JOURNAL OF MEMORY AND LANGUAGE **39,** 347–370 (1998)
ARTICLE NO. ML982571

Is Compensation for Coarticulation Mediated by the Lexicon?

Mark A. Pitt
Ohio State University

and

James M. McQueen
Max Planck Institute for Psycholinguistics, Nijmegen, The Netherlands

Ambiguous stops between /t/ and /k/ tend to be heard as /k/ after /s/-final words and as /t/ after /ʃ/-final words. Elman and McClelland (1988, *Journal of Memory and Language,* **27,** 143–165) obtained this compensation for coarticulation effect when the word-final fricatives were replaced with an ambiguous phoneme and argued that this indicated top-down lexical involvement in a prelexical process, as predicted by interactive models of speech perception. But autonomous models, which have no top-down processing but which are sensitive to the transitional probabilities between speech sounds, can also account for the effect. This study tested these two accounts. In Experiments 1, 2, and 3, listeners categorized fricatives at the ends of nonwords and words and the immediately following word-initial stops. When nonwords were the context stimuli, categorization of both the fricatives and the stops was influenced by the transitional probabilities of the vowels into the fricatives. When words were used, the transitional probabilities into the fricatives were matched. No compensation on following stops was found, even though the fricatives tended to be labeled in a lexically consistent manner. In Experiment 4, where listeners simply categorized stops at the ends of nonwords, further evidence of sensitivity to the transitional probabilities of these consonants was obtained. These results challenge interactive models, but are accounted for most parsimoniously by autonomous models in which transitional probabilities are represented independently of lexical information.

A perennial question in the study of spoken word recognition is whether or not higher-level sources of information can influence processing at lower levels, in a top-down, interactive way. This question has frequently been asked of the lexical and prelexical levels. That is, does lexical knowledge influence the prelexical processing involved in lexical access, as assumed in the interactive model Trace (McClelland & Elman, 1986; McClelland, 1991), or is information flow strictly bottom-up, with no top-down flow from the lexicon to prelexical processing, as assumed in autonomous models (such as the Race model, Cutler & Norris, 1979; Cutler, Mehler, Norris, & Segui, 1987; the Fuzzy Logical Model of Perception (FLMP), Oden and Massaro, 1978; the Shortlist model, Norris, 1994; and the Merge model, McQueen, Norris, & Cutler, submitted; Norris, McQueen, & Cutler, submitted)?

Lexical effects in the phonetic categorization task (Ganong, 1980) have fueled this debate. When listeners are asked to categorize tokens of a phonetic continuum, ranging, for example, from /d/ to /t/, and placed at the onsets of strings which could either form words or not (e.g., *dice–tice* and *dipe–type*), they are more likely to label ambiguous phonemes (i.e., those in the

---

The authors contributed equally to this project. This research was supported by Grant R29 DC01774 from the National Institute on Deafness and Other Communication Disorders, National Institute of Health, by the Human Frontiers Science Project, and by the Max Planck Society, with a visiting fellowship to MP. Some of this work was reported at the 134th meeting of the Acoustical Society of America, San Diego, California, December 1997. We thank Lisa Shoaf for diligently testing the subjects in all four experiments and Paul Luce and an anonymous reviewer for their comments on an earlier version of this paper.

Address correspondence and reprint requests to M. A. Pitt, Department of Psychology, Ohio State University, 1885 Neil Avenue, Columbus, Ohio OH 43210. E-mail: pitt.2@osu.edu, or J. M. McQueen, Max Planck Institute for Psycholinguistics, Wundtlaan 1, 6525 XD Nijmegen, The Netherlands. E-mail: james.mcqueen@mpi.nl.

0749-596X/98 $25.00

middle of the continuum) in a lexically consistent manner (e.g., more /d/ responses to the *dice–tice* continuum and more /t/ responses to the *dipe–type* continuum). There has been considerable discussion in the literature as to whether such lexical effects in phonetic categorization are more consistent with interactive models or with autonomous models (Burton, Baum, & Blumstein, 1989; Connine & Clifton, 1987; Fox, 1984; Massaro & Oden, 1995; McQueen, 1991a; Pitt, 1995; Pitt & Samuel, 1993).

Most of the results from phonetic categorization, as well as those from other tasks requiring phonetic decisions, such as phoneme monitoring (Cutler et al., 1987) and phonemic restoration (Samuel, 1981, 1996), can be explained by both classes of model. Both types of model assume that lexical knowledge can influence perceptual decisions, but they differ with respect to how this influence is exerted. In interactive models, lexical involvement in categorization follows from the claim about the architecture of the recognition system which defines such models: that the lexicon influences the operation of lower-level perceptual processes. In autonomous models, lexical involvement in categorization follows from the claim that perceptual decisions can be made postlexically: there is no need to postulate top-down connections to lower levels if output from the lexicon can directly influence perceptual decision-making.

One particular set of results from the phonetic categorization literature, on apparent lexical involvement in compensation for coarticulation, has been particularly important in the theoretical debate between interactive and autonomous models and is the focus of this study. Although the effect was originally taken to be powerful evidence in support of interactive models (Elman and McClelland, 1988), the present results challenge this view.

Mann and Repp (1981; see also Repp & Mann, 1981, 1982) showed that in fricative–stop sequences, where the stops varied along a place of articulation continuum, the fricative influenced the categorization of ambiguous stops. In the context of the alveolar fricative /s/, ambiguous steps on a stop continuum ranging from /t/ to /k/ were more often categorized as velar (/k/). But following the palatal fricative /ʃ/, the ambiguous stops were more often labeled as alveolar (/t/).

The explanation which Mann and Repp gave for this finding was that the perceptual system compensates for coarticulation. During production of a fricative–stop pair, the place of articulation of the stop shifts toward that of the fricative. Thus, following an /s/, the formant transitions of a velar stop, which provide one cue to place of articulation, will cue a place more anterior than in a neutral context. According to Mann and Repp (1981), the perceptual system compensates for this coarticulation, shifting the category boundary between alveolar and velar stops such that a more ‘‘alveolar’’-sounding (i.e., anterior) ambiguous stop will nonetheless be labeled as velar (i.e., /k/) in an /s/ context. Likewise, following an /ʃ/, the formant transitions in alveolar stops will cue a more posterior place of articulation than normal. Therefore, in an /ʃ/ context, perceptual compensation shifts the category boundary between /t/ and /k/ such that more posterior (velar) ambiguous stops will still be labeled as alveolar (i.e., /t/).

