, correcting my attempts to repeat what they are saying.

Many speakers are interested in their language, and willing to work for nothing. But it is a good practice to insist that *it is work*, and they should be paid – if they want to, they can always give the money to charity. I find it much easier if there is a regular professional relationship between the speakers and myself. They are then more willing to provide the data I want and not to digress. Speakers are often reassured about accepting payment if you can point out that it is not your money, but money that has been given to you so that you can do this work. On some occasions a gift to the community might be more appropriate – I once bought a cow for a group of Kalahari Bushmen who were trying to start a herd. Many groups have a cultural center that would welcome support.

A good basis for determining the rate of pay is to offer twice the wage that they might otherwise be earning. If they are not in regular employment, then twice the wage that a laborer in the fields might get seems fair. Paying at a much higher rate is unfair in another way. Other linguists working in the area will not bless you if you have inflated the local economy by paying too much. But you should, of course, always offer to provide something to eat and drink; and, alas, tobacco is often welcome. In communities that would like a linguist's help, I am delighted to offer what services I can, such as providing recordings and annotated word lists. But I don't like it when people expect too much. I've known cases when people think they are offering me something that is worth more than ten times the hourly rate for part-time work by university undergraduates. I won't deal with people like that, and just move on.

Another question that I am often asked is: 'How many speakers do I need to record?' Ideally you want about half a dozen speakers of each sex. There may be systematic differences between male and female speech. In Pirahã, a language spoken in the Amazonian rain forest, women always use /h/ where men have /s/. If you can eventually find 12 or even 20 members of each sex, so much the better. When working on an endangered language, this may not be possible. I've

sometimes had to be satisfied with the four or five speakers that I can find. In any case, particularly if you are working on an endangered language, or in a country in which the language is not ordinarily spoken, you should check that each speaker is considered by the others to be a good speaker of the same dialect of the language. Sometimes this requires a certain amount of tact, but you can usually take one of the group aside and ask about the capabilities of the others.

I (and other linguists) used to describe the phonetics of a language on the basis of information from a single speaker. Clearly, a single speaker is no longer enough to satisfy modern standards of phonetic description. Such speakers may not speak in what others consider to be the normal way. They may be more literate, and have their own ideas on how the language 'should' (in their view) be spoken. Groups of ordinary speakers are needed to reflect the true phonetic characteristics of the language.

A final very important point in connection with choosing and working with speakers: before you go into the field consult your institution's Office for the Protection of Research Subjects. In the United States, anyone who is affiliated with any university or other institution that receives federal funds must get permission from their Institutional Review Board (IRB) for any research that involves human subjects. This regulation applies irrespective of where in the world you are working or whether you are also affiliated with an institution in another country. The IRB will want to be assured that your language consultants know all the procedures you are using, and what their rights are. Usually they include being able to withdraw their participation at any time, with no questions asked, and to retain their anonymity, unless they explicitly sign away that right. The IRB will want to know about rates of pay, translation of the agreement that is made with participants, whether minors (such as the high-school students mentioned above) will be involved, how permission for their participation will be obtained, how permission for using photographs of field situations (such as those in this book) will be sought, and many other points.

All good fieldworkers will look after their consultants carefully and not do anything that is harmful in any way. But when out in the field it is sometimes hard to get people to sign a written agreement. Getting signatures, or even a mark on a page, is virtually impossible when dealing with people who have never held a pen or signed a document of any kind in their entire life. When working in the Amazonian rain



forest, for example, I've been with people who have had very little contact with the outside world. They are willing to sit and talk for a while, but are reluctant to do anything more. Similarly, in Africa I've met people who are suspicious of formal arrangements. If you are likely to be in these circumstances, you should ask your IRB whether it would be possible to substitute a recording for a written record. You can then make a recording of the procedures being explained to the consultants in their own language or the local trade language. If your consultants show by their comments that they have understood what is going on, and are willing to participate, your IRB may be willing to consider this as informed consent.

