Perception & Psychophysics
1998, 60 (6), 1022-1031

Effects of phoneme repertoire on phoneme decision

ALBERT COSTA
University of Barcelona, Barcelona, Spain

ANNE CUTLER
Max Planck Institute for Psycholinguistics, Nijmegen, The Netherlands

and

NURIA SEBASTIÁN-GALLÉS
University of Barcelona, Barcelona, Spain

In three experiments, listeners detected vowel or consonant targets in lists of CV syllables constructed from five vowels and five consonants. Responses were faster in a predictable context (e.g., listening for a vowel target in a list of syllables all beginning with the same consonant) than in an unpredictable context (e.g., listening for a vowel target in a list of syllables beginning with different consonants). In Experiment 1, the listeners' native language was Dutch, in which vowel and consonant repertoires are similar in size. The difference between predictable and unpredictable contexts was comparable for vowel and consonant targets. In Experiments 2 and 3, the listeners' native language was Spanish, which has four times as many consonants as vowels; here effects of an unpredictable consonant context on vowel detection were significantly greater than effects of an unpredictable vowel context on consonant detection. This finding suggests that listeners' processing of phonemes takes into account the constitution of their language's phonemic repertoire and the implications that this has for contextual variability.

Languages differ considerably in terms of the size of their phonemic repertoires. Although about two thirds of the world's languages make do with an inventory in the range of two to four dozen segments, the range of repertoire size is actually very extensive. At one end occur Polynesian languages, with fewer than 20 phonemes, and at the other end, African languages, with 140 or more (Maddieson, 1984). The relationship between different kinds of segments within the inventories also varies. Again, there is a preferred pattern whereby languages tend to contain almost twice as many consonants as vowels; but the vowel/consonant ratio ranges from .065 to over 1.3.

Are there perceptual consequences of cross-linguistic differences in phonemic inventory size? Assuming that languages are roughly equivalent in the number of word contrasts that they encode, it is reasonable to expect that languages with small inventories, allowing relatively few phoneme contrasts, will tend to have a larger number of homophonic word forms, or greater word length. Languages with larger inventories, supporting relatively many contrasts, will be able to get by with shorter words and fewer homophones. Thus it can be argued that a reduction in the number of phonetic contrasts to be processed trades off against an increase in the average number of phonemes to be processed per word, or an increase in the number of ambiguous words that must be distinguished via contextual information; hence, phoneme inventory size itself should not carry implications for perceptual complexity.

The relative distribution of phonemes within the phonetic repertoire raises, however, different issues. Roughly speaking, consonants and vowels tend to alternate in word forms. Accordingly, the absolute size of a language's consonant inventory could provide an indication of the number of different transition types that will be realized for any given vowel, while the size of the vowel inventory could likewise give a rough guide to how many different transitions will be realized for each consonant. If the relative number of vowels and consonants in a language is unbalanced, the potential for variability of vowels and of consonants, respectively, will likewise be unbalanced.

Needless to say, this potential may not be realized; other factors may operate to ensure that vowel and consonant variability is similar across languages. Consider, for instance, the well-known argument of Lindblom (e.g., 1986, 1988) that phonetic contrasts will be realized as clearly as is needed to ensure successful communication, but no more so. One of the corollaries of this argument is that in languages with few phonemes (or few vowels, or few con-

This research was supported by a grant from the Human Frontier Scientific Program, Grant AI96-26 from Acciones Integradas 1996 and Grant PB94-126 from the Spanish Ministerio de Educación y Ciencia. We thank Mark Scholten and Mariëtte Koster for assistance with the experiment in Dutch, C. Cortes for assistance with the Spanish experiments, Christophe Pallier and James McQueen for useful discussions and comments on the paper, and Ann Bradlow and Björn Lindblom for helpful reviews. Correspondence should be addressed to A. Cutler, Max Planck Institute for Psycholinguistics, P.O. Box 310, 6500 AH Nijmegen, the Netherlands (e-mail: anne.cutler@mpi.nl).

sonants), the smaller set of available contrasts will be inherently more distinctive and hence will not need to be as clearly realized. In an unbalanced inventory, under this account, the smaller set would be realized with less precision, or effectively with greater variability, which in turn would introduce greater variability into the set of transitions available for each member of the larger set. This would cancel out the potential differences in variability offered by the inventory size imbalance.

Studies of this issue have so far been few, and have principally addressed issues of speech production. Vowel inventory size has played a role in studies of the acoustic realization of vowels, in which languages with small vowel inventories have been compared with languages with large vowel inventories. Thus Manuel (1990) compared the amount of vowel-to-vowel coarticulation in CVC sequences in three Bantu languages and found, for one of the two vowels studied, greater anticipatory coarticulation in the two languages with five-vowel inventories than in the third language, which had a larger number of vowels. This result supports the prediction that vowels will be less clearly realized when there are fewer of them. However, the other vowel that she studied showed no such differences. Flege (1989) measured the precision of tongue positioning in the production of English and Spanish point vowels; he found no difference in articulatory precision, but three of the English point vowels (/i, u, a/) were realized at more extreme positions than their Spanish equivalents, which Flege explained in terms of the greater need to preserve contrast with neighbors in the more crowded English vowel space. Several studies have shown that the vowel spaces of languages with larger vowel inventories are expanded relative to those of languages with smaller vowel inventories (e.g., Bradlow, 1995, for English vs. Spanish; Jongman, Fourakis, & Sereno, 1989, for English and German vs. Greek). Bradlow (1996) compared the discrimination of similar vowel contrasts in English and Spanish; she found that identical stimuli were judged by the two speaker populations relative to their native vowel systems, so that boundaries reported in the continua reflected the boundaries in the natural systems—in the more expanded English and more restricted Spanish vowel spaces, respectively.

There is thus little direct evidence so far that clearly supports the proposal that variability in processing reflects phoneme inventory size. Moreover, the available studies have tended to focus just on the size of the vowel inventory, rather than the vowel/consonant ratio. In the present study, we asked whether the perception of vowels and of consonants varies as a function of vowel/consonant ratio in the listener's language. We addressed this question via the issue of variability just raised: Are effects of variability proportional to the size of the set from which the varying phonemes are taken? If listeners are indeed sensitive to the simple size of their language's available repertoire of vowels and consonants, respectively, then a larger set within which variability can potentially occur should exercise relatively greater perceptual effects than a smaller set.

