Eentaalpsychologie is geen taalpsychologie

(Part II)

AFSCHEIDSREDE DOOR PROF. DR. E. A. CUTLER

Radboud Universiteit Nijmegen

AFSCHEIDSREDE

PROF. DR. E. A. CUTLER

Het eerste deel van deze titel vormde ook de titel van Anne Cutlers inaugurele rede aan de Radboud Universiteit, in 1996. De strekking van het toenmalige betoog was dat de psycholinguistiek vergelijkend te werk moet gaan; alleen door de menselijke verwerking van verschillende talen in kaart te brengen, en de verwerking van de ene taal met de verwerking van de andere taal te vergelijken, is de menselijke verwerking van taal in het algemeen te doorgronden. Inmiddels is zulke taalvergelijking een wijdverbreide psycholinguistische methode. Vergelijking tussen het gebruik van twee talen tegenover die van één taal is in dit vak ook goed vertegenwoordigd. En sinds kort staat ook de vergelijking tussen de taalverwerking van de ene tegenover de andere individuele taalgebruiker op de wetenschappelijke agenda. Hier wordt voor nog een andere soort vergelijking gepleit. Nieuwe onderzoekstechnieken bieden mogelijkheden om bij dezelfde gebruikers de verwerking van de ene (bijvoorbeeld perfect beheerste) taal tegenover de andere (bijvoorbeeld bijna vergeten) taal te vergelijken. Deze afscheidsrede is dus geen einde: eerder een nieuw begin.

Anne Cutler (Melbourne, Australia, 1945) neemt met deze rede afscheid als hoogleraar Vergelijkende psycholinguïstiek aan de Radboud Universiteit Nijmegen. Vanaf 1993 was zij tevens een van de directeuren van het Max Planck Institute for Psycholinguistics in Nijmegen. In 1999 ontving Cutler een Spinozapremie van de Nederlandse organisatie voor Wetenschappelijk Onderzoek (NWO). Met die premie maakte zij de totstandkoming van het baby research center, verbonden aan de Radboud Universiteit, mogelijk.

Radboud Universiteit Nijmegen

EENTAALPSYCHOLOGIE IS GEEN TAALPSYCHOLOGIE: PART II

Eentaalpsychologie is geen taalpsychologie (Part II)

Rede uitgesproken bij het afscheid als hoogleraar Vergelijkende taalpsychologie aan de Faculteit der Sociale Wetenschappen van de Radboud Universiteit Nijmegen op donderdag 20 september 2012

door Prof. dr. E. A. Cutler

Vormgeving en opmaak: Nies en Partners bno, Nijmegen
Fotografie omslag: Bert Beelen
Drukwerk: Van Eck & Oosterink

Mijnheer de rector, dames en heren,

Vreemd genoeg staat voor u een titel deels in het Nederlands terwijl de rede zelf geheel in het Engels wordt gesproken (behalve deze zin dan).

The first part of this title was also the title of my inaugural lecture, delivered in this auditorium 16 years ago this month. At the time I had put a lot of effort into composing and delivering this inaugural address in Dutch, my new second language, and I was particularly proud of the title, simply because it could not be translated into English without the loss of a good deal of its meaning (and especially the loss of any associations with eggs which some of you may involuntarily call up). Translating it word by word ('One-language psychology is no language psychology') produces a superb example of *steenkolenengels* which would be quite opaque to English speakers with no knowledge of Germanic languages. All English titles that I could think of seemed pale in comparison to the original; the best was 'Why psycholinguistics must be comparative'. As a translation this is highly impoverished (it even fails to stress that what we must compare across is languages!). But at least it captures the central argument of the inaugural lecture, that basing psycholinguistic research and theorizing on evidence from only a single language, as has so frequently been done, will often simply lead to a wrong conclusion or to only a partial truth. You can find a version of that argument serving as the introductory chapter in my book *Native Listening*, which was published just this month.[1]

Today I won't repeat that argument, but I will give you some new examples of the importance of cross-language comparison. These examples are taken from the past 19 years, the years during which I have been fortunate enough to hold the position of director at the MPI and a chair at this university. In fact, it's easy to make this case even by looking at the simplest examples, i.e., the building blocks of spoken language.

My research focuses on listening to spoken language, as the title of the book makes clear. I regard listening to speech as an operation that is continuously influenced by one's native language. More than that, listening is exquisitely tailored to the native language, which is the main reason why it is so extraordinarily efficient and so wonderfully flexible and adaptable.

THE BUILDING BLOCKS OF SPEECH

The bottom line for listeners is, of course, that we want to understand what people say to us. What people say we call utterances, but utterances are made up of words. Take an example utterance (e.g., by a minister): *This government has been good for science*. And then take another utterance (perhaps by a working scientist): *This government has not been good for science*. There is a crucial difference between the two, and it consists of just one word. One small word makes all the difference, and we cannot understand utterances without paying close attention to all the words they contain.

Words, in their turn, are made up of sounds. Thus there is a difference between *knot* and *net*[2], *not* and *knit*, *not* and *knock*, *not* and *nod* (not that all Dutch listeners attend to that latter word-final voicing difference, which is not part of their native language experience). Similarly *not* differs from *rot, lot, what* and *got*. All of these differences can be vital for understanding. *What a chance!* versus *Not a chance! Niet gewonnen!* versus *net gewonnen!*

Sometimes, in some circumstances, we may understand utterances as a whole. A predictable commonplace in its expected context may be taken for granted by the listener. Indeed you can teach a dog to respond with appropriate behaviour if you say *fetch the paper!* But never let anyone tell you that this is the whole story of listening to speech, because that is far from the case. The most extraordinary aspect of human listening is how we deal with the unexpected: we seemingly instantly understand new utterances, we laugh at jokes, we groan at puns. This ability is what psycholinguists need to explain.