Most standard IRB regulations require that the identity of research participants be kept confidential. You should ask your consultants whether they mind their names being revealed. Virtually everybody I've worked with has always been pleased to let it be known that they helped describe their language. It is something that they are proud of doing. You should be sure that your consultants understand they are giving you permission to identify them. Then, in any publication, you can acknowledge them by name. This is a good practice, not only because one should give thanks where thanks are due, but also because it may help other researchers. When in the field, I've often appreciated being able to find consultants that other linguists have worked with.

### 1.3 Recording Systems

There are many ways of making recordings. A common technique is to use a cassette recorder, but DAT (Digital Audio Tape) and CD recorders are widely available. In addition, systems for recording directly onto a computer offer great advantages. My own preference is to use a portable DAT recorder that is small and inconspicuous. People often get worried when you set up a recording machine or a laptop computer and start checking it out. With a portable DAT recorder you can see that it is working properly before you begin, and walk into a room with it already running in your pocket. Then you simply have to bring out a microphone when it comes to the moment to record. (One could even record without people seeing the microphone, but I always ask permission and let people know they are being recorded. Making surreptitious recordings of anybody



anywhere is not normally approved by my (or, probably, any) university's Institutional Review Board.)

Comparing the different systems for making recordings leads to no easy conclusions. There are four properties that one wants from a good recording system, each of which will be considered in turn: (1) a good frequency response (roughly speaking, the range of pitches that the system can record); (2) a good signal/noise ratio (the range of loudness); (3) reliability and user-friendliness; and (4) the possibility of using the recordings for a long time.

The frequency response of a system is a measure of the extent to which it faithfully records and reproduces each frequency. We want pitches that go in at a certain relative loudness to come out at the same relative loudness. We can represent the frequency response of a system by a graph showing the difference in dB (decibels, the unit of intensity, the acoustic correlate of loudness) between the input and the output. The solid horizontal line in figure 1.1 shows a perfect frequency response.

A perfect frequency response is almost what one gets using a DAT recorder or a direct recording onto a computer or CD with a good microphone and Analog to Digital sound system. The thinner solid line in figure 1.1 shows the response of a professional-quality cassette




Figure 1.1 The frequency response curves of a perfect recording system (horizontal solid line), a professional cassette recorder (solid curve), and an ordinary cassette recorder (dashed line).

recorder that has been kept in good condition by cleaning the recording and playback heads. It will record within  $\pm 2$  dB all the frequencies (roughly the pitches) between 80 Hz (the pitch of a fairly deep bass voice) and 11,000 Hz (almost the highest frequency components in an s). This pitch range is completely satisfactory for nearly all speech projects. The dashed curve is that of a cheaper cassette recorder. It will cut off quite a bit of the bass components in speech, and will also have an uneven response including a severe drop-off in the higher frequencies.

The second factor we have to consider is the background noise in a recording. The amount of noise can be measured by what is called the signal/noise ratio. This is the difference between the signal (the sound you want to record) and the noise (everything else, including both the inherent noise of the system and all the background sounds such as running water, distant TVs, refrigerators, or fans that you want to exclude). The signal/noise ratio is stated in dB. Roughly speaking, when two sounds differ by one dB, there is a just noticeable difference in loudness, and when one sound is twice as loud as another, there is a 5 dB difference. A quiet room often has quite a lot of background noise, making it perhaps 30 dB above the threshold of hearing (the quietest sound that you can just hear). A 40 dB signal/noise ratio will occur if the speech you are trying to record is 80 dB above the threshold of hearing and the background is about 40 dB above that level. CD quality, or sounds recorded directly onto a computer, can have a signal/noise ratio of up to 96 dB, although they seldom get above 84 dB, as there is always noise from the system itself.