Elman and McClelland (1988) replicated the compensation effect in the labeling of word-initial stops when the stops followed fricative-final words such as *christmas* and *foolish.* There were more /k/ responses in the ambiguous region of the stop continuum (tapes–capes) following *christmas,* and more /t/ responses after *foolish.* The crucial manipulation was the replacement of the word-final fricatives with an ambiguous fricative (/?/) midway between /s/ and /ʃ/. Compensatory effects in the categorization of the stops were still observed. Importantly, these effects depended on the lexical status of the fricative-final string. When the /s/ in *christmas* was replaced with /?/, there were more /k/ responses in the ambiguous region of the continuum, as if listeners had heard an unambiguous /s/. When the /ʃ/ in *foolish* was replaced with the same ambiguous fricative, there were more /t/ responses to the ambiguous stops, as if listeners had heard an unambiguous /ʃ/.

Elman and McClelland argued that this was powerful evidence in favor of interactive models like Trace. They assumed that the coarticulatory compensation process has its locus at a prelexical level. The fricative information which triggers this process can be provided either by the signal (when the fricative is unambiguous) or by the lexicon, via top-down connections (when the fricative is ambiguous). The reason these data have been taken to be such clear evidence in favor of interactive models is that the lexicon appears to influence a process (i.e., compensation) that is assumed to operate at a prelexical level. These data seem problematic for autonomous models, since lexical involvement in a prelexical process appears to be direct evidence against the claim that information flow is strictly bottom-up. But there are, however, autonomous accounts of the effect.

Norris (1993) has shown how a recurrent network, in which information flow during recognition is strictly bottom-up, can simulate the compensation for coarticulation effect. The network learned sequential dependencies in featural input and stored this information in its recurrent connections. This knowledge was sufficient for the model both to learn the difference between /t/ after /s/ and /t/ after /ʃ/ (and similarly for /k/) and to learn that /s/ was more likely after one vowel (such as the /ə/ in *christmas*) and that /ʃ/ was more likely after another vowel (such as the /ɪ/ in *foolish*). Even in the case where the network had no lexical knowledge whatsoever, a compensation for coarticulation effect following an ambiguous fricative was observed.

But does the English vocabulary have the sequential dependencies necessary for compensation to occur after ambiguous fricatives in the contexts used by Elman and McClelland (1988)? Chater, Shillcock, Cairns, and Levy (submitted; see also, Cairns, Shillcock, Chater, & Levy, 1995, and Shillcock, Lindsey, Levy & Chater, 1992) have shown that this is the case. They trained a recurrent network on a sequence of two million phonemes from the London–Lund corpus of conversational English (Svartvik & Quirk, 1980). Their simulations showed that the compensation effect after an ambiguous fricative can occur in a bottom-up model with no lexical knowledge. In principle, therefore, any model in which the prelexical level is sensitive to sequential dependencies between segments can account for the compensation effect without the need to postulate top-down connections.

The simulations from both recurrent networks demonstrate that models with no lexical knowledge can simulate the Elman and McClelland (1988) data by learning the transitional probabilities (TPs) between speech sounds instead of by using lexical context. The reason these models can successfully simulate the compensation data is because lexical context was confounded with transitional probability in Elman and McClelland (1988). (The probabilities from the London–Lund corpus are provided in Cairns et al., 1995.) That is, /s/ is more likely than /ʃ/ after /ə/ in the full English vocabulary, and not only specifically at the end of the lexical entry beginning *christma.* Similarly, /ʃ/ is more likely than /s/ after /ɪ/, not only specifically at the end of *fooli.*

Because autonomous models with sensitivity to TPs can account for the compensation effect, we decided to test experimentally whether TPs can account for compensation in listeners' categorization performance. There is now considerable evidence to suggest that listeners are indeed sensitive to the statistical regularities in their language, and to sequential likelihoods between speech sounds in particular. For example, they play a role in the segmentation of continuous speech into words. Listeners learning an artificial language show sensitivity to the relative likelihoods of different sound sequences in that language (Saffran, Newport, & Aslin, 1996). Consonant clusters with a TP of zero (that is, phonotactically illegal clusters) provide listeners with a cue that can be used for lexical segmentation (McQueen, 1998). Even nine-month-old infants appear to be sensitive to the statistical properties of their native language (Jusczyk, Luce, & Charles-Luce, 1994).

There is also evidence to suggest that TP can influence phoneme identification. In phoneme monitoring, targets can be detected more easily when they occur in high-frequency syllables (i.e., syllables with high TPs) than when they occur in low-frequency syllables (i.e., those with low TPs; Pitt & Samuel, 1995). McQueen and Pitt (1996) found weak effects of TP when listeners were required to monitor for the final phoneme in consonant–vowel–consonant (CVC) nonwords. Stronger TP effects were observed when the target was the first consonant of the coda cluster in CVCC nonwords, where target detection was, overall, more difficult than in the CVCs. McQueen and Pitt (1996) argued that these results are likely to be most robust when target detection is difficult. If this prediction is correct, TP effects in phoneme identification should be strongest when the information specifying a phoneme is ambiguous, as in the phonetic categorization paradigm. Data from two studies that used this procedure can be interpreted as demonstrating TP effects. Massaro and Cohen (1983) and Pitt (in press) found that the labeling of an ambiguous phoneme between /l/ and /r/ was influenced by the phonotactic permissibility of the utterance. In /t?i/, there were more /r/ than /l/ responses because only /tri/ is legal in English. In /s?i/, just the opposite labeling bias was obtained (more /l/ than /r/ responses). Since the phonotactic legality of a sequence can be viewed as an extreme form of TP (impermissible clusters such as /tl/ and /sl/ have word-initial TPs of effectively zero), these results suggest that listeners are indeed sensitive to TP when they categorize ambiguous phonemes. However, less extreme TP differences between clusters did not yield a labeling bias (Pitt, in press).

EXPERIMENT 1

Although the preceding results are consistent with the claim that compensation for coarticulation is due to TP, there has not yet been a direct test of the influence of TP on the compensation effect. Are the apparent effects of lexical involvement in compensation in fact due to a lexical influence, to sensitivity to TP, or to both of these factors? There is a clear way to address this issue. Experiment 1 tested compensation for coarticulation following fricative-final nonwords, where the fricatives had TP biases, and following fricative-final words, where the TPs of the fricatives were matched.

If the effect is due to TP alone, it should be possible to observe compensation for coarticulation in the categorization of stop consonants when the stops are preceded by an ambiguous fricative, placed at the end of two nonwords, one in which the TPs make /s/ more probable and one in which /ʃ/ is more probable. If the effect is instead due to lexical involvement alone, no effect should be observed following these nonwords. But if the effect is purely lexical, it should be possible to observe the effect when stops are preceded by the same ambiguous fricative, placed at the end of two words, in which the TPs are controlled, such that /s/ and /ʃ/ are more or less equally likely in both words. Compensation should be observed following words balanced in TP and following nonwords with TP biases if both lexical and TP contexts are responsible for the effect. The present experiment tested these predictions.