Listeners know well that individual sounds matter – that *not* and *net* are different words and that they differ just in that minimal way. If just one sound in a word is changed, so that a meaningless non-word results, listeners can tell that the word has been changed, but can still usually work out very rapidly what it was supposed to be. Take for instance the non-words *eltimate* and *teeble*, which differ from real English words in just one sound. Or *kossa* and *podaal*, which are similarly transformed Dutch words; or *pecto* and *cefra* which have their basis in Spanish words. If listeners are asked to do this 'word reconstruction' task as quickly as they can, they usually succeed in coming up with the responses in less than a second: *ultimate, table, kassa, pedaal, pacto,* and *cifra*.

But in fact all of these non-words actually differ minimally from two real words. They could also be *estimate, feeble, komma, modaal, recto* and *cebra*. Speech sounds come in two varieties: vowels and consonants. The first set of words required a vowel to be changed, the second set needed a change in a consonant. Across these three languages, listeners were consistently more likely to find a response that differed from the non-word in a vowel rather than in a consonant. If they were required to alter a vowel, their correct responses were made more rapidly than if they had to alter a consonant; further, when they were required to alter a consonant, they were not only slower, but they were also more likely to respond with the vowel-change option by mistake. In other words, it seemed to be much easier to find a word neighbour by changing a vowel than by changing a consonant. Across languages, these two types of sound behave differently in this task.

Why did we do these experiments in more than one language? The cross-language comparison tells us that the vowel-consonant difference has nothing to do with (a) the relative size or (b) the constitution of the speech sound inventory of a language, because (a) British English has nearly twice as many speech sounds as Spanish overall, and (b) Dutch has almost the same number of vowels and consonants, but Spanish has four times as many consonants as vowels. It also tells us that the vowel-consonant difference in this task has nothing to do with what varies most across dialects of a language, because

that too is ruled out by this comparison: British dialects differ more in vowels than in consonants (Southern British English, *luck, look* and *Luke* all have different vowels, while in Scottish English, *look* and *Luke* are homophones, and in Yorkshire English *look* and *luck* are homophones), but Spanish dialects have salient consonant differences (the medial consonant in *gracias* is [s] in the Canary Islands and parts of Andalusia, but is the dental fricative [θ] in Castilian Spanish spoken in most of central and northern Spain). Without doing the experiment across these different languages we would never have known how to interpret the difference we found. Since it appears in essentially the same way in all of these languages, we know that it must be due to some aspect of the vowel-consonant difference itself, which makes a vowel change easier for listeners to do.[3]

SO WHAT IS THE DIFFERENCE BETWEEN VOWELS AND CONSONANTS?

Mainly it is articulatory: whether articulation is uninterrupted (vowels), or is interrupted by the vocal tract being closed off or constricted (consonants). Stopping or constricting the vocal tract is a speech gesture, with a trajectory that can affect the sounds before and after it. In general, such a gesture affects the vocal tract positions before and after it more than these positions in turn affect the gesture. Vowels and consonants tend to alternate more than they tend to cluster together (indeed, some languages require vowels and consonants to alternate with one another, while no languages require either vowels or consonants to cluster). The alternation plus the difference of articulatory type produces a situation where consonants effectively alter the nature of vowels more than vowels alter the nature of consonants. The findings from word reconstruction show how we as listeners take this into account all the time, every day. We know that vowels tend to be a little unreliable, so we have accumulated a lifetime of experience changing our first impression of what a vowel sound is. That also makes it easier to do in this experiment.

Thus speech sounds are not always equal. In this particular respect, consonants seem to be more solid than vowels – they seem to give us more information about word identity. If this seems familiar, perhaps it's because you have come across something

Table 1: utterances reduced to just their consonants or just their vowels.

(a) GFLCTRD MT J VRJRDAG
(b) E E I I EE E E E AA A
(c) GD SV TH KNG
(d) O A E I
(e) Y
(f) Y T T
(g) N S GN

similar in writing, such as a card inscribed with example (a) in Table 1. Somehow providing only the vowels, as in (b), would be much less informative here. (An English example works just as well: example (c) in the table is interpretable, while (d) isn't.)

But, as I am sure everybody in the audience realises full well, these are carefully chosen examples and it is just as easy to find cases that don't work. Example (e) in the table, for instance, is a very well-known three-word English utterance, but what is it? With four more words added, in example (f), it becomes a seven-word utterance, but the consonants are no more helpful; only with the vowels added can you know that the three words are *I owe you*, and the seven words *I owe you a tea or two!* A Dutch example again works equally well; consider (g) in the table. I hope you found *Een ei is geen ei* immediately! Please don't say that you were not expecting this – after all, I reminded you about eggs right at the outset!

These wordplay games make an important point. Words like *eye, owe,* or *awe* in English, or *ei, ui, aai* in Dutch, have no consonant sounds at all. The reverse doesn't exist in English or Dutch: you cannot have words without vowels in them. This is simply because vowels are more continuant sounds and the consonants are largely the transitions between them, as we said.