Before beginning a recording you should check the signal/noise ratio. You will need to record a short piece first, to see that your recording level is set as high as possible without overloading. Then, when no one is saying anything, if you are recording onto a tape recorder of any kind you should check that the sound level indicator is as low as possible. It is difficult to quantify this on most tape recorders as the recording indicator will not be steady, but it may be hovering around  $-40$  dB. If you are recording directly onto a computer, look at the recording on the screen and run the cursor over it. Most systems will provide a record of the voltage measured in bits. The maximum variation in voltage that you can record will probably be 96 dB, which is equivalent to  $\pm 32,768$  bits.

As I sit in my living room, using a computer, but with the microphone unplugged, the system is recording a baseline variation of  $\pm 1$  bits,

which means that there is 6 dB system noise. With the microphone plugged in but turned off, this increases to  $\pm 2$  bits, making the total machine noise 12 dB. With the microphone on but nobody talking, the baseline variation goes up to  $\pm 8$  bits with an occasional spike up to  $\pm 32$ , giving a background noise level of 24 to 36 dB. This still leaves me a 66 dB signal/noise ratio when I record up to the full 96 dB that the system allows. In practice, because I want to stay well below the overload level, I have a usable range on this computer system of 50–60 dB.

Any tape recorder has a range below that, as there is inevitably considerable noise from the tape itself. A professional tape recorder using high-quality tape may have a signal/noise ratio of about 45–50 dB. (I've often heard claims for more, but seldom found them true.) Many cheaper cassette recorders have only a 30 dB range. This is sufficient if you simply want to record the words that are being used, but it is not sufficient for a good acoustic analysis, nor will you be able to use it for a narrow phonetic transcription.

Whether you are using a tape recorder or a computer, it is important to use a high-quality microphone. Built-in microphones, whether on a tape recorder or a computer, are seldom high-quality – they usually cut both the bass and the treble. A condenser microphone, with its own power supply (a small battery, such as those used in watches or hearing aids), will provide the best frequency range. It should have a covering over it, to protect it from the wind and from direct puffs of air from the speaker's mouth. It should also be directional so that it records sound from the front better than from the rear.

The third factor we should consider in comparing recording systems is their reliability and user-friendliness. Cassette recorders are comparatively easy to maintain, but DAT recorders are more complex. When they go wrong, they will probably need to be returned to the manufacturer. Computer and CD recording systems are fairly rugged. User-friendliness from the phonetician's point of view is best summed up by assessing how easy it is to find particular words and to repeat small parts of a recording over and over again, the major listening tasks for phoneticians describing languages. Cassette recorders are fairly easy to use in this way. DAT recorders are more difficult, although they do have the advantage that each section between stops and starts on the original recording is identified by a number, and the time of each event is displayed in real minutes and seconds, rather than by an arbitrary number. Computers win hands down for providing easy

My first paid position in the field of phonetics was as a lab technician, at Edinburgh University. Much of my job consisted of making phonograph records – 78 rpm recordings on vinyl disks. Not many people have the equipment to play these disks now.

finding and repeating of phrases. Selected portions of a waveform can be played repeatedly at the touch of a key without degrading the recording, and a transcription can be put on the screen, directly above the sound wave. Older computer systems were tiresome to use in that one had to stop recording frequently to save files, but this is no longer true, and one can now record reliably for long periods onto gigabyte hard drives.

Finally we must consider how long the recordings are going to last. There are two aspects to this: firstly, how long a particular type of tape or other recording medium will last, and secondly, how long there will be systems on which they can be played. Old-fashioned reel-to-reel tape recordings are still playable after many years, as long as proper precautions have been taken. The main problems that arise are that the tape becomes brittle, and requires careful handling, and there may be some print-through – the transference of the magnetism on one part of the tape to the piece of the tape immediately below it on the reel. The possibility of print-through can be lessened by winding the tape back and then winding it forward before storing it. In 2000 I found that there was only a slight increase in the level of the noise on some reel-to-reel recordings I had made in 1954 (recordings of cardinal vowels by Daniel Jones, now on the web). My only real problem was finding a machine to play them on.