Method

Participants. Sixteen Ohio State University undergraduates participated for pay or course credit. None reported hearing difficulties. The data from one listener were discarded because that listener responded only to variation in the fricative.

Design. Three context-related variables were crossed: biasing context, fricative bias, and fricative ambiguity. *Biasing context* refers to the two types of context, that with lexical bias and that with TP bias. *Fricative bias* refers to whether /s/ or /ʃ/ was most likely to occur based on the context preceding the fricative. One stimulus was biased toward /s/, another toward /ʃ/. Fricative bias was manipulated in both lexical and TP biasing contexts (using the words *juice,* /ʤus/, and *bush*, /bUʃ/, and the nonwords *ders,* /dɜ˞s/, and *naish,* /neɪʃ/, respectively). *Fricative ambiguity* refers to the manipulation of the fricative ending the context utterance. The fricative was perceptually unambiguous (a clear token of /s/ or /ʃ/) or perceptually ambiguous (/?/), halfway between /s/ and /ʃ/. Fricative ambiguity was manipulated to measure compensation with both unambiguous and ambiguous fricatives. Labeling in the latter case is of primary interest because it is here that frication was held constant across fricative bias (e.g., *jui?* vs *bu?*) and we were looking for differences to emerge

TABLE 1

Design and Context Materials for Experiments 1–3

		Fricative ambiguity		
		Unambiguous		Ambiguous
Biasing context	Fricative bias	/s/	/ʃ/	/?/
Lexical context	/s/-bias	juice	juish	jui?
	/ʃ/-bias	bus	bush	bu?
TP context	/s/-bias	ders/mees	dersh/meesh	der?/mee?
	/ʃ/-bias	nais	naish	nai?

Note. The /s/-bias TP context *der* was used in Experiments 1 and 2, and the context *mee* was used in Experiment 3. The other contexts remained the same in all three experiments.

as a function of the biasing context (lexical vs TP). This experimental design (summarized in Table 1) generates twelve contexts: /s/- and /ʃ/-biased contexts as determined by either lexical or TP biases, with both unambiguous and ambiguous fricatives. Following Elman and McClelland (1988), contexts preceded a word-initial /t/–/k/ continuum, *tapes–capes.*

Note that in four contexts the unambiguous fricatives were inconsistent with the contextual biases. For both the lexical and TP contexts, the unambiguous fricatives were swapped over contexts to create the nonwords *juish* (/ʤuʃ/) and *bus* (/bUs/, that is, not the word ‘‘bus,’’ /bʌs/), and the nonwords *dersh* (/dɝʃ/) and *nais* (/neɪs/).

Stimuli. A *tapes–capes* continuum was created using natural tokens of these words. The initial 128 ms (burst plus aspiration plus the first three (for /t/) or five (for /k/) pitch periods of the transition into the following vowel) were spliced from the words and digitally blended (by averaging samples; see McQueen, 1991a; Repp, 1981) in the following proportions to produce an eight-step /t/–/k/ continuum: .65/.35, .625/.375, .60/.40, .575/.425, .55/.45, .525/.475, .50/.50, .475/.525. The first value of each pair of numbers corresponds to the proportion of /t/, the second to the proportion of /k/. All steps were combined with the remaining portion of /eɪps/ (654 ms) from *capes* to yield the complete *tapes–capes* continuum.

Selection of the /t/–/k/ blending proportions was based on piloting, the goal of which was to create a continuum whose middle steps were sufficiently ambiguous to produce compensation in the ambiguous and unambiguous fricative contexts. To do this while at the same time maintaining a constant blending distance between steps required sacrificing the quality of the endpoint tokens. These items were somewhat ambiguous. (Elman and McClelland, 1988, p. 152, made a similar remark about their continua.) As a consequence, the labeling functions did not span the full range of the dependent measure across the stop continuum (from 0 to 1).

The /ʃ/-/s/ continuum was made using the parallel branch of the Klatt (1980) synthesizer. First, a 240-ms /ʃ/ was made. Next, /s/ was then created from /ʃ/ by altering the amplitudes of the formants. F3, F4, and F5 were decreased in amplitude; F6 was increased. The bandwidth of F3 was also reduced slightly for /s/. All other parameters were held constant. An eight-step continuum was created by interpolating between the parameter values of the two endpoint tokens in equal steps. The fricatives were then equated for perceived loudness. The synthesis parameters for the endpoints are listed in the Appendix.

Four additional utterances (two words and two nonwords) served as the biasing contexts that preceded the stop-initial word. The two words *juice* and *bush* have word frequencies of 11 and 14, respectively (Kucera & Francis, 1967). The words were chosen because neither /ʤuʃ/ nor /bUs/ are

TABLE 2

Transitional Probabilities[a]

			Probability conditional on	
Context	Vowel	Consonant	Vowel	Consonant
		Experiments 1—3		
Lexical contexts	/u/	/s/	$p(\text{us}\mid\text{u}) = 0.019$	$p(\text{us}\mid\text{s}) = 0.009$
(*juice/bush*)		/ʃ/	$p(\text{uʃ}\mid\text{u}) = 0.010$	$p(\text{uʃ}\mid\text{ʃ}) = 0.022$
	/U/	/s/	$p(\text{Us}\mid\text{U}) = 0.004$	$p(\text{Us}\mid\text{s}) = 0.000$
		/ʃ/	$p(\text{Uʃ}\mid\text{U}) = 0.004$	$p(\text{Uʃ}\mid\text{ʃ}) = 0.002$
TP-bias contexts	/ɝ/	/s/	$p(\text{ɝs}\mid\text{ɝ}) = 0.058$	$p(\text{ɝs}\mid\text{s}) = 0.028$
(*ders/naish/mees*)		/ʃ/	$p(\text{ɝʃ}\mid\text{ɝ}) = 0.007$	$p(\text{ɝʃ}\mid\text{ʃ}) = 0.019$
	/eɪ/	/s/	$p(\text{eɪs}\mid\text{eɪ}) = 0.064$	$p(\text{eɪs}\mid\text{s}) = 0.022$
		/ʃ/	$p(\text{eɪʃ}\mid\text{eɪ}) = 0.139$	$p(\text{eɪʃ}\mid\text{ʃ}) = 0.222$
	/i/	/s/	$p(\text{is}\mid\text{i}) = 0.021$	$p(\text{is}\mid\text{s}) = 0.014$
		/ʃ/	$p(\text{iʃ}\mid\text{i}) = 0.002$	$p(\text{iʃ}\mid\text{ʃ}) = 0.001$
		Experiment 4		
Stop contexts	/i/	/p/	$p(\text{ip}\mid\text{i}) = 0.020$	$p(\text{ip}\mid\text{p}) = 0.030$
(*yeep/chait*)		/t/	$p(\text{it}\mid\text{i}) = 0.025$	$p(\text{it}\mid\text{t}) = 0.011$
	/eɪ/	/p/	$p(\text{eɪp}\mid\text{eɪ}) = 0.015$	$p(\text{eɪp}\mid\text{p}) = 0.012$
		/t/	$p(\text{eɪt}\mid\text{eɪ}) = 0.151$	$p(\text{eɪt}\mid\text{t}) = 0.034$