Effects of contextual variability on phoneme perceptual measures may be easily observed in the laboratory. For example, uncertainty arising from the number of available phonetic environments slows phonetic decisions in the selective attention paradigm (Garner, 1974). In selective attention experiments, listeners perform forced-choice categorization along one dimension of an input under conditions in which another dimension either is held constant or varies. Wood and Day (1975) found that variation in vowels slows the categorization of the initial stop consonant of a CV syllable, and variation in the consonant likewise slows categorization of the vowel (see also Tomiak, Mullenix, & Sawusch, 1987). Shand (1976) showed that variation in the second syllable of CVCV stimuli had a similar effect.

The two-way categorization task is not the only paradigm in which effects of contextual uncertainty appear. In the present study, we used another task in which uncertainty effects have been repeatedly observed—phoneme detection, in which subjects hear lists of syllables (or other speech stimuli) and respond as soon as they detect an occurrence of a specified phoneme. This task now has a venerable history in psycholinguistics and has been shown to display a large range of effects that also occur in other laboratory tasks and in more naturalistic perceptual situations (see Connine & Titone, 1996, for a review); contextual uncertainty effects with this task are robust. Swinney and Prather (1980) found that detection of an initial /b/ in a list of CVC syllables was slower if the vowel in the syllables varied among four possible vowels than if it remained constant across all syllables, and slower again if it varied among eight possible vowels than if there were only four possibilities. Similar to this effect is the effect of matching target specification observed by Mills (1980) on detection of initial consonants in lists of syllables. Subjects were supposed to ignore the vowel in the target specification, but apparently they could not; when the target was given as /be/, for example, /b/ was detected faster in /be/ than in /bo/. Further, Dijkstra, Schreuder, and Frauenfelder (1989) found that a visually presented letter representing a consonant had a similar effect on vowel detection in spoken CV syllables: Again, the subjects were not supposed to consider the consonant letter in the target specification, but nevertheless /a/ was detected faster in /pa/ if the vowel target had been specified with the written example PA than with KA. (Note that the effectiveness of a visual target specification here suggests that the match effect does not occur simply at an acoustic level; see Cutler, Butterfield, & Williams, 1987, for further evidence that matching effects in phoneme detection involve relatively abstract representations.)

Thus this task can be relied on to produce effects of contextual uncertainty, and we can therefore use it to examine whether such effects differ across languages in proportion to differences in vowel versus consonant repertoire size. Unfortunately, none of the experiments cited above compared detection of vowel versus consonant targets; thus we have as yet no basis for predicting whether the effects of vowel unpredictability on consonant detection, and of consonant unpredictability on vowel detection, should underlyingly be equivalent. Note that a comparison across languages with different vowel/consonant ratios would not in principle be vitiated by differences in the underlying size of the effects, since our predictions concern only differences in the size of each effect separately as a function of differences in this ratio. Nevertheless, we began the present series of experiments in a language with a balanced vowel/consonant ratio, with the aim of comparing the results obtained there to those obtained in a language with an unbalanced vowel/consonant ratio. On the default assumption that the effects of vowel unpredictability on consonant detection and of consonant unpredictability on vowel detection are, ceteris paribus, equivalent, we predicted that these effects would be equal in the language with the balanced ratio of vowels to consonants. In the language with the unbalanced vowel/consonant ratio, on the other hand, we predicted that listeners would show greater effects of contextual unpredictability when detecting targets from the smaller set (where the unpredictability occurs within the larger set) than when detecting targets from the larger set (where the unpredictability occurs within the smaller set).

In order to compare the differing phonemic repertoires in as simple a manner as possible, we did not (as, for example, Swinney and Prather, 1980, did) include differing levels of variability within the experiment. In each experiment reported below, the contexts within which a target phoneme was detected were either predictable or unpredictable; when they were unpredictable, they varied over four possible phonemes. The same five target vowels and five target consonants were chosen for each language, and they were realized in the same 20 CV syllables. Thus the contextual variability within the experiment was matched across languages. Any differences in the relative effects of vowel versus consonant unpredictability between the language with the balanced vowel/consonant ratio and the language with the unbalanced ratio could then be ascribed to the underlying distributional difference between the languages, rather than to differences in the experimental situation itself.

In Experiment 1, then, we investigated the effect of unpredictable vowel information on consonant detection and the effect of unpredictable consonant information on vowel detection in Dutch. This language has 16 vowels and 19 consonants (Booij, 1995) and thus provided as near to equivalent vowel and consonant inventories as was easily available to us (languages with exactly balanced inventories, such as the Papuan language Dani [Maddieson, 1984] do exist, but they are less accessible to psycholinguistic laboratories). Listeners performed a phoneme detection task with five possible vowel targets and five possible consonant targets; the target item appeared in varying positions in a list of CV syllables. The choice of the 10 phonemes that we used was motivated by their availability in the phonological repertoires of both Dutch and Spanish (the language of Experiment 2); although the acoustic realizations in the two languages are not absolutely identical, they are certainly similar, and the phonemic transcriptions in the International Phonetic Alphabet (IPA) are the same in each language. All 10 phonemes could furthermore occur in CV syllables in both languages; CV was the only syllable structure with which we could construct 20 different monosyllabic stimuli with identical IPA transcriptions in the two languages. As in Swinney and Prather's (1980) study, the nontarget dimension of the CV syllables was either predictable across all items of a given list or was unpredictable. However, our experiment differed from Swinney and Prather's in that (1) there was only one either/or predictability comparison, (2) we used more than one phoneme as target, (3) our target set included both vowels and consonants, and (4) the specified target varied from trial to trial.

EXPERIMENT 1
Dutch

Method

Subjects. Forty undergraduates from the University of Nijmegen, all native speakers of Dutch with normal hearing, took part in this experiment. They received a small payment for their participation.

Materials. Five vowels and five consonants were selected as targets. The five vowels were /a/, /e/, /i/, /o/, and /u/. The five consonants were /p/, /t/, /k/, /l/, and /s/. From these 10 phonemes, 20 CV syllables were constructed: /pe/, /pi/, /po/, /pu/, /ta/, /te/, /to/, /tu/, /ka/, /ki/, /ko/, /ku/, /sa/, /se/, /si/, /su/, /la/, /le/, /li/, and /lo/; it can be seen that each of the 10 (target) phonemes occurred in four different contexts. These 20 syllables were recorded by a female native Dutch speaker.