Cassette recordings are probably about as stable as old-fashioned reel-to-reel recordings. They, too, will become brittle over the years. We should also note that every time a tape is copied some noise is added. This is not the case for DAT tapes or CDs, which can be copied with no loss of sound quality. Nobody knows how long DAT recordings will last, or how long they will be popular enough to be supported. Nor do we know the lifetime of CDs if they are properly kept, but we do know that they can be easily scratched and damaged. Computer systems are changing all the time, and it is a good idea to make backup recordings of different kinds to be safe. My best guess is that recording directly onto a computer and making a backup copy on a CD-ROM

will, for some time, be the best way of making a high-quality recording and preserving it so that it can be played for many years.

I always try to have a whole backup system for making recordings in the field. You never know when someone is going to spill a mug of water on your machine or a goat will try to eat it. In addition to a DAT recorder and a computer, I travel with a video camera (for reasons that will be discussed in the next chapter), which has excellent sound recording capabilities and can be used as a slightly less convenient backup system when you've unfortunately dropped all your other equipment over the side of the boat.

## 1.4 Making a Recording

The main problem in making a technically good recording is the elimination of background noise. This is largely a matter of placing the microphone correctly and finding a good, quiet place. The best place for the microphone is as close to the speaker as possible. The place could be a quiet living room (make sure all the doors, especially that to the kitchen, are shut) or somewhere outdoors away from the clamor of the village, with no waterfalls, rustling trees, pounding waves, chickens, chirping cicadas or other animal noises.

Try to keep the level of the recording constant and as high as possible without overloading. In this way the signal (what your speaker is saying) will have a high level relative to the noise (everything else). If you can, it is good to work as a team, with one person ensuring that the proper words are recorded, and the other keeping the level steady while listening for background noise – the refrigerator that has suddenly switched on, or the noisy children drifting closer. It is hard to have to pay attention to the speaker and at the same time look after the recording. As we will see in the next two chapters, field research usually works best when there is somebody watching and controlling the environment leaving someone else to concentrate on the speaker.

I once spoiled an otherwise excellent recording made in the wet season in Nigeria. I was so used to the noise of the rain on the roof that I just didn't realize how loud it was. When I listened to the recording I found I couldn't distinguish some words because of the noise of the rain.

## 22 Recording the Sounds of a Language

The best way to keep a high signal/noise ratio is to have the speaker as close as possible to the microphone without actually blowing into it. The intensity of a recorded sound depends on the square of the distance between the source of sound and the microphone. Chickens or cars that are 10 m away will not be a problem if the microphone is only 2 cm from the speaker's mouth. Even if they had equal intensities as sources of sounds, at the microphone the speaker would have 250,000 times more intensity than the chickens. You can achieve ratios of this kind and avoid problems due to speakers moving around by using a head-mounted microphone placed just to the side of the lips so as to avoid the direct rush of air in fricatives and stop bursts. If it is also a directional microphone that records sound from the front better than from the rear, the signal/noise ratio will be even better.

There are advantages in recording people one at a time using a head-mounted microphone to control the noise. But there are also advantages to recording people in groups, with one person saying the word or phrase and the rest of the group repeating it. One has to make sure that the leader is a good speaker of the language, so that the others, if they are simply imitating this pronunciation, are at least using a pronunciation that is accepted as being representative of



Recording a group of !Xóõ speakers in the Kalahari Desert. (Photograph by Tony Traill.)

the language. This type of recording takes less time, and has more homogeneity. When I have managed to get a group of people together, I am often unwilling to let them disperse while I record them individually, in case some of them don't come back.

Using a highly directional microphone it is possible to get a good recording by pointing the microphone at the speakers one at a time. This both signals that it is their turn to speak and allows for some adjustment in the intensity of the recording. The microphone can be placed closer to those who speak more softly, and further away from the loud-mouthed types. The disadvantage of this technique is that it involves holding the microphone and moving it – circumstances that are apt to produce unwanted noises. You may be able to place a microphone near the center of a group (preferably so that it is closer to the quieter speakers) and then signal by a hand gesture when it is each person's turn to speak. But in fieldwork situations such control is often not possible.