THE INDISPENSABILITY OF VOWELS

This really important vowel-consonant difference – words can't be words without vowels, so vowels are indispensable while consonants are not – explains a lot about the efficiency of recognizing words in speech. Here is another experiment (using a simple task called word spotting) that shows how users take account of this dispensability difference. If I ask you to listen for any real word somewhere in a string of nonsense, for instance in:

obzel crinthish bookving fegg ooble

then the evidence (from experiment after experiment) suggests that you will find *book* with ease, but it's rather unlikely that you will find *egg*, and certainly unlikely that you will find it as quickly as you found *book*. This isn't due to their relative position in the word, because suppose that I had said

obzel crinthish bookving maffegg ooble

then the same experiments tell us that you would be much more likely to have found *egg*, and indeed to have found it quite rapidly. The same holds true in Dutch: *wonen* is often missed, or is spotted only slowly, in *dwonen*, but it is spotted faster and more accurately in *lewonen*.

What listeners are doing here – to produce this difference in speed and accuracy of detection – is making use of the vowel/consonant difference of dispensability to deal

with a major problem that we face every time we listen to speech: the problem of getting rid of words that are accidentally present. No language has enough speech sounds to make words really different from one another (languages in general have a few dozen speech sounds, but over a hundred thousand distinguishable words); thus when we listen to speech we are always going to hear lots of words that are present accidentally, because they occur inside other words.

We can call this *the egg in the leg* problem; every time we hear the word *leg* we also hear the word *egg*; how do we know to recognize *leg* and not *egg*? If it seems obvious that *leg* starts first, then we can also call it the *bee on the beach* problem, because it's the same there too; every time you say *beach* you say *bee* first, and how do listeners know not to recognize the first word first? Even more interestingly: How come we don't even notice the presence of embedded words like *egg* or *bee* or the rest?

The reason is exactly that we can ignore many accidentally embedded words, including *egg* in *leg* and *bee* in *beach*, because what they leave behind is a residue that could not itself be a word. It's no accident that *l* and *ch* are not words in English (or Dutch); crucially, *l* and *ch* alone could not possibly be words in English or Dutch. They are just consonants and consonants alone cannot be words. Vowels on their own can indeed be words (*eye; aai*), but consonants cannot.

Acting on this firm principle, listeners rightly conclude that the *l* must belong with the *egg*, and the *ch* with the *bee*, to make longer words. This is why it's hard to detect *egg* in *fegg*, even though *fegg* isn't a word and it's easy to find *egg* in *maffegg*, even though neither *maffegg* nor *maff* is an English word – *maff* might have been a word (given that *mat* and *muff* and *gaff* are all existing English words). But *f* simply could not be a word, so it is treated as if it must be part of a longer word.

Of course, we don't do this consciously at all, and the easiest way to prove this is to show that even babies do it! It is possible to show that very young babies easily find individual words in continuous speech. (They need to be able to do this or they would never start to learn words, because nearly all of the speech we address to babies is continuous; we don't start to teach individual words to babies until they are old enough to try to say them back, at around one year old. It's actually between six and nine months of age that babies start working on the problem of how to find words in speech; but nobody leans over a six to nine-month-old baby in a pram saying *dimple! dimple! dimple!* We're much more likely to say *what a cute little dimple you've got!*)

In a typical continuous-speech experiment, babies sit on their mother's lap and hear a word spoken several times in the way we typically talk to babies, with exaggerated intonation and higher than usual pitch. Then they hear some sentences which may contain the word they had just heard. If the word they heard was *dimple*, they might then hear: *She worried about the dimple in her chin. That dimple is inherited. Has her friend got a dimple too? Not everyone has a dimple.* Or, the sentences might be: *People talked about the hammock in the yard. The hammock was her grandfather's. Many campsites have*

a hammock. Swinging in a hammock is nice. Typically, the babies who had heard *dimple* beforehand listen longer to *dimple* sentences before they show signs of getting bored, while other babies who heard *hammock* first tend to listen longer to *hammock* sentences than to *dimple* sentences. This shows that the babies can spot the words they had just heard when these words occur in sentences. (The mother, by way, listens to music over headphones while all this is going on, so she has no idea of what the baby is hearing.)

In our experiment on vowel indispensability (in English; this experiment was run in Baltimore) we first let babies hear a word such as *rest* or *win* spoken several times. Then they heard sentences containing not *rest*, but *caressed* or *suppressed*, and not *win*, but *window* or *winsome*. They responded just as if *rest* or *win* had been there, listening longer to these sets of sentences than to sentences that did not contain the words they had heard. But when we did the same thing but used sentences with *pressed, dressed* or *wind* and *wince* (instead of sentences with *caressed* or *window* etc.), then we got a different result; the babies behaved as if none of the sentences contained the words they had heard. That is, they did find *rest* in *caressed* and *win* in *window*, but they did not find *rest* in *pressed* and *win* in *wind*. This is exactly the same as the adults in the word-spotting experiments who find *egg* in *maffegg* but not in *fegg*, or *bee* in *beeshub* but not in *beesh*. For the babies just as for the adults, an acceptable word must not strand a residue consisting of consonants only; *egg* is not noticed if it leaves just *f* behind, *rest* is not noticed if it leaves just *p* behind. The indispensable presence of a vowel for anything to count as a word is thus a really powerful effect of the vowel-consonant difference. Even as a baby, before you have had the chance to construct a stock of words, you can realise, just from hearing what speech is like, that some types of speech sound are more essential than others, and that anything without that indispensable vocalic component is probably not a word.[4]

THE INDISPENSABILITY OF CROSS-LANGUAGE COMPARISON

The word-spotting experiments produced the same result across many languages. As in the original experiment in English (*egg* is hard to find in *fegg*, etc.), so it was in Dutch (*wonen* in *dwonen*), French (*lac* in *flac*), or German (*Rose* in *trose*). And babies' use of much the same useful rule of thumb suggests that this way of distinguishing really present words from accidentally present words might be universal across languages.