[a] For each vowel, probabilities of the consonants following that vowel are given, both conditional on the probability of the vowel (e.g., $p(\text{us}\mid\text{u})$ for /u/ and /s/) and conditional on the probability of the consonant (e.g., $p(\text{us}\mid\text{s})$ for /u/ and /s/).

words and because the TP from the vowel into both fricatives is similar in the two cases. TPs were estimated using phoneme counts (weighted for word frequency) from an on-line American English dictionary consisting of words that most university undergraduates are likely to know. Simple diphone probabilities were computed since they appeared to require the fewest theoretical assumptions. They are not position-specific, nor do they entail assumptions about the coding of phonological structure (e.g., syllabic structure). Note, however, that diphone probabilities are necessarily confounded with syllable frequency: syllables which occur more often will contain diphone transitions which are more common. The transitional probability of a pair of phonemes was determined by summing the frequency of their joint occurrence and then dividing this value either by the total number of occurrences of the vowel (i.e., conditional on the vowel) or by the total number of occurrences of the fricative (i.e., conditional on the fricative). Words were selected for which the vowel-fricative transitions were as closely matched as possible on both counts. The diphone conditional probabilities for the vowels /u/ and /U/ into /s/ and /ʃ/ are given in Table 2. The use of *juice* and *bush* as the lexical contexts enabled us to assess the influence of lexical status on compensation without significant variation in the transitional probability between the vowel and fricative. The two nonwords *ders* and *naish* were created to maximize transitional probability from the vowel into one fricative but not the other. /ɝ/ was /s/-biased and /eɪ/ was /ʃ/-biased, conditional both on the fricative and the vowel. The conditional probabilities for these vowels are also given in Table 2. The use of these nonwords as the TP contexts enabled us to observe the influence of TPs on compensation independently of the influence of specific lexical entries.

The context utterances were combinations of natural and synthetic speech (Mann & Repp, 1980; Repp & Mann, 1981). Tokens of the strings *juice, bush, ders,* and *naish* were recorded. The final fricative was removed from each token at frication onset. Durations of the four initial portions were 247, 193, 186, and 267 ms, respectively. The two endpoint fricatives (/s/ and /ʃ/) from the synthetic continuum plus a step from the middle of the continuum (/?/) were then appended to three copies of each of the four natural strings. The two copies with the endpoints (/s/ and /ʃ/) served as the context items in the unambiguous fricative conditions. The copy with /?/ served as the context in the ambiguous fricative condition.

The ambiguous fricative was chosen on a by-participant basis to ensure that the fricative would be maximally ambiguous. Participants were pretested on the fricative continuum, labeling each step 12 times. The fricatives were not presented in isolation, but appended to a natural token of /a/ to make the frication phonetically meaningful. The step whose /ʃ/-labeling proportion was closest to .5 was used as the ambiguous fricative. This was always step 4 or 5 on the eight-step continuum.

The 12 context utterances (4 initial portions × 3 fricatives—two unambiguous, one ambiguous) were then prepended to each step of the *tapes–capes* continuum (separated by 20 ms of silence) to form the entire set of stimuli that were presented to listeners. In the four ambiguous fricative conditions (32 stimuli), each stimulus was presented 16 times, for a total of 512 trials. In the eight unambiguous fricative conditions (64 stimuli), each was presented eight times (512 trials).

Procedure. Listeners were tested individually in a sound-attenuated cubicle. After the pretest to select the ambiguous fricative, listeners were instructed on participation in the main experiment. They were told to listen closely to the pair of utterances presented on each trial and classify the last consonant of the first utterance as either /s/ or /ʃ/ and the first consonant of the second utterance as either /t/ or /k/. Use of this two-response procedure (Mann & Repp, 1981) provides a direct measure of the effect of context on the identification of the fricative as well as a measure of any compensatory effect that the context (lexical or TP) might have on categorization of the following stops. Indeed, pilot work showed that larger compensation effects were obtained when listeners were required to classify the fricatives as well as the stops. Once the listener made the classification decisions, one of four response buttons had to be pressed, each labeled with a different combination of the four consonants (i.e., sk, st, shk, sht). Fast responding was emphasized.

Stimuli were presented in blocks of 128 trials, pseudorandomly ordered. All 64 unambiguous stimuli were presented once per block. Each of the 32 ambiguous stimuli were presented twice. There was a 1.5-s pause between trials and a 4-s window in which listeners had to respond after stimulus offset. Four blocks were presented in each of two 45-min testing sessions, which were separated by one or two days. Rest breaks were provided after every block. Sixteen practice trials preceded the first block in each session.

Equipment. Stimuli were recorded by MP onto audio tape and then digitized onto magnetic disk at 10 kHz sampling rate (low-pass filtered at 4.8 kHz), where they were edited and combined for playback. Synthesis was also carried out at 10 kHz. Stimuli were presented binaurally over headphones at a comfortable listening level. Stimulus presentation and response collection were controlled by a microcomputer.

Results

Stop labeling. The stop data were divided and analyzed in two parts as a function of whether the preceding fricative was unambiguous or ambiguous. The proportion of /t/ responses was calculated for each listener across all steps of the continuum in each condition. The averaged data in the eight unambiguous conditions are plotted in the two graphs at the top of Fig. 1. The graph on the left contains data from the lexical biasing context, that on the right from the TP context. Within each graph, the data in each fricative bias condition (e.g., *jui, bu*) are plotted as a function of whether the unambiguous fricative was /s/ (solid lines) or /ʃ/ (dashed lines).