When recording a group of speakers it is important to make sure that the lead speaker is well respected. It is, for example, inappropriate in most countries to have a young woman as the lead speaker. Whatever your feelings about the status of women, don't try to impose them on a group of young African men, who will only laugh or play the fool when you expect them to take their lead from a young woman. On the other hand, a well-respected older woman is often an excellent choice to lead a group. Sometimes it is easier to get women speakers than men. When working on Montana Salish we had a wonderful group of elderly women, but I had to cajole a number of men into joining us by suggesting that they surely did not want their language to be represented just by women.

In a country in which the speakers are literate, it may be possible to ask speakers to read the list, but this is seldom a good idea. Even educated speakers are apt to read with a different pronunciation from that in their normal speech. You can usually get a more natural pronunciation by giving a prompt in English, or an equivalent in the contact language being used, and then having the speakers respond by saying the required word in their own language. (You should use a technique like this even when working on English.) Other useful elicitation techniques include naming objects in pictures. One linguist I know always travels with books about the birds, mammals and insects in the region. Children's reading books often have useful illustrations. But you should note that elderly speakers in some countries may

not be accustomed to looking at pictures. I remember one old man in Kenya who had no trouble naming the animals represented by crude carvings that are sold to tourists, but could not recognize those same animals in well-drawn pictures.

There are some excellent materials for speech pathologists that can be used equally well for eliciting words from speakers who are not in need of therapy. Most of the web sites I am familiar with were devised for speakers of English, but they can often be used with speakers of other languages. For example, Black Sheep Press, <http://www.blacksheep-epress.com/pages/freebies>, has freely downloadable pictures that can be used for eliciting verbs, adjectives, prepositions, emotions, etc. Another site with many free pictures and stories designed for speech pathologists is: [http://www.speechteach.co.uk/p\\_general/downloads.htm](http://www.speechteach.co.uk/p_general/downloads.htm). You can also search for images at <http://www.google.com/> (click on 'Images' at the top of the page, instead of accepting the default 'Web'), although you should remember when doing this that many of the images you find may be copyright. The web is full of sites that you can use to build up elicitation materials.

Illustrations can often be built into a story. Alternatively, if you want to make sure you get certain words pronounced, a map task is useful. Given a map such as that in figure 1.2, a request to explain how to get to the Holiday Inn on the route shown will produce

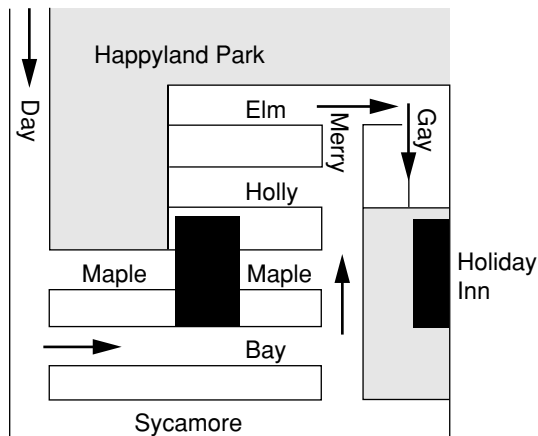


Figure 1.2 An example of a map task for eliciting specific words.



utterances like: 'Go down Day Street, turn left on Bay, left on Merry, right on Elm, and right on Gay.' This includes the words *bay*, *day*, *gay* in roughly comparable contexts (they would be more comparable if *gay* were not paragraph final).

Whatever the technique that is used for making a recording, each session should begin with an announcement of the date and time, the name of the person responsible for the recording, the place and topic of the recording session, an identification of the first speaker, and the speaker giving permission to be recorded (and permission to be identified, if that is the case). Whenever there is a new speaker this person should be identified, permissions recorded, and the time noted. All this information should also be written on the box for the tape, and, in an abbreviated form, on a label on the tape itself. Remember that labels come off, and that tapes get put in the wrong boxes, so make a point of recording all the information at the beginning of the tape. If anything changes during the course of the recording session, make a note of that as you go along, and sum up all the changes in a recorded comment at the end. Say what is happening clearly and distinctly into the microphone, so that the tape itself provides all the information you will need about the session. You can always edit out unwanted comments later.