But cross-language comparison is still a good idea, because there are languages that raise questions about the rule. For instance there are languages which only allow words with two or more syllables, so that perhaps more than a single vowel would have to remain for a word to be acceptable in such languages (since in such languages, *maff* could not be a word any more than *f* could). We went to Africa, to Lesotho, to test that out because the Bantu languages (including Sesotho, the language of Lesotho) are like that. We found that the effect worked in just the same way in Sesotho as in the European

languages where we had first seen it. Also in languages where many vowels effectively disappear in most speech, e.g., in Portuguese or Japanese, the effect was the same: if a vowel was left over, words were easy to spot, whereas if there was no vowel, they were hard to detect.

It gets still more complicated, though: there are, among the weird and wonderful linguistic resources of our world, some languages that do actually allow words to escape having vowels. For instance, Czech and Slovak have a few prepositions like *in* or *to* that are just consonants, and French has a form of the definite article that is just *l'*. But again, we found that the basic vowel-consonant rule worked in French and Slovak too. By this time just this one single line of research had taken us round a good part of the world, producing a map with almost as many countries represented as on the fieldwork maps of our colleagues in MPI's Language and Cognition Department! But the end has not yet been reached: consider this sentence (the *S* represents the sound beginning *she* and the *x* the sound at the end of *Bach*):

tsskSftstt tftxtstt

The utterance means 'you dried it and rolled it', in Berber. There are not a lot of vowels in that utterance! In fact there are none. Berber languages, along with a tiny handful of other languages (e.g., in the northwest of North America and in the southern Caucasus) allow any words – nouns and verbs as well as function words – to consist of consonants only. So we have to ask: how could the vowel indispensability rule be at all useful in such a language? It would be counter-productive, as it would rule out words that actually exist.

So we went to northern Africa and ran the experiment in (two varieties of) Berber in Morocco; and here we got our first negative result. The rule was not used in Berber. A word was spotted as easily when it left just a consonant behind (*fad*, meaning 'bread' in *ghfad*) as when it left a whole syllable with a vowel in it.

Babies are not born with language, they acquire it. So it is experience with their language that has taught speakers of Berber: this rule would not be helpful. What we cannot as yet know is what Berber babies do. Since babies acquiring English apparently develop this rule simply by being exposed to speech, perhaps Berber babies try it out too, but abandon it once they have acquired enough vocabulary to know that it is not helpful for their language. Or perhaps they never develop it at all. We cannot know until we can repeat the baby experiment with infant Berber learners.

However, this single line of comparative psycholinguistic research has shown how not even something as simple as the difference between vowels and consonants can be understood without comparing evidence across languages.[5]

There have been many such comparative research lines in the past 20 years of my group's efforts in Nijmegen (besides English, Dutch, German, French, Japanese, Slovak, Portuguese, Sesotho and Berber, there has been work on Cantonese and Korean and

Arabic and Finnish and Telugu…). Comparison across languages is the only way to understand how humans listen to and understand spoken utterances in any language.

PART II

Now the 'Part II' in the title becomes important. There is more than one way to do comparative psycholinguistics.

Experimental psychology, including psycholinguistics, always analyses behaviour by comparing across participants to discover what is typical of the group as a whole. The type of research described here additionally compares groups of listeners, where the groups differ in their native language. Another established approach is to compare participants who are bilingual with those who are not. This is a well-developed line of research in Nijmegen, and it's very close to second language research, in which we look for instance at whether your native language interferes when you are listening to your second language. Then, there is a new field opening up that is again sure to deliver great Nijmegen achievements: comparing across individual listeners, to discover, perhaps, the genetic basis of why some people are particularly good at language skills. That makes four different ways to look at the psycholinguistics of listening to speech, three of which can properly be called comparative psycholinguistics.

I believe there is also a fourth type of comparative psycholinguistics. It's this: Comparing across different languages within individuals.

This differs from the way in which bilingualism or second language research is typically conducted, in that what is primarily at issue is not which languages are involved or whether the listener is a native speaker of the language(s) or a recent learner of the language. Instead, the focus is on the comparison: what are the dimensions along which listening in different languages varies? Can we map the gradients of listening ability? Can we, ultimately, arrive at an understanding of what constitutes native performance?

This new approach is timely because, first, we have new evidence pointing to tractable issues that can be addressed concerning such inter-language differences, and second, we also have new techniques that offer a way to address these issues. (Both the new evidence and the new techniques have been developed here in my group!)

The new evidence concerns some very well-known ways in which the first language has an advantage. One such area involves identifying talkers' voices; for a very long time it has been known that this is easier in the listeners' native language. It might seem that this is due to being able to understand what is said; if you can understand what a voice is saying, it is easier to identify who the voice belongs to. But, although this implication may indeed be true, it's not the real reason why voice discrimination and identification is easier in the native language. We know that it isn't at all necessary to understand the content of what is said in order to have this native-language advantage and we know this because we found out that seven-month-old babies already have it! How do we test whether babies can tell the difference between individual talkers? We

do this by presenting a set of stimuli for a while. Inevitably, after a while a baby listener gets bored; at that point, we change the input in some way, and then we can ascertain whether the change is noticed (if the baby stays bored, the change has not been noticed, but if the baby perks up and shows interest again, the difference has been noted).