Two main bivalent variables were included in the experiment: type of target (vowel or consonant) and type of context (predictable or unpredictable). Crossing these two variables, the four experimental conditions were obtained: consonant detection with unpredictable context, consonant detection with predictable context, vowel detection with unpredictable context, and vowel detection with predictable context. In unpredictable-context lists, the syllables preceding the target-bearing syllable could include all the above types of consonants and vowels (with the exception of the target phoneme); thus an unpredictable-context list for either /p/ or /u/ could be /ko se to si pu/. In predictable-context lists, for consonant detection all syllables contained the same vowel as that of the target (e.g., for /p/: /ku su tu su pu/); likewise for vowel detection, all syllables in the lists contained the same consonant as that of the target (thus for /u/: /po pe po pi pu/).

Experimental lists varied in length between five and nine syllables, and the target was always positioned at the end of the list. Each potential target syllable occurred as target once in each condition; thus there were 20 experimental lists in each condition. To avoid strategies of pressing the response key at the end of every list, there were also five filler lists per condition; in these filler lists, the target appeared elsewhere than at the end. The filler lists were between three and five syllables long, and the target could appear in any of the first four positions. The filler lists were constructed from the same set of 20 syllables as the experimental lists. Ten practice lists were constructed, and all four experimental conditions were represented in this set.

Procedure. The four experimental conditions were presented blocked, and each subject was tested with all conditions. The four blocks (conditions) were presented in four different orders, following a Latin square design; thus, there were four possible orders of presentation. Equal numbers of subjects were tested in each order.

Table 1
Mean Reaction Times (RTs, in Milliseconds) and Percentage of Errors (PEs) in Experiment 1 (Dutch)

Consonant Detection				Vowel Detection			
Fixed Context		Varied Context		Fixed Context		Varied Context	
RT	PE	RT	PE	RT	PE	RT	PE
395	2.23	468	5.65	369	6.58	444	8.68

Each block contained an equal number of different list lengths (four lists per length). The list presentation inside each block was randomized for each subject with the following restrictions: First, no more than six experimental lists could appear in a row; second, no more than three lists of the same length could appear in a row; and third, the same target phoneme never appeared more than twice in a row.

Each trial had the following sequence: A letter representing the phoneme target appeared on the screen for 1,000 msec, and then syllables in the list were played one after the other with an intersyllable interval of 500 msec. After the last syllable of the list had been presented, subjects had up to 2,500 msec to respond before the next trial started. After each block there was a pause of 4,000 msec. During the practice trials, the subjects were given feedback if they made a mistake: A false alarm response, or failure to detect a target, caused the word ERROR to be printed on the center of the screen after the end of the trial for 1,000 msec. Reaction time (RT) was measured from the onset of the target syllable.

Subjects were encouraged to use an acoustic representation of the phoneme target and they were instructed not to use the orthographic code. The choice of cross-linguistically matched materials raised one problem for the Dutch experiment in that the target sound /u/ is represented in Dutch orthography by the bigraph OE. Since we did not wish to make the visual specification of any one target sound more complex than any other, we used the letter U to represent /u/, although in Dutch orthography U most often corresponds to the phoneme /y/. Subjects were specifically instructed that given this target specification they should listen for the sound /u/, and they appeared to experience no difficulty with this. (Note that a /u/–U mapping does occur in Dutch in some loan words, such as *computer*, and in previous experiments by Otake, Yoneyama, Cutler, & van der Lugt, 1996, Dutch subjects experienced no difficulty monitoring in spoken Japanese words for the sound /g/ given the specification G, although G in Dutch orthography represents the velar fricative /x/ in native words and the stop /g/ only in loan words such as *goal*.)

Subjects were instructed to make their responses in each trial as quickly and as accurately as possible. They were also asked to pay attention to the whole list of syllables because the target could appear at any position in the list. Subjects were tested individually in a sound-attenuated room; the stimuli were presented binaurally over Sennheiser HD-222 headphones. The experimental session was controlled by a personal computer using the program NESU. The experiment lasted for about 30 min.

Results and Discussion

RTs were measured from the onset of the target syllable (i.e., RTs for consonant targets were measured at target onset, but RTs for vowel targets were measured at the onset of the preceding consonant). RTs longer than 1,500 msec and shorter than 150 msec were discarded from the analyses. One list (in the unpredictable-context consonant detection condition) had to be excluded because the target inadvertently appeared twice (i.e., a syllable other than the intended target-bearing syllable contained the same consonant). To keep the design balanced, the same target syllable was discarded from the analyses of the remaining three conditions. Therefore, there were only 19 responses per subject per condition. Mean RTs and mean missing data rates for each of the four experimental conditions, averaged across subjects and items, are shown in Table 1.

The data were submitted to separate analyses of variance (ANOVAs) across subjects ($F1$) and items ($F2$), with type of target (vowel/consonant) and type of context (predictable/unpredictable) as main factors. Full details of the statistical analysis are reported only for those effects that reached significance on both analyses.

Analyses of the missing data rates showed that subjects made significantly more errors in the unpredictable context (7.17%) than in the predictable context [4.405%; $F1(1,39) = 7.61$, $p < .01$, $F2(1,8) = 10.8$, $p < .02$]. Subjects also made more errors responding to vowels than to consonants (7.63% vs. 3.945%); however, this effect reached significance only in the subjects analysis, not in the items analysis. The interaction of the two factors was not significant (both Fs < 1).

The analysis of RTs showed the same pattern of results. The context type effect was significant in both analyses, with responses in the predictable context (382 msec) being faster than those in the unpredictable context [456 msec; $F1(1,39) = 103.53$, $p < .001$, $F2(1,8) = 87.08$, $p < .001$]. The effect of target type was again significant in the subjects analysis (subjects responded more rapidly to vowels than to consonants: 407 vs. 432 msec) but failed to reach significance in the items analysis. Again the interaction between target type and context type was not significant (both Fs < 1).

In this experiment, context effects were observed for all targets, in contrast to the findings of Swinney and Prather (1980), who found significant context effects for the target /b/ but not for the target /s/. Further, the present results demonstrate that context effects appear in a standard phoneme detection experiment in which the nature of the target varies from trial to trial, and not only when, as in Swinney and Prather's study, there is only one target per block. Indeed, the effects that we have observed were much larger than the mean 15-msec context effect reported in the earlier study. Moreover, we found no consistent difference in detection times or error rates for vowel versus consonant targets, and we observed that vowel detection was just as sensitive to consonantal unpredictability as consonant detection was to vowel unpredictability. Every one of the 10 target phonemes showed a substantial context effect, and, in this language, in which inventory size for consonants and vowels is almost equivalent, the sizes of the average effects for consonants and for vowels were hardly distinguishable (73 vs. 76 msec).