When a recording session is over, or the tape is full, there are two things that should be done promptly, in addition to completing the documentation. First, secure your recordings, making it impossible to record anything else on top by accident. If you are using a cassette tape recorder, break the tabs on the back edge of the cassette. With a DAT recorder you have to move the slider to the locked position. In the case of a recording made directly onto a computer, make sure that the file is locked. Secondly, make a backup of the whole session, and store it in a safe place, preferably in a different location from the original recording. CDs on which you can write only once are invaluable for this purpose; but it really does not matter in what way you make a backup copy, as long as you do it some way.

## 1.5 Digital Recording

We should follow the previous section on making recordings by a short note on the choices that are available when using a DAT recorder or recording onto a computer. This is not the place for a full discussion

on the nature of digital recording – you can find that in my book *Elements of Acoustic Phonetics*. All we need to note here is that computers, and to some extent DAT recorders, let you choose the range of frequencies you record. The highest frequency (pitch) that you can record depends on the sample rate – the number of points on the sound wave that are recorded as digital values every second. If you record a sound wave by writing its value onto the computer 22,000 times a second (that is, if the sample rate is 22,000 Hz), then you will (at least in theory) be recording frequencies up to 11,000 Hz. If your sample rate is 44,000 Hz (the same as a hi-fi CD), the theoretical limit is 22,000 Hz, which is above the limits of hearing for nearly everybody. As we saw earlier, speech recordings should include all frequencies up to 11,000 Hz (a sample rate of 22,000 Hz) but there is seldom any need to go higher.

My normal practice is to use a sample rate of 22,000 Hz when recording directly onto a computer for making a backup copy. When transferring data onto a computer for acoustic analysis a lower sample rate is often preferable, both because it saves space on the computer and because, curiously enough (see *Elements of Acoustic Phonetics*), it may enable you to make *more* accurate analyses of the most important phonetic aspects of sounds. All the information in vowels is below 5,000 Hz, which can be successfully recorded with a sample rate of 11,000 Hz; furthermore, the best analyses of the pitch of the voice can be made using this rate and eliminating the higher frequencies.

## 1.6 Listening to Recordings

I have already mentioned that it is best to make a rough transcription, perhaps in orthography, as soon as possible. Even if you don't have time to make a full transcription, at least listen to the whole recording, and make notes of what happened at particular times. Nowadays I do all my listening on a computer. There are many different programs that can be used to record and play speech on a computer. Some of them are available free or come with the operating system, but the better ones are fully supported commercial programs. In this book I will be referring to the one I use most often, PCquirer/Macquirer (the same program for different computers).

If the original recording was not made directly onto the computer, begin by copying sections of a minute or so onto your hard disk. You

When Daniel Jones, the greatest phonetician of the first part of the twentieth century, was setting out on a fieldwork trip, a reporter asked him, 'Professor Jones, what instruments are you taking with you?' He pointed to his ears and said, 'Only these.' There is no doubt that the ultimate authority in all phonetic questions is the human ear. But nowadays instrumental aids can often illuminate particular points, acting like a magnifying glass when we need to distinguish between two similar sounds.

can then listen to a word or a phrase at a time over headphones (computer loudspeakers are seldom very high quality), playing it over and over again with a single keystroke. If your computer system will allow you to listen to the recording at half speed without halving the pitch (as it should), then listen to vowels and diphthongs in this way. Consonants are best heard by playing the syllables containing them over and over again, but vowel quality can be fully appreciated when played at half speed without distortion. You should also try playing a selected portion of a sound in reverse. This allows you to focus more easily on offglides, which will have been made to sound like onsets.