In the voice experiment, the input was just ordinary sentences spoken in an ordinary conversational manner (i.e., not in a specially baby-directed way) by three different speakers. These were absolutely not utterances that a seven-month-old would make anything of (an example Dutch sentence: *Een gevoel van enorme opluchting maakte zich van hem meester*). The babies heard either three Dutch speakers saying such sentences, or three Japanese speakers saying similar sentences in Japanese, or three Italian speakers saying such sentences in Italian. All the speakers were young women with similar voices. Once the baby listener got bored, a fourth speaker was added to the mix. This new speaker was either speaking another language, or the same language.

Unsurprisingly, every baby noticed when a new language appeared – that always caused them to show renewed interest. Japanese and Italian sound quite different from Dutch, and babies notice such inter-language differences from a very early age. But the test of voice discrimination was whether the babies also perked up when the language stayed the same and only the talkers changed – did they notice that a fourth speaker had joined the previous three? The answer is that they only noticed this when the language being spoken was the language they were exposed to in their everyday environment (see Figure 1).

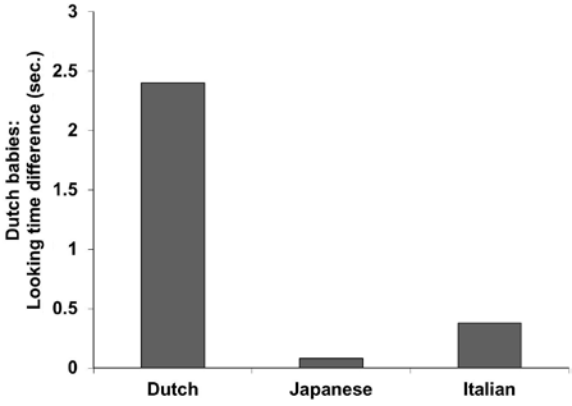

Figure 1 Seven-month-old babies recognise when a new speaker is added to a set of (previously) three unknown female speakers, but only when the language being spoken is the babies' environmental language. The figure shows the difference in attention paid by the babies (as measured by looking time to an accompanying visual display) before versus after the switch from the set of three to the set of four speakers, for Dutch babies listening to Dutch versus to Japanese or Italian (data from Johnson et al., 2011). The difference is significant only when the language being spoken is Dutch.

These 7-month-olds are not yet able to understand anything of what they hear in this experiment. But what they can recognise is the familiar sound of the language that is spoken around them every day. The sounds of the other languages are not familiar. So, it is not understanding that causes this first language advantage in talker identification; it's familiarity with how the speech sounds. It is, in other words, knowledge of the language phonology. That is the basis of the first language advantage in voice discrimination.[6]

The second long-known effect is listening in noise. No matter how proficient we become in a second language, there are still ways in which the native language has an advantage, and one of them is obvious every time we are having a conversation in a noisy cafe or party, or trying to understand an announcement at a railway station or airport: noise makes listening harder in any language, but we suffer more in a second language than in our native language. This is something we all, as second-language users,

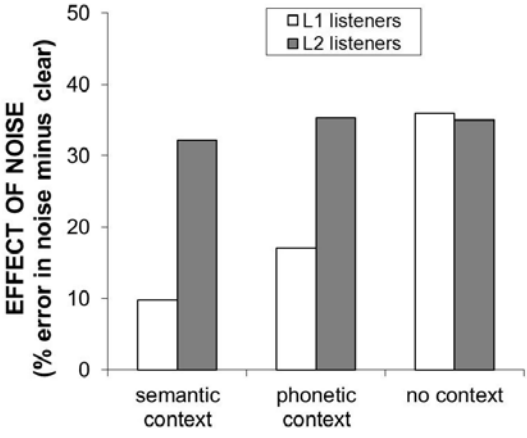

Figure 2 Listening in noise is, under normal circumstances, significantly harder in a second language (L2) than in the native language (L1). This figure shows the difference in the percentage of listening errors when there is noise (multi-speaker babble; 16dB signal-to-noise-ratio) versus errors made under clear listening conditions, and it can be seen that this difference is much greater for L2 listeners (the left column represents the normal case, with ordinary meaningful utterances). When there is no context at all (meaningless syllables differing in length are centrally embedded in a second of noise), however, L1 and L2 listeners are equally affected (right column). Even a constant phonetic context (e.g., consonants vary but are always surrounded by the same vowels: *aba, ana…*; middle column) helps L1 listeners more than L2. Noise is always detrimental, but it is easier to recover from this in the native language by drawing on phonological (and other) knowledge.

are very familiar with (interestingly, this has not stopped scientists demonstrating it repeatedly in the laboratory!) The left-hand column in Figure 2 shows a typical result from such an experiment: the columns show the effect of noise (i.e., how high is the proportion of listening errors in a noisy compared with a quiet environment). The difference is much greater for second-language (L2) listeners than for native (L1) listeners. This is what we usually experience and it is the result we would expect with normal listening to meaningful sentences.