The vowel/consonant ratio within our experiment was, of course, 1:1. We used five vowels and five consonants, and the materials included four alternative contexts for each individual phoneme target. Thus it is arguable that

Table 2
Mean Reaction Times (RTs, in Milliseconds) and Percentage of Errors (PEs) in Experiment 2 (Spanish)

Consonant Detection				Vowel Detection			
Fixed Context		Varied Context		Fixed Context		Varied Context	
RT	PE	RT	PE	RT	PE	RT	PE
400	2.10	460	3.42	408	1.96	499	2.89

it was this feature of the design, rather than the relatively balanced background vowel/consonant ratio in Dutch, that resulted in uncertainty effects of equivalent size for each target type. In Experiment 2, we maintained the matched within-experiment ratio, but induced a cross-experimental comparison of phoneme repertoire in that we switched the language of the experiment to one with an unbalanced vowel/consonant repertoire. Spanish has a simple five-vowel repertoire: /a/, /e/, /i/, /o/, and /u/, with the members maintaining a high degree of acoustic distinctiveness (Skelton, 1969). The consonant repertoire of Spanish, in contrast, is much more varied, with 20 members in the variant in which this experiment was conducted (Castilian; see Harris, 1969; Maddieson, 1984; Stockwell & Bowen, 1965). Experiment 1 having established that both vowel detection and consonant detection are subject to effects of contextual uncertainty, Experiment 2 allowed us to investigate effects of phoneme repertoire size, while keeping all other features of the experimental situation constant. We predicted that if subjects are indeed sensitive to the phonetic choices inherent in their language, then the size of the phonetic population within which the variation occurs should determine the size of the effect of the introduced uncertainty—the greater the population of potential alternatives, the greater the uncertainty and hence the greater the effect. In other words, in this language, with a 1:4 vowel/consonant ratio, there should be a significantly greater effect of consonant unpredictability on the detection of vowels than of vowel unpredictability on the detection of consonants.

EXPERIMENT 2
Spanish

Method

Subjects. Forty native Spanish speakers without early (i.e., before school) exposure to another language[1] were tested in this experiment. They were all undergraduate psychology students at the University of Barcelona, and they received extra course credits for their participation.

Materials and Procedure. A new recording of the syllables was made, by a female native Spanish speaker, thus ensuring that the listeners were presented with tokens that evoked native language processing. The program controlling the experiment was EXPE (Pallier, Dupoux, & Jeannin, 1997). In all other respects the materials and procedure were exactly the same as in Experiment 1. Subjects were again instructed not to use an orthographic code in monitoring for the target. In the Spanish, as in the Dutch materials, one target phoneme was particularly susceptible to grapheme–phoneme mismatch effects; in this case, it was the use of the letter K as target specification for the phoneme /k/, which would usually be represented in Spanish orthography by the letter C or the letter Q. Subjects were instructed to listen for /k/ given the target specification K, and again they experienced no difficulty with this mapping. Subjects were tested individually in a sound-attenuated booth; the materials were presented binaurally, via Sennheiser HMD-224 headphones.

Results and Discussion

Because the same design was employed as that in Experiment 1, and indeed the experiment in Barcelona was conducted concurrently with the experiment in Nijmegen, the same list as in the previous experiment contained an inadvertent fault and had to be discarded. The same criteria to eliminate data were employed in this experiment. Mean RTs and mean missing data rates for each condition, averaged across subjects and items, are shown in Table 2.

All data were again submitted to ANOVAs across subjects and across items. In the missing data analysis, the only significant effect was the context type effect, in the subjects analysis only; as in Experiment 1, subjects missed more target phonemes in the unpredictable-context conditions (3.16%) than in the predictable-context conditions (2.04%), but here the error rate was in general so low that the difference did not reach significance. All other effects in the subjects analysis of the pattern of missing data, and all effects in the items analysis, were nonsignificant.

In the RT analysis, the effect of context type was highly significant. Subjects' responses in the unpredictable-context conditions were significantly slower than those in the predictable-context conditions [404 vs. 476 msec; $F1(1,39) = 96.85$, $p < .001$, $F2(1,8) = 163.0$, $p < .001$]. The difference between RTs to consonants (430 msec) and vowels (454 msec) was again significant in the subjects analysis but not in the items analysis. Thus the results of Experiment 2 confirm those of Experiment 1: Both vowel detection and consonant detection show strong effects of contextual uncertainty, and standard phoneme detection procedures elicit these effects.

Also as in Experiment 1, every one of the 10 target phonemes showed a substantial context effect. In this case, however, the size of the average effect for vowel targets (91 msec) was greater than that for consonants (60 msec). The interaction of the target type and context type effects here reached significance [$F1(1,39) = 5.62$, $p < .03$, $F2(1,8) = 6.18$, $p < .04$]. Separate post hoc analyses showed that the effect of the context manipulation was significant for both consonant targets [$t1(39) = 6.06$, $p < .001$, $t2(4) = 5.76$, $p < .005$] and vowel targets [$t1(39) = 8.88$, $p < .001$, $t2(4) = 8.67$, $p < .001$]; nevertheless, the interaction indicates that the effect of consonant uncertainty on vowel targets was significantly greater than the effect of vowel uncertainty on consonant targets.

This is exactly the predicted result; in this language, with its imbalance in the size of the consonant and vowel inventories, there is an equivalent imbalance in the uncertainty effects exercised by each type of phonetic context. Despite the fact that the vowel/consonant ratio in the experiment itself was 1:1, the asymmetry in the background ratio in the language was reflected in the asymmetric uncertainty effects.