When you think that you know how to represent a particular piece of a recording, transcribe it directly on the computer screen above the waveform, as well as adding it to the text, which you should be continually building up. I also try to get a group of consultants to listen to as much of the recording as possible. They can be a great help in checking the data. They are apt to say things like 'Listen to Charlie and how he says [the word for "grass"]. None of the rest of us say it like that.'

While you are checking your recordings, remember the object of the whole enterprise – to make a description of the sounds of a group of speakers. As a result of making a list of items to record you will know a great deal about the language. Now is the time to think about how you are going to write up your results. What are the most interesting points about this language? What are you going to discuss first?

It may seem odd to be thinking that far ahead, and no doubt your original outline will bear only a faint resemblance to what you finally write about the language. But all the time you are working you should remember the ultimate aim. You are not making recordings for your own collection, but planning how to describe the phonetic structures

of a language. To do this effectively, write everything up as you go, getting as much down in well-ordered thoughts as you can, and certainly having a draft of the main points well before finishing the fieldwork.

## 1.7 Making Field Notes

Consideration of writing leads naturally to the final point to be discussed in this chapter. Coming back from the field with a set of well-labeled recordings is not enough. You must also have a written record of all that occurred. Before the days of computers it was easy to say how that should be accomplished. Buy the best notebook you can afford and write down everything that happened in pen. As soon as I was doing funded research, I bought leather-bound notebooks with numbered pages. Everything went in there: the date and time, the names of the speakers and other people present, their ages and relevant details of their linguistic background, and the word lists with glosses and transcriptions, always using the International Phonetic Alphabet, with varying degrees of detail being shown, and extra comments clarifying the precise use of the symbols. If some notes got made on scraps of paper, they were stuck in securely. Photographs of the speakers, photocopies of data from elsewhere, everything was dated and pasted in, so that there was a complete record of all that I was learning and observing. When I went back and made further notes on the data, I used a different color pen, leaving the original entries readable and adding a date for the notes in a different color.

To keep this book honest, I must admit that I tried to follow these procedures, but I was not always as thorough as I should have been.

Shortly after I went to UCLA I was hired as a consultant on the movie *My Fair Lady*. My job was to help Rex Harrison act like a Professor of Phonetics, pointing to the correct symbols and making transcriptions in appropriate notebooks. I explained that I was only a poor Assistant Professor of Phonetics, but if I were rich enough to have a butler and three singing maids, I would have handmade notebooks bound in green leather. Six such notebooks were made to my specification. Rex Harrison took one look at them and decided they were too big to go into his pocket. So I got a set of good notebooks for my own use.

I've just been looking back at my notebooks, which date from 1957 to the present. I see that there are errors – missing dates, full names of speakers not given, and some things not stuck in as well as they should have been. But my intentions were always to keep as full a record as possible.

Even with their failings, my notebooks have proved invaluable. When someone walks into my office and asks if I know anything about a particular language, say Shona, I can recall that I was in Zimbabwe in 1965 (it was called Southern Rhodesia then) and find what I had recorded – amongst other things, the word for a sugar ant, [s<sup>w</sup>os<sup>w</sup>e], with its unusual whistling fricatives that at the time I symbolized [s<sup>w</sup>] with some additional notes and pictures.

With the advent of computers it is not so clear how one should keep written records of a fieldwork investigation. I find it easiest to develop a word list on a computer, using a spreadsheet or a standard word-processing program. I like the ability to have transcriptions in an IPA font in one column, English glosses in a standard font in another, Swahili or whatever inter-language is being used in a third, and additional notes as needed elsewhere. That way I can alphabetize columns as needed, search for particular items, and add extra columns when working with a different speaker. But speakers sometimes find it distracting if you work on a computer while obtaining data. You need to make it clear to your speakers that you are fully focused on what they are saying. Making notes using pen and paper seems less intrusive. My compromise solution is to prepare word lists and notes on a computer, print them out, and then annotate them by hand while working with a speaker. At the end of the day I rename the data file (so that I always retain my older work), type in the new data that I have obtained, and print out the new files. Fortunately small printers are now cheap and can easily be used in the field.