But if we take away all contextual support of any kind and make the task just identification of individual speech sounds, then we get a different result. This is shown in the right-hand column in the figure; the task here was to identify either the first or the second sound in a meaningless syllable such as *za* or *ol* or *vu*, and in this case, L1 and L2 listeners suffered equally from the noise interference. Crucially, if we make just the phonetic context predictable (for instance, by asking for identification of consonants in a constant vowel context: *aza, ala, ava*), then much of the native-language advantage is again apparent (as can be seen in the middle column in the figure). In other words, the native-language advantage is not (solely) due to being able to understand the speech better; it is to a considerable degree due to familiarity with the phonological structure and what speech sounds are like when they are next to other speech sounds. Again we see that listening in the native language has an advantage at the level of phonological processing.[7]

These new insights could be used, for example, to explore the comparative size of the native language advantage across the various languages a person knows. We might in one case compare talker identification as a function of the degree of phonological similarity between the known languages; or we can rank the languages a person knows and examine the extent of recovery from noise-masking as a function of phonological similarity to the native language.

And, as noted above, we also have powerful techniques for carrying out this new endeavour, in both cases: talker identification and adaptability to noisy conditions. The flexibility of native listening in each case has been extensively studied by my group.

One of our major discoveries, in fact, has been the process of perceptual learning by which we listeners adapt so rapidly and easily to new talkers whom we have never heard before. The process is simple: if you hear something ambiguous – for instance, a sound which might have been either an *s* or an *f* – you use whatever information you have in the context to tell you what it is. That is most likely to be the word in which the sound is heard; if the sound occurs as the end of *gira-* it is presumably *f,* simply because *giraffe* is an existing word, while *girasse* is not. This disambiguation allows you to alter your knowledge about that particular speech sound for that particular speaker, so that if the same person later says *ni-* with the same odd sound, you immediately take it for *knife* if you heard *giraffe* before, but for *nice* if you heard *horse* before. We have explored this type of learning quite exhaustively[8] and it will certainly lend itself as a useful

research tool for the new comparative direction. How rapidly do we adapt to a new talker in our native language, our second language, our third language…?

Likewise, the adaptability of listening in noisy conditions has been explored. Normally we listeners are extremely quick in making decisions that we heard one word and not another similar word. If we hear *Not a chance!* we do not usually continue to consider the possibility that it might have been *What a chance!* However, we now know that just a little bit of noise can cause us to do exactly that – to consider other possibilities for longer. In effect, we reduce the speed and confidence with which we make use of incoming acoustic information and exercise a certain degree of caution in our decisions about speech sound identity.[9] In the native language, this adjustment to noisy or otherwise unreliable input proceeds automatically. But how fast and how automatic is it in our second language? In our third? Again, we have identified a clear path to follow in this new type of cross-language comparison within individuals.

This is thus my view of part of the future of comparative psycholinguistics. If it sounds rather programmatic – as befits an inaugural lecture more than a valedictory one – that's because it is; one of the reasons that this lecture is in English is that it will be doing double duty by also serving as my inaugural lecture at the University of Western Sydney in a few months' time!

TENSLOTTE

My world is about to change and, after 36 years in Europe, it's going to take some getting used to. One of the principal ways in which I feel myself anchored in Europe is that all my PhD students so far have graduated here in Europe. There are 43 in all, nine who worked with me in the UK before I came to the Netherlands, 26 who have completed their PhD dissertation at this university, another seven who are just finishing or hope to finish in the next couple of years (and the first one in Sydney is coming along too!).

So I have an enormous amount to be grateful for, here in Europe and here in this country and here at this university.

Without doubt, the most magnificent thing this university did for me was nominate me – successfully – for the Spinoza prize in 1999. The resulting project was called *Native and non-native listening* (yes, a lot of the book is an account of what we found out). A great deal of the cross-language work that I have talked about today was supported by this project, and a great deal more that I didn't have time to talk about – much of the work in Japan, all the work in India and Morocco, much of the work in Korea, and of course a huge amount in Europe, too.

Besides all that research on listening by adults, we also made a major effort to elucidate the origins of native listening – how is it that a baby is born with the ability to become a listener to any language, but then turns into a native listener of Dutch, or English, or Berber…? To do this we set up two baby labs at this university, one for work of the kind you have heard about today on vowels and consonants, and on voice

recognition, and another for studying babies' brain responses while listening to speech. Our baby labs were soon imitated all over this country, as a result of which the Netherlands has become a leader in this type of research; I am enormously grateful for the opportunity to initiate this process!

It is clear that it's because of the Max Planck Institute – and especially because of Pim Levelt who set it up here – that I started on the 19-year voyage that is now coming to an end. Be in no doubt: I have a great deal to thank Pim and all my Max Planck colleagues and the Max Planck Society for. But today is not the MPI farewell and thank-you day; that day is happening next week. And this is not a special thank-you day for my husband Bill Sloman either (every day is a day for thanking Bill).

Today is something else; the principal reason for undertaking today's lecture was in fact to create an occasion for expressing my gratitude to all those outside the Max Planck world to whom I also owe so much after nearly two decades in this country.

Dat wil zeggen, ik wil vandaag vooral mijn hele Nederlandse gemeenschap bedanken.

In eerste instantie gaat het om deze vooraanstaande en vooruitstrevende universiteit. Ik dank het college van bestuur en het decanaat van de Faculteit der Sociale Wetenschappen waaraan mijn leerstoel verbonden is, en de Stichting Nijmeegs Universiteitsfonds die tot en met 2008 dezelfde leerstoel onder haar hoede heeft genomen. Deze universiteit heeft zich tot een wereldleider gemaakt op het gebied van de cognitieve wetenschappen en de cognitieve neurowetenschappen, en heeft verder zich voorbeeldig bewezen in de bevordering van vrouwen. Ik heb me altijd uiterst prettig gevoeld in deze universitaire gemeenschap. En uiteraard blijf ik voor altijd dankbaar voor de Spinoza-nominatie.