Table 3
Mean Reaction Times (RTs, in Milliseconds) and Percentage of Errors (PEs) in Experiment 3 (Spanish Listeners, Dutch Recording)

Consonant Detection				Vowel Detection			
Fixed Context		Varied Context		Fixed Context		Varied Context	
RT	PE	RT	PE	RT	PE	RT	PE
417	3.38	464	5.25	450	21.63	531	22.38

The results of Experiments 1 and 2 thus clearly suggest that Spanish listeners, but not Dutch listeners, are more sensitive to consonantal variation when detecting vowels than to vowel variation when detecting consonants. However, it is of course the case that not only the listeners but also the recording of the CV syllables differed across the two experiments. This was an important aspect of the experimental design in that it was necessary for listeners to perform the task in a processing mode as near as possible to that normally used for their native language. Nevertheless, it is conceivable that in some way the actual acoustic variability within the 20 syllables used in the experiment might have been greater in the Spanish materials set than in the Dutch set. Although only a single recording of each syllable was used in each materials set, it is, for instance, possible that individual phonemes might have been realized with greater variability in their four syllables (e.g., /l/ in /la, le, li, lo/) in the one as opposed to the other set. In our final experiment, therefore, we undertook a control comparison by presenting the Spanish listeners, whom we claim to have greater sensitivity to consonantal variation than to vowel variation, with the Dutch recording of the syllables, which in Experiment 1 produced no difference in the size of uncertainty effects for consonants versus vowels.

One might expect that the Spanish listeners would in general have more difficulty performing the task with the nonnative materials. However, our claim is that the uncertainty effects observed in Experiment 2 reflect the Spanish listeners' accumulated experience with the asymmetric phonemic repertoire of their native language and the potential for phonetic uncertainty offered by such asymmetry. Thus we expected that these listeners would maintain their response pattern despite the fact that the input may not have mapped so easily onto their native phonology. In other words, they should manifest the differential sensitivity to consonantal as compared with vowel uncertainty, which they showed with the Experiment 2 materials, to an equivalent degree with the Experiment 1 materials, even though the former were recorded by a speaker of their own language while the latter were recorded by a speaker of another language.

EXPERIMENT 3
Spanish Listeners, Dutch Recording

Method

Subjects. Forty Spanish-speaking subjects from the same population as that used for Experiment 2 were tested in this experiment; none had taken part in Experiment 2.

Materials and Procedure. The materials used were the syllables as recorded for Experiment 1. The experimental list generation program was corrected to remove the inadvertent double occurrence of a target phoneme in one list. The procedure was as in Experiment 2. Because the acoustic targets for the phonemes used in our study were not identical in Dutch and in Spanish, the subjects were informed that the syllables were pronounced by a foreign speaker and might sound a little unusual.

Results and Discussion

Mean RTs and mean missing data rates for each condition, averaged across subjects and items, are shown in Table 3. ANOVAs were carried out in the same manner as for the preceding experiments.

In the missing data analysis, no effect reached significance in either the subjects or the items analysis. However, it can immediately be seen that the proportion of missed targets was much higher in this experiment than in Experiments 1 and 2. Although this might suggest that subjects found the foreign pronunciation in general harder to process, in fact there were quite localized sources of difficulty; two of the vowels, as realized by the Dutch speaker, caused particular problems. The vowel /e/ sounded to these Spanish listeners like the diphthong /ei/, and the vowel /o/ sounded like the diphthong /ou/. As a consequence, when subjects were listening for the target /i/ they frequently responded in error to an occurrence of /e/ earlier in the list, and when listening for /u/, they frequently responded in error to an earlier occurrence of /o/; 85.5% of the errors to vowel targets were to /i/ or to /u/.

Accordingly, we conducted two separate analyses of the RT data—one across the whole data set and one omitting the eight target items containing the vowels /u/ and /i/. Across the whole data set, the effect of context type was highly significant: Subjects' responses were significantly slower in unpredictable contexts (497 msec) than in predictable contexts [434 msec; $F1(1,39) = 74.07$, $p < .001$, $F2(1,8) = 104.19$, $p < .001$]. Responses to consonants (441 msec) were significantly faster than those to vowels [490 msec; $F1(1,39) = 27.9$, $p < .001$, $F2(1,8) = 9.17$, $p < .02$]. Every one of the 10 target phonemes again showed faster responses in predictable than in unpredictable contexts, and, as in Experiment 2, the average effect for vowel targets (81 msec) was substantially larger than that for consonant targets (47 msec). The interaction of target type and context type effects was significant [$F1(1,39) = 7.32$, $p < .01$, $F2(1,8) = 4.36$, $p < .07$]. Post hoc analyses showed that once again the difference between predictable and unpredictable contexts was significant both for consonant targets [$t1(39) = 5.31$, $p < .001$, $t2(4) = 5.76$, $p < .005$] and for vowel targets [$t1(39) = 8.0$, $p < .001$, $t2(4) = 8.65$, $p < .001$].

When the target syllables containing /i/ and /u/ were dropped (both in the cases when their vowel formed the target and in the matched cases when their consonant formed the target), an analysis of the RTs across subjects showed a significant advantage for predictable over unpredictable contexts [$F1(1,39) = 27.2$, $p < .001$], a significant advantage for consonant over vowel targets [$F1(1,39) = 79.72$, $p < .001$], and again a significant interaction between these two effects [$F1(1,39) = 5.91$, $p < .02$]. The mean size of the context effect for vowels (82 msec) was still substantially greater than the mean effect for consonants (44 msec). An analysis of the missing data patterns for this subset of the data revealed no significant effects.

Despite the fact that the Spanish listeners clearly did not perceive the tokens in this materials set as native productions, they nevertheless once again produced a response pattern displaying greater sensitivity to consonantal uncertainty than to vowel uncertainty, in conformity with the distributional patterns of their native phonemic repertoire.

In a final statistical analysis, we compared the context effects (RT to targets in unpredictable context minus RT to targets in predictable contexts) in all three experiments. There was no main effect of experiment; overall context effects in all three studies were comparable. There was a main effect of phoneme type, with greater average context effects for vowels than for consonants [$F1(1,117) = 9.3$, $p < .005$, $F2(1,8) = 7.55$, $p < .05$], but t tests revealed that this effect was significant for Experiment 2 [$t1(39) = 3.42$, $p < .001$, $t2(8) = 3.21$, $p < .02$] and for Experiment 3 [$t1(39) = 2.43$, $p < .02$, $t2(8) = 2.09$, $p < .07$], but not for Experiment 1. A subanalysis of the context effects comparing just Experiments 2 and 3 revealed a significant effect of phoneme type only [$F1(1,78) = 12.84$, $p < .001$, $F2(1,8) = 11.83$, $p < .01$], with no interaction between experiment and phoneme type ($F1$ and $F2 < 1$); that is, the context effects shown by the Spanish listeners in their native language and in the Dutch materials did not differ in statistical significance.