Recall that compensation for coarticulation causes listeners to report alveolar stops (i.e., /t/) following /ʃ/ and velar stops (i.e., /k/) following /s/. This effect shows up as a separation of the labeling functions in the middle of the /t/–/k/ continuum, where the stop is perceptually am-

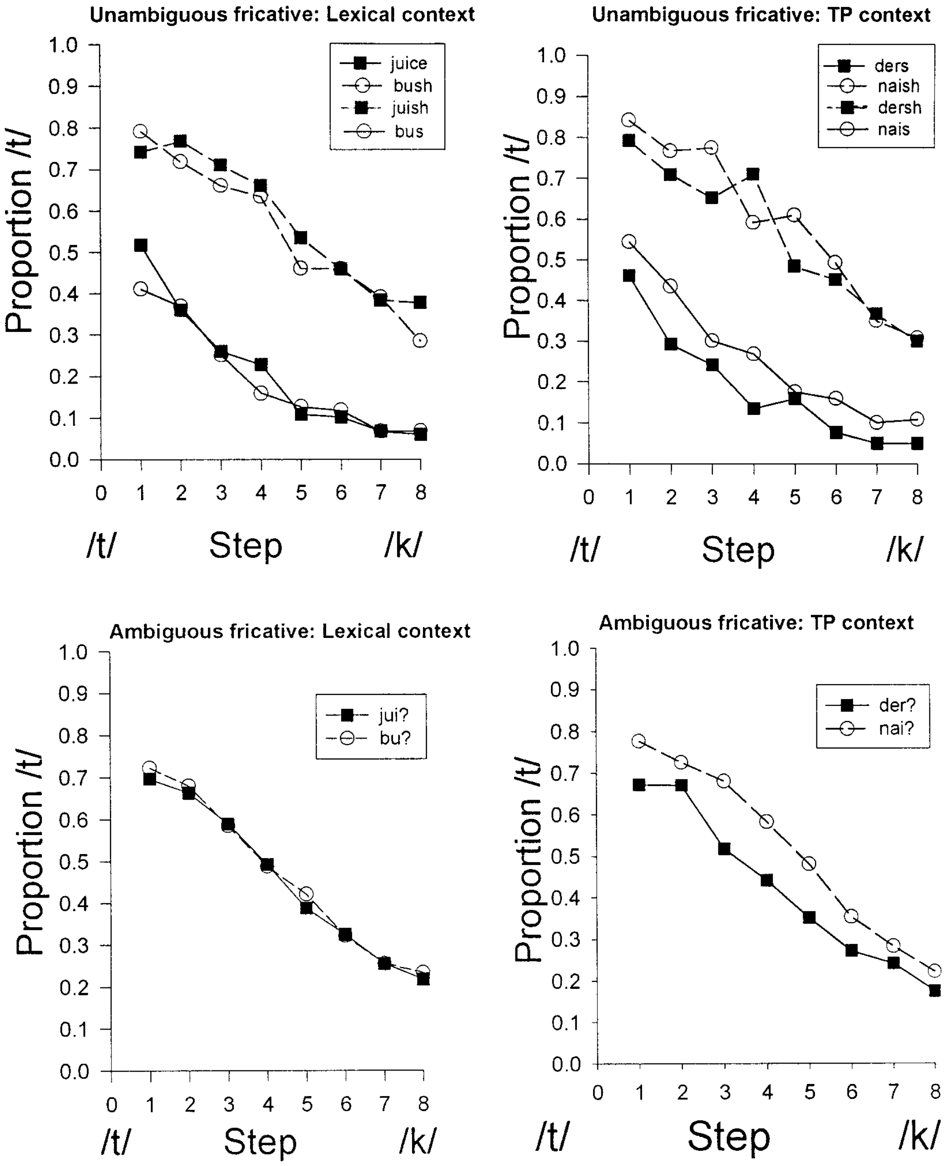

FIG. 1. Experiment 1. Overall proportion of /t/ responses to each step of the eight-step /t/–/k/ (*tapes–capes*) continuum. The upper panel shows categorization following unambiguous fricatives in both lexical- and TP-bias contexts. The dashed lines show labeling after /ʃ/, and the solid lines show labeling after /s/. Filled squares show contexts with an /s/ bias (*jui* and *der*); open circles show contexts with an /ʃ/ bias (*bu* and *nai*). The lower panel shows categorization following ambiguous fricatives (/?/) in both lexical- and TP-bias contexts. Filled squares and solid lines show contexts with an /s/ bias (*jui* and *der*); open circles and dashed lines show contexts with an /ʃ/ bias (*bu* and *nai*).

biguous and, therefore, contextual (i.e., fricative) influences on labeling are likely to be largest. The results in the two biasing conditions are similar and clear cut: Fricative identity (/s/ or /ʃ/) had a large influence on labeling whereas fricative bias (/s/- or /ʃ/-bias in both lexical and TP contexts) had none. Large compensation effects are evident in both graphs, with stops preceded by /ʃ/ being classified as /t/ far more often than stops preceded by /s/. That the two /s/ and two /ʃ/ functions within each graph overlap indicates that the context within which the fricative was embedded (e.g., /ʤu/ or /bU/) had no effect on stop labeling.

These observations were supported by the results of two-way ANOVAs (fricative bias × fricative identity) performed separately on the data in each biasing context. Because compensation was found over a range of steps in the middle of the continuum, compensation magnitude was measured by calculating the area between pairs of functions from step 3 through step 6 (see Pitt & Samuel, 1993, for additional discussion of this procedure). The data used in the analyses were listeners' average proportion /t/ responses across steps 3–6 in each condition. The only reliable outcomes were the main compensation effects (Lexical context: $F[1,14] = 36.71$, $p < .001$; TP context: $F[1,14] = 33.92$, $p < .001$). Compensation magnitude was virtually identical in both biasing contexts: .40 units (i.e., units of area between the curves) in the lexical context and .41 units in the TP context.

The results with the unambiguous fricatives thus replicate what others have found (Elman & McClelland, 1988; McQueen, 1991b; Repp & Mann, 1980; Mann & Repp, 1981). Robust compensation was obtained along with no influence of fricative bias. Neither lexical nor TP contexts influenced compensation when the fricative was a clear token of its category.

This outcome did not hold when the fricative was ambiguous. The results in this condition are of primary interest because they enable us to assess the independent contributions of lexicality and transitional probability on compensation. The labeling data are in the graphs in the lower part of Fig. 1. The results are again clear cut: Lexical bias produced no compensation (the left graph), but TP did (the right graph). In the lexical context, fricative bias had absolutely no influence on fricative labeling. The nearly complete overlap of the /s/-bias and /ʃ/-bias functions across the /t/–/k/ continuum is striking. The magnitude of compensation in this case was .005 units ($F < 1$), due solely to the difference at step 5.

TABLE 3

Mean Proportion /s/ Classification of the Unambiguous and Ambiguous Fricatives in the Four Context Utterances[a]

Context	Unambiguous fricative (/s/ and /ʃ/ combined)	Ambiguous fricative (/?/)
Experiment 1		
jui	.53	.61
bu	.46	.22
der	.48	.34
nai	.47	.18
Experiment 2		
jui	.63	.68
bu	.50	.27
der	.53	.49
nai	.47	.12
Experiment 3		
jui	.42	.36
bu	.49	.19
mee	.53	.46
nai	.50	.29

[a] In the unambiguous condition, /s/ and /ʃ/ followed the context equally often; a value of .50 represents unbiased classification.

In contrast, a sizeable compensation effect is evident across the middle of the continuum when TP served as the biasing context. Compensation magnitude was .13 units, which was reliable, $F(1,14) = 6.54$, $p < .05$.