Dat gevoel geldt ook voor de gemeenschap die gevormd is door mijn Nederlandse collega's in het algemeen. De psycholinguïstiek verbindt een alfawetenschap, de taalkunde, en een gammawetenschap, de psychologie. In beide vakken staat Nederland wereldwijd op de eerste rang. Er kan zeker gesproken worden van zwaartekracht op dit vakgebied. Geen wonder dan dat ik al voordat ik in Nederland kwam werken een uitgebreid netwerk Nederlandse collega's had. Inmiddels zijn er natuurlijk veel meer bijgekomen. Ik heb me in niet mindere mate zeer thuis gevoeld in de Nederlandse wetenschappelijke gemeenschap.

Ik zei al twee keer dat ik deze universiteit ontzettend dankbaar ben dat zij mij genomineerd heeft voor de Spinoza-prijs, en drie keer is menens! Maar ook naar NWO gaat mijn dank uit voor deze prijs.

En als we het over geldschieters hebben, kom ik naar het cadeau van vandaag. Niets tegen alle soorten cadeaus eigenlijk, maar als je hele inboedel al weg is naar Australië en je zelf over een week of wat in het vliegtuig stapt, dan heb je aan zeg maar een leuke fles wijn helaas weinig. Van dit vreselijke dilemma hebben de collega's van het Landelijk Netwerk Vrouwelijke Hoogleraren mij op fantastische wijze verlost, door een

fonds in het leven te roepen dat de Distinguished Women Scientists Fund gaat heten. Vanaf vandaag is op de website van LNVH te lezen:

'In 2012 heeft het Landelijk Netwerk Vrouwelijke Hoogleraren het Distinguished Women Scientists Fund opgericht. Het bestuur van het LNVH heeft dit in dank aanvaard en besloten het ontvangen bedrag te verdubbelen en in de vorm van een reisbeurs ter beschikking te stellen aan vrouwelijke in Nederland werkende postdocs…'

Dit vind ik het summum, echt het cadeau aller cadeaus, hiervoor ben ik oneindig dankbaar. Nooit zo'n mooi cadeau gekregen!

Universiteit, NWO, LNVH, en alle vrienden en collega's van heel Nederland: van harte bedankt!

Ik heb gezegd.

NOTES

1 Cutler, 2012. You could have found that in the bibliography without assistance! But later footnotes serve to prevent the text being cluttered with references.

2 These examples are of course carefully chosen to underline the point that spelling has nothing to do with any part of this argument!!! This is all about speech sounds; the letters that represent them are not (definitely knot!) relevant.

3 The cross-language word reconstruction study: Cutler, Sebastian, Soler & Van Ooijen, 2000.

4 The first reports of the egg-in-the-leg phenomenon: Norris, McQueen, Cutler & Butterfield, 1997 (English), McQueen & Cutler, 1998 (Dutch). With babies: Johnson, Jusczyk, Cutler & Norris, 2003.

5 Some eggs-in-the-legs across languages: Cutler, Demuth & McQueen, 2002 (Sesotho), McQueen, Otake & Cutler, 2001, and Cutler, Otake & McQueen, 2009 (Japanese), Hanulíková, McQueen & Mitterer, 2010 (Slovak), Hanulíková, Mitterer & McQueen, 2011 (German), El Aissati, McQueen & Cutler, 2012 (Berber).

6 Voice discrimination at seven months: Johnson, Westrek, Nazzi & Cutler, 2011.

7 Listening in noise: Cutler, Weber, Smits & Cooper, 2004, Cutler, Garcia Lecumberri & Cooke, 2008, Garcia Lecumberri, Cooke & Cutler, 2010.

8 The perceptual learning saga: Norris, McQueen & Cutler, 2003; Eisner & McQueen, 2005, 2006; McQueen, Cutler & Norris, 2006; Norris, Butterfield, McQueen & Cutler, 2006; Cutler, McQueen, Butterfield & Norris, 2008; Escudero, Hayes-Harb & Mitterer, 2008; Sjerps & McQueen, 2010; Cutler, Eisner, McQueen & Norris, 2010; Mitterer, Chen & Zhou, 2011; Scharenborg, Mitterer & McQueen, 2011; Scharenborg, Janse & Weber, 2012; McQueen, Tyler & Cutler, 2012; Reinisch, Weber & Mitterer, 2012.