GENERAL DISCUSSION

Our cross-linguistic comparison has provided clear evidence that the effects of contextual unpredictability on phoneme detection alter as a function of the relative size of the phonemic repertoire within which the context may in principle vary. Vowel detection was always slower under conditions of consonantal unpredictability, and consonant detection was always slower under conditions of vowel unpredictability. However, when the comparison was carried out in Dutch, a language with approximately balanced inventories of vowels and consonants, approximately equivalent uncertainty effects were observed for vowel targets and for consonant targets. By contrast, when the comparison was performed in Spanish, a language with four times as many consonants as vowels, the effects of consonant uncertainty on vowel detection were significantly greater than the effects of vowel uncertainty on consonant detection. The Spanish listeners produced the same asymmetric effects even with the same materials with which the Dutch listeners had shown symmetric effects. Thus a larger population of potential phonetic contexts within which variation could occur seemed to be directly translated into larger deleterious effects of the contextual unpredictability on the speed of subjects' responses.

These experiments have also answered a number of questions arising from earlier findings. Contextual uncertainty effects in phoneme detection are robust: It is always more difficult to detect a phoneme target if the surrounding phonetic context varies than if it is held constant. Uncertainty effects are observed for vowel targets when consonantal context varies, and for consonant targets when vowel context varies; there is no categorical difference between the two phoneme classes in their susceptibility to context effects. Moreover, the effects are observed in a standard phoneme detection design with a range of potential targets; they are in fact stronger in this situation than in the highly constrained single-target design used by Swinney and Prather (1980). Nevertheless, the effects are subject to variation, as a function of how the language's phonemic repertoire is made up.

Recall that our repertoire size manipulation was not carried out within the experimental materials. The sets of target phonemes in Experiments 1 and 2 were identical, as were the lists of syllables presented to the listeners. Only the background repertoire of the language—the sets from which the five vowels and five consonants that served both as targets and as contexts were drawn—differed across Experiment 1 on the one hand and Experiments 2 and 3 on the other. Thus the listeners' knowledge of the phonemic repertoire of their language appeared to play a role in the amount of uncertainty engendered in them by an unpredictable phonetic context. Spanish listeners found an unpredictable context that could range over four potential and actual vowel alternatives less problematic than an unpredictable context that in fact ranged over 4 consonantal alternatives but in principle could have ranged over 20. Dutch listeners found unpredictable contexts of the same actual complexity and larger but comparable notional complexity comparably difficult. Thus the listeners were unable to restrict their expectations of contextual variability to that which actually occurred in the experiment; this actual variability was equivalent for the two listener groups, but their response patterns were not. That the listeners seemed unable to restrict their expectations in conformity with the experimental context is also revealed by another aspect of the results. It is in fact not the case that every phoneme of each language could occur in a CV syllable. Both languages have the consonant [ŋ], for example, which cannot appear in an onset; and 4 of the 16 vowels of Dutch cannot occur in open syllables. Thus if listeners had been able to confine their expectations to open syllables, they might, when detecting an onset consonant, have been able to expect only variation among the 12 vowels that could follow the onset of an open syllable, rather than among the 16 vowels that could follow the total population of onsets. This would have resulted in a slightly asymmetric 12:18 expected ratio for Dutch; yet the Dutch listeners' response patterns showed no trace of asymmetry at all.

Thus this simple laboratory task has allowed us to observe effects of listeners' background expectations of phonetic variability based on their experience with their native language. In the predictable context, listeners were presumably able to form precise expectations of the form of the target given the target specification plus the information, provided by the first syllables in the incoming sequence, as to the nature of the vowel. In the unpredictable context, they had only the target specification itself to work with; in theory, they could have constructed precise expectations of the four relevant syllables in the experiment, but in practice they apparently did not (note that phoneme detection becomes enormously harder if subjects have to retain more than one target specification in memory; Foss & Dowell, 1971; Steinheiser & Burrows, 1973). In other words, we assume that the context effects observed in our study (and its predecessors) are more accurately described as facilitation of detection in the predictable-context condition, rather than inhibition of detection in the unpredictable-context condition. In the unpredictable-context condition, listeners' knowledge of how varied the realization of the target phonemes could potentially be became relevant. Note once again that actual variability in the target realization across the two contexts could play no role; not only were there just four actual realizations in each unpredictable context in all experiments, but also the target syllables used in the predictable- and unpredictable-context conditions were physically identical. The effect must be located at the level of listener expectation; clearly, expectations differed across the experiments, in just such a way as to reflect the underlying potential for variability in the language. Even in the context of a constrained laboratory task, listeners thus could not prevent their responses from being influenced by the phonological structure of their native language—specifically, by their accumulated experience of whether variability for vowels and for consonants is symmetric.

Recent research has suggested a number of ways in which the processing demanded by vowels and consonants may differ. Note first that vowel and consonant processing can be interdependent. Whalen (1989, on the basis of earlier work by Mermelstein, 1978) examined the perceptual interpretation of vowel duration in CVC syllables for which the vowel duration was relevant to both a vowel and a consonant judgment—for example, the voicing distinction between syllable-final [d] and [t] and the distinction between the (longer) vowel [æ] and the (shorter) vowel [ɛ]. Listeners heard synthesized tokens in which only vowel duration was varied and categorized each token as "bad," "bed," "bat," or "bet." Judgments about the vowel and consonant were not made independently; for example, choosing a [d] could make the choice of [æ] less likely. Whalen further found interaction for the acoustic shape of fricative noise, which can serve both as a cue regarding whether the fricative is [s] or [ʃ] and as a cue regarding whether the following vowel is rounded or not. Thus vowel and consonant processing may be required to draw on the same input information. Nonetheless, differences in response patterns for vowels and for consonants have been observed in a number of perceptual tasks. Van Ooijen (1994, 1996), for instance, found that when listeners are asked to turn a nonword into a real English word by changing just one phoneme, they more readily alter a vowel than a consonant. Vowels, proposed van Ooijen, are treated by listeners as inherently mutable; listeners appear to expect that vowels will show more extensive effects of surrounding phonetic context and of other factors such as position in the word than consonants will. Van Ooijen's experiment was carried out in English, a language with a very large vowel repertoire but an even larger consonant repertoire.