Fricative labeling. One concern with the stop labeling data might be that the TP and lexical biasing manipulations were not equally effective. Although the TP contexts were sufficiently biased to induce a compensation effect, the lexical contexts might not have been. That is, perhaps the lexical items (i.e., the initial CVs) had no effect on fricative perception. Analysis of the fricative identification data shows that this concern is unfounded. To examine biases in fricative labeling, mean proportion /s/ responses in the two fricative ambiguity conditions were calculated for each subject. These values were then averaged over listeners and are listed at the top of Table 3. A value of .5 (an equal number of /s/ and /ʃ/ responses) indicates no bias in labeling. This was the case for the unambiguous fricatives. Proportions in all four contexts hover around .50. A two-way ANOVA on the unambiguous fricative responses with biasing context (lexical, TP) and fricative bias (/s/, /ʃ/) as variables yielded no statistically reliable differences. Note that the averages near .50 reflect accurate responding: overall, listeners almost always labeled /s/ as /s/ and /ʃ/ as /ʃ/.

The effects of these variables were quite different in the ambiguous fricative condition. Overall, there was a tendency to report the fricative as /ʃ/, except for *juice,* which showed a slight bias toward /s/. In contrast to the unambiguous condition, fricative bias yielded robust effects in the two biasing contexts. In the lexical context, /s/ responses dropped by .39 from the /s/-bias (*juice*) to the /ʃ/-bias (*bush*) context. The drop between the two corresponding items in the TP context, although still large, was less than half of this value (.16). A two-way ANOVA produced reliable main effects of biasing context (more /s/ responses, overall, in the lexical than in the TP bias condition; $F(1,14) = 13.89$, $p < .01$) and of fricative bias (more /s/ responses in the /s/-biased than in the /ʃ/-biased contexts; $F(1,14) = 27.68$, $p < .001$). The interaction of these variables was significant ($F(1,14) = 5.77$, $p < .05$), showing that the effect of lexical context was larger than that of TP context. Further tests showed that these two effects were reliable in themselves (lexical: $F(1,14) = 20.97$, $p < .001$; TP: $F(1,14) = 10.04$, $p < .01$).

These data clearly demonstrate that the lexical and TP contexts were both powerful enough to influence listeners' labeling of the ambiguous fricative. These results thus contrast with the stop data, where the TP contexts influenced identification but the lexical contexts did not.[1]

Discussion

The present results suggest that TP-biased contexts can cause compensation for coarticulation following ambiguous fricatives. When transitional probability was manipulated while lexical context was held constant, a sizeable compensation effect (.13 units) was found. Because the context utterances in this case were nonwords (*ders* and *naish*), it is highly unlikely that the effect originated from specific lexical entries. The current findings also suggest that lexical influences are not necessary to observe perceptual compensation following an ambiguous fricative. When transitional probability between the vowel and fricative was held constant on words, compensation was not found. It should have been if the effects were due solely to lexical feedback from specific lexical entries.

Note that the effects we have observed are due to the influence of the fricatives on identification of the stops and not vice versa. Mann and Repp (1981) did find effects of stop identification on fricative judgments, but they were much smaller than the effect of fricatives on stop judgments. In the present experiment, an effect of stop judgments on fricative judgments is ruled out, since on this account, one would expect the effects in the lexical and TP bias conditions to be equivalent. With the ambiguous fricative in the lexical context there was no shift in stop identification, but there was nonetheless an effect on the fricative judgments. The effects on fricative judgments are thus due to the contexts preceding the fricatives.

Although interpretation of a null effect is

---

[1] A concern with Experiment 1 might be that the endpoints of the *tapes–capes* continuum were not fully unambiguous tokens of /t/ and /k/ and thus that listeners' responses were not fully anchored. This concern was addressed by rerunning Experiment 1, but using a more widely spaced eight-step continuum, including endpoints which were identified almost always as either /t/ or /k/. This continuum necessarily contained a narrower ambiguous region. The results nevertheless patterned like those of Experiment 1. The compensation effect following /?/ was again present in the TP context, but, as one would expect when there are fewer ambiguous stops, it was reduced in magnitude (.065 units, as opposed to .13 units in Experiment 1) and was not significant ($F(1,9) = 1.82$, $p < .21$). There was again no indication of a compensation effect in the lexical context. In fricative identification, however, context effects were actually larger than in Experiment 1 (lexical: 0.60, $F(1,9) = 57.89$, $p < .001$; TP: 0.27, $F(1,9) = 8.01$, $p < .02$). The difference between these two effects was also significant ($F(1,9) = 5.72$, $p < .05$). These results show that the effects of lexical and TP bias on fricative identification are reliable. The effect of TP bias on stop identification with unambiguous endpoints was present, but, because of the limited range of ambiguous phonemes, was not significant. In subsequent experiments we therefore resorted to the use of a *tapes–capes* continuum with a wider ambiguous region (and less extreme endpoints).

complicated by the fact that its cause is difficult to pinpoint, the failure to find a compensation effect in the lexical conditions is not due to a weak manipulation of lexical context itself, which would have stacked the deck against finding compensation. This explanation might at first glance seem plausible given that the words were only three phonemes long, providing little context (or time) for biasing influences on fricative perception to emerge. However, the fricative classification data suggest just the opposite (Table 3). Classification of the ambiguous fricative differed by .39 between the two lexical contexts (*juice* and *bush*). It is this large shift in fricative labeling that convinces us that lexical influences are not necessary to observe compensation following ambiguous fricatives.

EXPERIMENT 2

The initial CV contexts used in Experiment 1 were taken from natural utterances. Thus, the *jui* used to make *juis, juish* and *jui?* came from a recording of *juice.* Likewise, *bu* came from *bush, nai* from *naish,* and *der* from *ders.* It is thus possible that the vowels in these CVs contained acoustic cues to the following fricative (such as, for example, formant transitions specifying place of articulation of the following fricative). Any of the apparent effects of lexical or TP bias which were observed could thus potentially be due to this confound, because in all four cases any acoustic cues would be likely to cue the contextually appropriate fricative. Note that although this possibility undermines the context effects observed in fricative labeling, it does not provide an account of the asymmetry in the stop labeling data between the lexical and TP context conditions. If the fricative biases were due solely to the acoustic confound, effects on subsequent stop identification should have been observed in both context conditions, with potentially larger effects on stop identification in the lexical context condition because the effect on fricative identification was largest. The acoustic cues in the initial vowels could nevertheless have contributed to the pattern of results observed in Experiment 1.

Experiment 2 was designed to deal with this concern. The initial CV portions were replaced with neutral tokens, in which there were no cues to a subsequent fricative. Natural CV tokens for each of the four contexts were recorded in isolation and spliced onto the fricatives and *tapes–capes* tokens used in Experiment 1. Other than this modification, Experiment 2 was a direct replication of Experiment 1.

Method

Participants. Twelve new undergraduates served in the experiment.