9 Lexical flexibility: Brouwer, Mitterer & Huettig, 2012; McQueen & Huettig, 2012.

REFERENTIES

- Brouwer, S., Mitterer, H., & Huettig, F. (2012). 'Speech reductions change the dynamics of competition during spoken word recognition'. *Language and Cognitive Processes, 27,* 539-571.
- Cutler, A. (2012). *Native listening: Language experience and the recognition of spoken words.* Cambridge, MA: MIT Press.
- Cutler, A., Demuth, K., & McQueen, J. M. (2002). 'Universality versus language-specificity in listening to running speech'. *Psychological Science, 13,* 258–262.
- Cutler, A., Eisner, F., McQueen, J. M., & Norris, D. (2010). 'How abstract phonemic categories are necessary for coping with speaker-related variation', in C. Fougeron, B. Kuhnert, M. D'Imperio & N. Vallée (Eds.), *Laboratory phonology 10* (pp. 91–111). Berlin: Mouton de Gruyter.
- Cutler, A., Garcia Lecumberri, M. L., & Cooke, M. (2008). 'Consonant identification in noise by native and non-native listeners: Effects of local context'. *Journal of the Acoustical Society of America, 124,* 1264-1268.
- Cutler, A., McQueen, J. M., Butterfield, S., & Norris, D. (2008). 'Prelexically-driven retuning of phoneme boundaries', in J. Fletcher, D. Loakes, R. Goecke, D. Burnham & M. Wagner (Eds.), *Proceedings of Interspeech 2008,* Brisbane (p. 2056). Adelaide, Australia: Causal Productions.
- Cutler, A., Otake, T., & McQueen, J. M. (2009). 'Vowel devoicing and the perception of spoken Japanese words'. *Journal of the Acoustical Society of America, 125,* 1693–1703.
- Cutler, A., Sebastian-Galles, N., Soler Vilageliu, O., & van Ooijen, B. (2000). 'Constraints of vowels and consonants on lexical selection: Cross-linguistic comparisons'. *Memory & Cognition, 28,* 746–755.
- Cutler, A., Weber, A., Smits, R., & Cooper, N. (2004). 'Patterns of English phoneme confusions by native and non-native listeners'. *Journal of the Acoustical Society of America, 116,* 3668–367.
- Eisner, F., & McQueen, J. M. (2005). 'The specificity of perceptual learning in speech processing'. *Perception & Psychophysics, 67,* 224–238.
- Eisner, F., & McQueen, J. M. (2006). 'Perceptual learning in speech: Stability over time'. *Journal of the Acoustical Society of America, 119,* 1950–1953.
- El Aissati, A., McQueen, J. M., & Cutler, A. (2012). 'Finding words in a language that allows words without vowels'. *Cognition, 124,* 79-84.
- Escudero, P., Hayes-Harb, R., & Mitterer, H. (2008). 'Novel second-language words and asymmetric lexical access'. *Journal of Phonetics, 36,* 345–360.
- Garcia Lecumberri, M. L., Cooke, M., & Cutler, A. (2010). 'Non-native speech perception in adverse conditions: A review'. *Speech Communication, 52,* 864–886.
- Hanulíková, A., McQueen, J. M., & Mitterer, H. (2010). 'Possible words and fixed stress in the segmentation of Slovak speech'. *Quarterly Journal of Experimental Psychology, 63,* 555–579.
- Hanulíková, A., Mitterer, H., & McQueen, J. M. (2011). 'Effects of first and second language on segmentation of non-native speech'. *Bilingualism: Language and Cognition, 14,* 506-521.
- Johnson, E. K., Jusczyk, P. W., Cutler, A., & Norris, D. (2003). 'Lexical viability constraints on speech segmentation by infants'. *Cognitive Psychology, 46,* 65–97.
- Johnson, E. K., Westrek, W., Nazzi, T., & Cutler, A. (2011). 'Infant ability to tell voices apart rests on language experience'. *Developmental Science, 14,* 1002–1011.

- McQueen, J. M., & Cutler, A. (1998). 'Spotting (different types of) words in (different types of) context', in R. A. Mannell & J. Robert-Ribes (Eds.), *Proceedings of the 5th International Conference on Spoken Language Processing, Sydney* (Vol. 6, pp. 2791–2794). Canberra: Australian Speech Science and Technology Association (ASSTA).
- McQueen, J. M., Cutler, A., & Norris, D. (2006). 'Phonological abstraction in the mental lexicon'. *Cognitive Science, 30,* 1113–1126.
- McQueen, J. M., & Huettig, F. (2012). 'Changing only the probability that spoken words will be distorted changes how they are recognized'. *Journal of the Acoustical Society of America, 131,* 509-517.
- McQueen, J. M., Otake, T., & Cutler, A. (2001). 'Rhythmic cues and possible-word constraints in Japanese speech segmentation'. *Journal of Memory and Language, 45,* 103–132.
- McQueen, J. M., Tyler, M., & Cutler, A. (2012). 'Lexical retuning of children's speech perception: Evidence for knowledge about words' component sounds'. *Language Learning and Development, 8,* 317-339.
- Mitterer, H., Chen, Y., & Zhou, X. (2011). 'Phonological abstraction in processing lexical-tone variation: Evidence from a learning paradigm'. *Cognitive Science, 35,* 184–197.
- Norris, D., Butterfield, S., McQueen, J. M., & Cutler, A. (2006). 'Lexically guided retuning of letter perception'. *Quarterly Journal of Experimental Psychology, 59,* 1505–1515.
- Norris, D., McQueen, J. M., & Cutler, A. (2003). 'Perceptual learning in speech'. *Cognitive Psychology, 47,* 204–238.
- Norris, D., McQueen, J. M., Cutler, A., & Butterfield, S. (1997). 'The possible-word constraint in the segmentation of continuous speech'. *Cognitive Psychology, 34,* 191–243.
- Reinisch, E., Weber, A., & Mitterer, H. (2012). 'Listeners retune phoneme categories across languages'. *Journal of Experimental Psychology: Human Perception and Performance.* Advance online publication. doi:10.1037/a0027979.
- Scharenborg, O., Janse, E., & Weber, A. (2012). 'Perceptual learning of /f/-/s/ by older listeners', in *Proceedings of Interspeech 2012,* Portland OR.
- Scharenborg, O., Mitterer, H., & McQueen, J. M. (2011). 'Perceptual learning of liquids', in *Proceedings of Interspeech 2011,* Florence, Italy (pp. 149-152).
- Sjerps, M. J., & McQueen, J. M. (2010). 'The bounds on flexibility in speech perception'. *Journal of Experimental Psychology: Human Perception and Performance, 36,* 195–211.