The disadvantage of this method is that I end up with masses of printouts, and other loose pieces of paper, something that I previously tried to avoid. I put everything connected with a particular investigation into a single binder, instead of into one notebook. At the moment I find the data is as easily accessible as in my old notebooks, but I am not sure if this is really the best solution. However I am sure that permanent records should be in the form of printouts, not computer disks or CDs. I see no reason to believe that our current computer media will be readily accessible in 50, or even 20, years' time. But printed matter in English has been around for well

over 500 years, and I expect it to be readable for many years in the future.

## 1.8 Instrumental Phonetic Techniques

For the first half of the last century phoneticians relied simply on their ears to make impressionistic descriptions of languages in the field. During the second half of the century their lives became easier because they could make tape recordings to support their auditory analyses. They could also take these recordings back into the lab for later acoustic analysis. But it was not until the latter part of the last century that it became possible to use a range of instrumental techniques in the field. Early phoneticians did wonderful work relying simply on their ears, but they could not find out what was going on inside a person's mouth, nor measure the amount of air flowing out of the nose in the ways we will discuss in the next two chapters. Nowadays you can take a basic phonetics lab – the kind of instrumentation that should be available to any student describing a language or dialect – anywhere you can carry a backpack. Instrumental phonetics has made it possible to document descriptions of languages more precisely. We can now

One of my tasks when working as an advisor on *My Fair Lady* was to help set up a 1910-style phonetics laboratory. We had an old-fashioned kymograph, a device with a revolving drum covered with white glossy paper that had been smoked so that there was a layer of soot on it. The subject of the experiment, Eliza Doolittle (Audrey Hepburn) spoke into a tube that led to a rubber drum that vibrated when a voiced sound was produced. A straw with a pointer on the end picked up these vibrations and made a trace on the soot-covered paper. It was quite an effective way of recording the duration of a voiced sound, but smoking the paper (and varnishing it afterwards so that it became fixed) was an elaborate procedure. I knew all about this technique as I am old enough to have used it as an undergraduate. I still have the records (one of which is reproduced below) of an undergraduate experiment in which I measured vowel length.



use scientific methodology to say how a previously undocumented contrast in a language is made. We can make hypotheses about the sounds of a language, and test them by making valid, reliable and significant measurements.

Let's consider what each of the terms, valid, reliable and significant, mean. Supposing, for example, that we had a hypothesis that there are two sets of vowels in English, one set exemplified by the vowels in *bait, beat, bite, boat, beaut(y)* and the other by the vowels in *bat, bet, bit, bob, bud*. (I've purposely used the way a schoolteacher might talk about the sounds of the letters 'A, E, I, O, U'.) We might make a hypothesis that the vowels in the first set are longer than those in the second. If we measured the vowel lengths in a recording of someone saying those words we would not have *valid* numbers for testing our hypothesis. These words differ in more than their vowels. They differ in their final consonants and number of syllables. We need an exact set of minimal pairs to make valid measurements that test the hypothesis, otherwise we can't be sure that the other differences between the words are not the reason for the differences in length.

Making *reliable* measurements means making them in accordance with a known procedure that others can follow. As we will see in chapter 4, measuring duration is not as simple a process as you might imagine. There are decisions to be made that have to be made explicit about where a sound begins and ends. Your measurements will not be considered reliable if others cannot follow your procedures and get the same answers.

Finally, for the hypothesis to be verified the measurements have to show that there is a *significant* difference. This requires showing that the differences are not due to chance. If you have measured a sufficient number of speakers to make sure that you are describing properties of the language, and not just the personal characteristics of one or two people, you will be able to report the statistical significance, the likelihood that the event is not due to chance. This book is not the appropriate place for a discussion of statistical techniques for measuring significance, but you should bear in mind that if you want to make claims about a language or a dialect, you should always subject them to statistical tests. Much of the rest of this book will be concerned with making measurements of speech by instrumental means. Try to make sure that these measurements are always valid, reliable, and statistically significant.