Stimuli. New tokens of the four context CVs (*jui, bu, der,* and *nai*) were recorded by MP without the final fricative. These items were then adjusted for amplitude and duration so that when combined with the synthetic fricative, the context utterances sounded as though they were spoken naturally. CV durations were as follows: *jui,* 264 ms; *bu,* 247 ms; *der,* 178 ms; *nai,* 317 ms. These CVs were on average 30 ms longer than those in Experiment 1. This slightly slower rate of speech also required lengthening the closure duration between the fricative and stop (from 20 to 40 ms) to make the two-word utterances sound natural. The stimuli were identical to those used in Experiment 1 in every other way.

Procedure. The experimental setup and the testing procedure were the same as in Experiment 1.

Results

Stop labeling. The /t/–/k/ labeling data are shown in Fig. 2 in the same manner as those in Fig. 1. Their similarity to those in Fig. 1 indicates that the compensation results found in the preceding experiment are not due to fricative cues being present in the original CV context utterances. Neutral CVs produced the same pattern of results.

When the fricatives were unambiguous, large compensation effects were obtained in both biasing contexts. Not only were biases similar in size across the lexical and TP contexts, but they were comparable to what was found in the Experiment 1. The effects in both contexts were statistically reliable, lexical context: .38 units, $F(1,11) = 80.82$, $p < .001$; TP context: .35

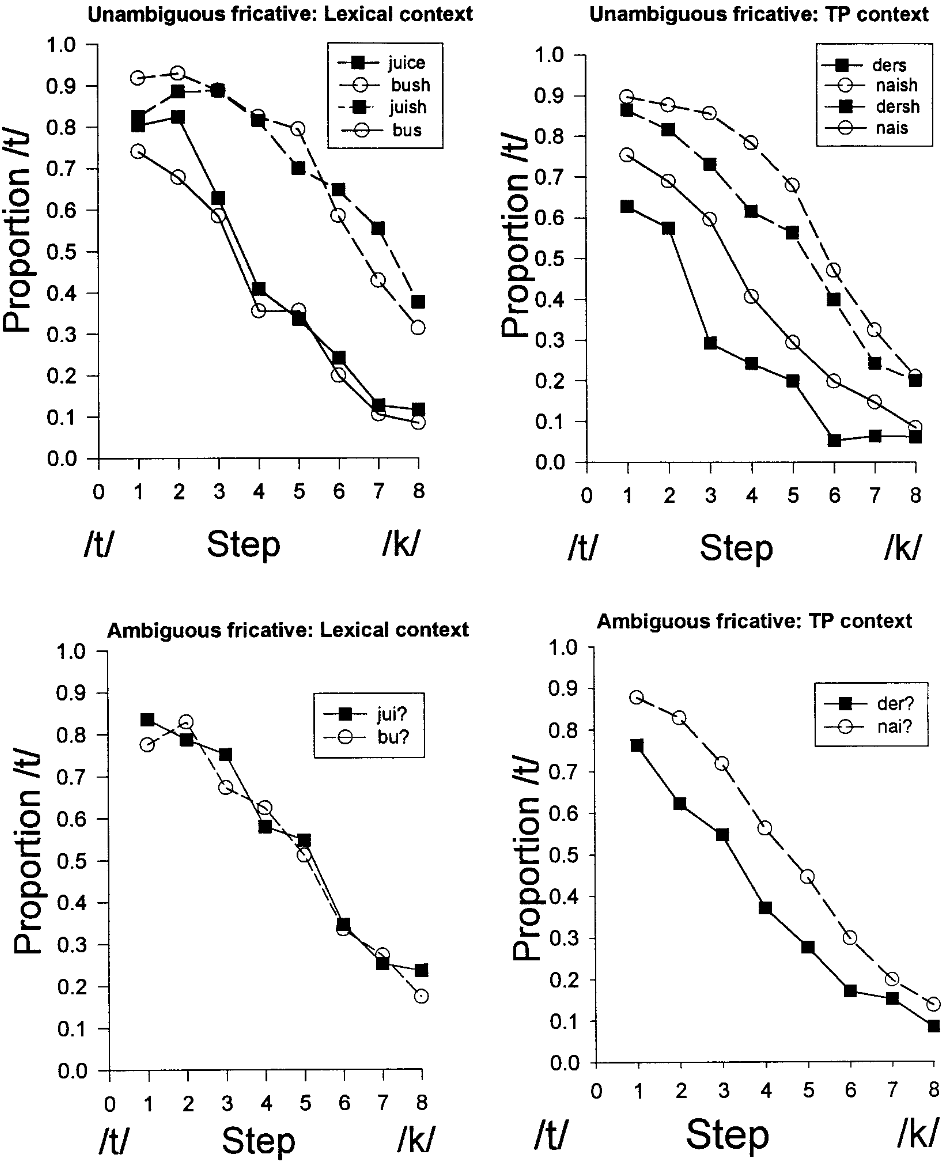

FIG. 2. Experiment 2. Overall proportion of /t/ responses to each step of the eight-step /t/–/k/ (*tapes–capes*) continuum. The upper panel shows categorization following unambiguous fricatives in both lexical- and TP-bias contexts. The dashed lines show labeling after /ʃ/, and the solid lines show labeling after /s/. Filled squares show contexts with an /s/ bias (*jui* and *der*); open circles show contexts with an /ʃ/ bias (*bu* and *nai*). The lower panel shows categorization following ambiguous fricatives (/?/) in both lexical- and TP-bias contexts. Filled squares and solid lines show contexts with an /s/ bias (*jui* and *der*); open circles and dashed lines show contexts with an /ʃ/ bias (*bu* and *nai*). The data-point for step 6 on the *der?* function is based on the three subjects who heard that token.

units, $F(1,11) = 54.73$, $p < .001$. Labeling in the TP context differed from the preceding experiment in that the *nais* and *naish* functions are shifted above their *der* counterparts (*ders* and *dersh*). This contextual biasing effect was reliable, $F(1,11) = 13.01$, $p < .004$ (.15 units) and suggests that the TP context affected stop labeling even in the presence of unambiguous fricatives. This effect is attributable to the unambiguous fricatives being somewhat ambiguous following the CV contexts in the present experiment (i.e., those with no formant transitions), that is, more ambiguous than following the more natural CVs used in Experiment 1. The shift between the *nai* and *der* functions following unambiguous fricatives thus appears to reflect an effect of TP bias on what were, in fact, not completely unambiguous fricatives. Support for this contention can be found in the fricative labeling data (discussed below).