Recent results with the phoneme detection task have likewise been interpreted as indicating that listeners are inherently cautious in making detection responses to vowel targets; in both English and Spanish, vowel detection time is inversely related to target duration (Cutler, van Ooijen, Norris, & Sánchez-Casas, 1996); that is, the shorter the vowel, the longer listeners take to detect it. This pattern does not appear with English consonants (van Ooijen, 1994). The effect could once again not be ascribed solely to the actual contextual variation realized in the experiment, since the English listeners in one of Cutler et al.'s (1996) experiments showed no difference in the pattern of the duration correlation as a function of whether the stimulus items in the experiment could contain any of the vowels of English or only vowels from a restricted set of five. Note that in the present study we observed weak effects indicating an overall advantage for consonant over vowel detection for the Spanish listeners in their own language (Experiment 2), but a slight effect in the opposite direction for the Dutch listeners in their own language (Experiment 1). However, the present study used highly constrained and simple materials, and thus it is hardly meaningful to compare RTs with those of these earlier studies, which used real words and word-like nonwords. Here, we focused particularly on the relative size of uncertainty effects and found them to be asymmetric.

The distribution of vowels versus consonants in speech is a function not only of a language's phonemic repertoire, but also of the permissible syllable structures of the language and their relative frequency of occurrence; nevertheless, the relative amount of vowel versus consonant processing that these factors combine to require for speech input in a given language embodies consequences for the speech processor. And differences in relative distinctiveness of one versus the other type of phoneme may cause the processor to direct attention to a greater extent to different processing routines.

Our findings are here consistent with evidence that languages differing in phoneme repertoire structure may consequently differ in terms of which aspects of language production demand greater versus lesser attention. Thus Nettle (1994) analyzed data on vocal amplitude during spontaneous conversation by speakers of 15 languages; he found a very high correlation between amplitude and the languages' vowel inventories. Average speech amplitude was higher in languages with fewer vowels (and hence a lower vowel/consonant ratio) and lower in languages in which the vowel/consonant ratio was higher and the functional load carried by vowels accordingly greater. Dutch and Spanish were included in Nettle's sample (although note that a more complicated measure of possible contrasts led him to a vowel/consonant ratio of 1.11 for Dutch and 0.43 for Spanish, rather than the ratios of 0.84 and 0.25 which absolute phoneme inventory counts warrant). With these two languages, Nettle found that the Spanish speakers spoke more loudly than the Dutch speakers. Nettle proposed that the greater the functional load carried by the vowels—which are the most sonorous portions—of a language, the less amplitude is required to convey speech contrasts efficiently to the listener; when a relatively greater proportion of the functional load falls on the (less sonorous) consonants, they require greater amplitude to be efficiently conveyed. Accordingly, Nettle suggested that languages differ in terms of the types of demands they make on the speaker. Manuel (1990) argued similarly on the basis of her cross-linguistic comparison of vowel-to-vowel coarticulation. Languages may not differ in terms of how intrinsically difficult they are to produce, but they certainly differ in terms of which aspects of the production process require more or less attention.

Similarly, languages differ regarding the types of demands their phonology makes on the listener. This is obvious insofar as languages encode information in phonological dimensions that are not exploited in other languages; thus languages that distinguish words from one another by tone require different processing from languages that do not; languages with vowel harmony can indirectly provide the listener with word-boundary information (Suomi, McQueen, & Cutler, 1997) of a sort unavailable in languages without vowel harmony, and so on. It is perhaps less obvious, but now well supported by empirical evidence, that languages differ in terms of which types of phonological information support a given aspect of processing—for instance, segmentation of utterances into words (Cutler, Mehler, Norris, & Seguí, 1986; Cutler & Norris, 1988; Otake, Hatano, Cutler, & Mehler, 1993; Sebastián-Gallés, Dupoux, Seguí, & Mehler, 1992). Since all languages contain both vowels and consonants, and the processing of the two types of phonemes is not qualitatively different, it is even less obvious that the phonemic repertoire itself should be the source of language-specific differences in processing demands. Nonetheless, the evidence now suggests that this is so. It has recently been reported that vowel discrimination in different languages can require different numbers of perceptual dimensions (Fox, Flege, & Munro, 1995). And it now appears that in the process of phoneme detection, listeners are sensitive to the repertoire of possibilities their language offers. Accumulated experience with the perception of the phonemes of their native language gives rise to accurate expectation of the potential scope for variation given a nonpredictable context; this expectation is not overridden by actual limitations in the experimental variability. Our findings thus provide yet one more piece of evidence for language-specific differences in the nature of the processing required by speech input.

The fact that the asymmetry was observed in the responses of Spanish listeners with both the materials recorded by a Spanish speaker and the materials recorded by a Dutch speaker suggests further that the effects of the native language repertoire can carry over to the processing of foreign languages. This finding joins a number of similar demonstrations; thus French listeners are sensitive to syllabic structure not only in their native language but also in English (Cutler et al., 1986) and Japanese (Otake et al., 1993), although the native listeners in neither language show the same response pattern as French listeners; Japanese listeners respond more rapidly to syllable-final than to syllable-initial nasal consonant targets not only in their own language but also in English, though English listeners again fail to show this pattern (Cutler & Otake, 1994); they also detect CV targets equivalently rapidly in open and closed syllables both in their own language and in French (Otake, Hatano, & Yoneyama, 1996), although French listeners respond otherwise (Cutler et al., 1986), and English listeners respond more rapidly and more accurately to nasal consonant targets than to vowel targets not only in English materials but also in Japanese materials, although Japanese listeners' responses in their native language again show no such effect (Cutler & Otake, 1994). Listening to both the native language and foreign languages, in other words, is conditioned by our native-language experience.

REFERENCES

Booij, G. (1995). *The phonology of Dutch*. Oxford: Oxford University Press.

Bosch, L., Costa, A., & Sebastián-Gallés, N. (1997). *The perceptual magnet effect in early bilinguals*. Manuscript submitted for publication.

Bradlow, A. R. (1995). A comparative acoustic study of English and Spanish vowels. *Journal of the Acoustical Society of America*, **97**, 1916-1924.

Bradlow, A. R. (1996). A perceptual comparison of the /i/–/e/ and /u/–/o/ contrasts in English and Spanish: Universal and language-specific aspects. *Phonetica*, **53**, 55-85.

Connine, C. M., & Titone, D. (1996). Phoneme monitoring. *Language & Cognitive Processes*, **11**, 635-645.

Cutler, A., Butterfield, S., & Williams, J. N. (1987). The perceptual integrity of syllabic onsets. *Journal of Memory & Language*, **26**, 406-418.

Cutler, A., Mehler, J., Norris, D. G., & Seguí, J. (1986). The syllable's differing role in the segmentation of French and English. *Journal of Memory & Language*, **25**, 385-400.

Cutler, A., & Norris, D. G. (1988). The role of strong syllables in segmentation for lexical access. *Journal of Experimental Psychology: Human Perception & Performance*, **14**, 113-121.

Cutler, A., van Ooijen, B., Norris, D., & Sánchez-Casas, R. (1996). Speeded detection of vowels: A cross-linguistic study. *Perception & Psychophysics*, **58**, 807-822.

Cutler, A., & Otake, T. (1994). Mora or phoneme? Further evidence for language-specific listening. *Journal of Memory & Language*, **33**, 824-844.

Dijkstra, T., Schreuder, R., & Frauenfelder, U. H. (1989). Grapheme context effects on phonemic processing. *Language & Speech*, **32**, 89-108.

Flege, J. E. (1989). Differences in inventory size affect the location but not the precision of tongue positioning in vowel production. *Language & Speech*, **32**, 123-147.
Foss, D. J., & Dowell, B. E. (1971). High-speed memory retrieval with auditorily presented stimuli. *Perception & Psychophysics*, **9**, 465-468.
Fox, R. A., Flege, J. E., & Munro, M. J. (1995). The perception of English and Spanish vowels by native English and Spanish listeners: A multidimensional scaling analysis. *Journal of the Acoustical Society of America*, **97**, 2540-2551.
Garner, W. R. (1974). *The processing of information and structure*. Hillsdale, NJ: Erlbaum.
Harris, J. W. (1969). *Spanish phonology*. Cambridge, MA: MIT Press.
Jongman, A., Fourakis, M., & Sereno, J. (1989). The acoustic vowel space of modern Greek and German. *Language & Speech*, **32**, 221-248.
Lindblom, B. (1986). Phonetic universals in vowel systems. In J. Ohala & J. Jaeger (Eds.), *Experimental phonology* (pp. 13-44). New York: Academic Press.
Lindblom, B. (1988). Phonetic invariance and the adaptive nature of speech. In B. A. G. Elsendoorn & H. Bouma (Eds.), *Working models of human perception* (pp. 139-173). London: Academic Press.
Maddieson, I. (1984). *Patterns of sounds*. Cambridge: Cambridge University Press.
Manuel, S. (1990). The role of contrast in limiting vowel-to-vowel coarticulation in different languages. *Journal of the Acoustical Society of America*, **88**, 1286-1298.
Mermelstein, P. (1978). On the relationship between vowel and consonant identification when cued by the same acoustic information. *Perception & Psychophysics*, **23**, 331-336.
Mills, C. B. (1980). Effects of context on reaction time to phonemes. *Journal of Verbal Learning & Verbal Behaviour*, **19**, 75-83.
Nettle, D. (1994). A behavioural correlate of phonological structure. *Language & Speech*, **37**, 425-429.
Otake, T., Hatano, G., Cutler, A., & Mehler, J. (1993). Mora or syllable? Speech segmentation in Japanese. *Journal of Memory & Language*, **32**, 358-378.
Otake, T., Hatano, G., & Yoneyama, K. (1996). Speech segmentation by Japanese listeners. In T. Otake & A. Cutler (Eds.), *Phonological structure and language processing: Cross-linguistic studies* (pp. 183-201). Berlin: Mouton de Gruyter.
Otake, T., Yoneyama, K., Cutler, A., & van der Lugt, A. (1996). The representation of Japanese moraic nasals. *Journal of the Acoustical Society of America*, **100**, 3831-3842.
Pallier, C., Bosch, L., & Sebastián-Gallés, N. (1997). A limit on behavioral plasticity in speech perception. *Cognition*, **64**, B9-B17.
Pallier, C., Dupoux, E., & Jeannin, X. (1997). EXPE: An expandable programming language for on-line psychological experiments. *Behavior Research Methods, Instruments, & Computers*, **29**, 322-327.
Sebastián-Gallés, N., Dupoux, E., Seguí, J., & Mehler, J. (1992). Contrasting syllabic effects in Catalan and Spanish. *Journal of Memory & Language*, **31**, 18-32.
Shand, M. A. (1976). Syllabic vs. segmental perception: On the inability to ignore "irrelevant" stimulus parameters. *Perception & Psychophysics*, **20**, 430-432.
Skelton, R. B. (1969). The pattern of Spanish vowel sounds. *International Review of Applied Linguistics in Language Teaching*, **7**, 231-237.
Steinheiser, F. H., & Burrows, D. J. (1973). Chronometric analysis of speech perception. *Perception & Psychophysics*, **13**, 426-430.
Stockwell, R. P., & Bowen, J. D. (1965). *The sounds of English and Spanish*. Chicago: University of Chicago Press.
Suomi, K., McQueen, J. M., & Cutler, A. (1997). Vowel harmony and speech segmentation in Finnish. *Journal of Memory & Language*, **36**, 422-444.
Swinney, D. A., & Prather, P. (1980). Phonemic identification in a phoneme monitoring experiment: The variable role of uncertainty about vowel contexts. *Perception & Psychophysics*, **27**, 104-110.
Tomiak, G. R., Mullennix, J. W., & Sawusch, J. R. (1987). Integral perception of phonemes: Evidence for a phonetic mode of perception. *Journal of the Acoustical Society of America*, **81**, 755-764.
van Ooijen, B. (1994). *Processing of vowels and consonants*. Doctoral dissertation, University of Leiden.
van Ooijen, B. (1996). Vowel mutability and lexical selection in English: Evidence from a word reconstruction task. *Memory & Cognition*, **24**, 573-583.
Whalen, D. H. (1989). Vowel and consonant judgments are not independent when cued by the same information. *Perception & Psychophysics*, **46**, 284-292.
Wood, C. C., & Day, R. S. (1975). Failure of selective attention to phonetic segments in consonant-vowel syllables. *Perception & Psychophysics*, **17**, 346-350.

NOTE

1. Like all students at the University of Barcelona, the subjects were also highly competent in Catalan. However, parallel research by Pallier, Bosch, and Sebastián-Gallés (1997) and Bosch, Costa, and Sebastián-Gallés (1997) shows that native Spanish-speaking subjects from this population without early exposure to Catalan have not developed accurate phonemic categories for Catalan. Even if they did have an effectively bilingual phonemic repertoire, however, that would operate against our experimental prediction, since the vowel repertoire, however, of Catalan is larger than that of Spanish.

(Manuscript received March 3, 1997;
revision accepted for publication August 30, 1997.)