Labeling in the two ambiguous conditions closely resembles what was found in Experiment 1: No effect of lexical bias was found, but there was an effect of TP bias. Due to an error in the experimental program, only three of the twelve subjects heard *der?* with stop 6 on the /t/–/k/ continuum. The analyses in the ambiguous fricative conditions was therefore based on the area between the curves for steps 3–5 (instead of from steps 3–6, as in all the other analyses). In the lexical context, the functions are intertwined across the continuum, producing a nonreliable shift ($-.02$ units), $F < 1$. The TP effect (.18 units) was significant: $F(1,11) = 5.99$, $p < .03$. The size of the TP bias was slightly larger than that found in Experiment 1 (.13 units).

Fricative labeling. The fricative labeling data are listed in the middle of Table 3. The results are in line with what was found in Experiment 1. When the fricative was unambiguous, labeling was near .50 in all contexts except *jui,* which showed a /s/-bias. Statistical analyses showed that the labeling difference (.13) between the two lexical contexts was reliable, $F(1,11) = 9.66$, $p < .01$, and between the two TP contexts (.06) was not. But note that this effect was larger than that in Experiment 1 (.01), consistent with the view that the unambiguous fricatives in the present contexts were less unambiguous than when they appeared in the contexts used in Experiment 1.

In the ambiguous fricative condition, the biasing context had a strong influence on labeling, with a large drop in /s/ responses after *bu* and *nai* compared with *jui* and *der.* The effects were of similar magnitude in the lexical and TP contexts (.41 and .37, respectively; lexical: $F(1,11) = 31.44$, $p < .001$; TP: $F(1,11) = 15.54$, $p < .002$). These effects did not differ reliably from each other.

Discussion

The results of Experiment 2 reinforce the two primary conclusions from Experiment 1. First, it appears that the transitional probability between adjacent phonemes can be responsible for compensation for coarticulation. Second, it appears that the lexicon, even when it is involved in fricative decisions, does not influence compensation for coarticulation in perception of the following stop. Note that the same pattern was observed in both the ambiguous and unambiguous fricative conditions in Experiment 2. In the TP conditions, there were shifts in stop labeling consistent with the TP bias both after the ambiguous and unambiguous fricatives; the unambiguous fricatives were clearly not completely unambiguous. But in the lexical conditions, there were no shifts in stop labeling due to lexical bias with either the ambiguous or unambiguous fricatives, in spite of the fact that in both cases the contextual involvement in fricative identification was larger than in the TP conditions.

EXPERIMENT 3

The basis of the compensation effects observed in Experiments 1 and 2 is that stop identification is influenced by fricative information preceding the ambiguous stops. It is also the case, however, that fricative identification can in turn be influenced by vocalic information preceding ambiguous fricatives. Experiment 3 addressed this latter effect, specifically, that fricative identification varies according to whether a neighboring vowel is rounded or unrounded. There tend to be more /s/ responses to an ambiguous fricative between /s/ and /ʃ/ in the context of a rounded vowel than in the context of an unrounded vowel. Although in most studies this effect has been found when the fricative preceded the vowel (Mann & Repp, 1980; Whalen, 1981, 1989), the effect has also been observed when the vowel preceded the fricative (Soli & Mann, 1982). Both these effects have been interpreted like that of fricative–stop compensation: the perceptual system acts to compensate for the acoustic consequences of fricative articulation in rounded versus unrounded contexts.

In the present situation, vowel rounding is not a concern in the lexical context conditions: both /u/ (in *juice*) and /U/ (in *bush*) are rounded vowels. But there is a potential confound in the TP context conditions. The vowel with the /s/ TP bias (/ɝ/) is a rounded vowel, while the vowel with the /ʃ/ TP bias (/eɪ/) is unrounded. The tendency to label /?/ as /s/ after /ɝ/ but as /ʃ/ after /eɪ/, and the subsequent compensation for coarticulation effect on the stops, could thus be due to the status of these vowels as rounded and unrounded, rather than because of TP biases.

This confound was removed in Experiment 3 by replacing /dɝ/ with /mi/. The vowel /i/, like /ɝ/, has a TP bias toward /s/ (see Table 2). But unlike /ɝ/, /i/ is unrounded. As shown in Table 1, the contexts *mees, mee?,* and *meesh* replaced *ders, der?,* and *dersh;* otherwise, the materials and design of Experiment 3 were identical to those of Experiments 1 and 2.

Method

Participants. Twelve new students from the same population as Experiment 1 served as participants.

Design. The experimental design was the same as Experiment 1.

Stimuli. The syllable /mi/ (253 ms) was recorded and replaced all occurrences of the syllable /dɝ/ in the stimuli from Experiment 2.

Procedure. The only procedural changes involved the fricative labeling experiment and the selection of the most perceptually ambiguous step for use in the ambiguous fricative condition. A pilot experiment in which fricative labeling was measured following each of the four vowels showed that the most ambiguous token tended to be about one step earlier in the continuum (i.e., fewer /s/ responses) for the two vowels in the two words (/u/ and /U/) than in the two nonwords (/i/ and /eɪ/). There were, however, no differences between the most ambiguous tokens within each pair. This was probably the effect of vowel rounding on fricative identification, since the vowels in the lexical context were both rounded and those in the TP context were both unrounded. To ensure that the most ambiguous fricative was being used in each biasing context in the main experiment, the fricative labeling pretest was modified so that if need be, a different fricative could be selected for the lexical and TP contexts. Instead of categorizing the fricatives only in the context of the vowel /a/, listeners categorized steps on the fricative continuum with two preceding vowels, one for lexical contexts (/u/) and one for TP contexts (/i/). The most ambiguous step with each vowel was used in the lexical and TP contexts, respectively. For eight listeners, step 4 was used in the lexical context and step 5 in the TP context. For the remaining four listeners, the same step was used in both contexts (three with step 4 and one with step 5).

Results and Discussion

Stop labeling. The /t/–/k/ data are displayed in Fig. 3. Inspection of the four graphs shows that the data tell the same story as those in Experiments 1 and 2. When the fricatives were unambiguous, robust compensation was found. With the lexical context, the effect size (again using the area-between-the-curves measure) was .25 units, which was statistically significant, $F(1,11) = 24.62$, $p < .001$. In the TP context, the effect was also .25 units, which was also significant, $F(1,11) = 15.97$, $p < .002$.

When the fricative was ambiguous, no lexical compensation effect was obtained. The /s/-bias and /ʃ/-bias functions overlap closely, particularly in the middle of the continuum (effect size = .01 units, $F < 1$). The magnitude of compensation in the TP context was .11, $F(1,11) = 9.04$, $p < .01$. Comparison of the two effect sizes across lexical and TP contexts showed that they were reliably different, $F(1,11) = 4.78$, $p < .051$.

Fricative labeling. The fricative labeling data are shown at the bottom of Table 3. As in Experiments 1 and 2, contextual influences were larger in the ambiguous than the unambiguous conditions. Although there was virtually no bias in classifying the unambiguous fricatives as /s/ in the TP context, a small, though statistically nonsignificant, effect

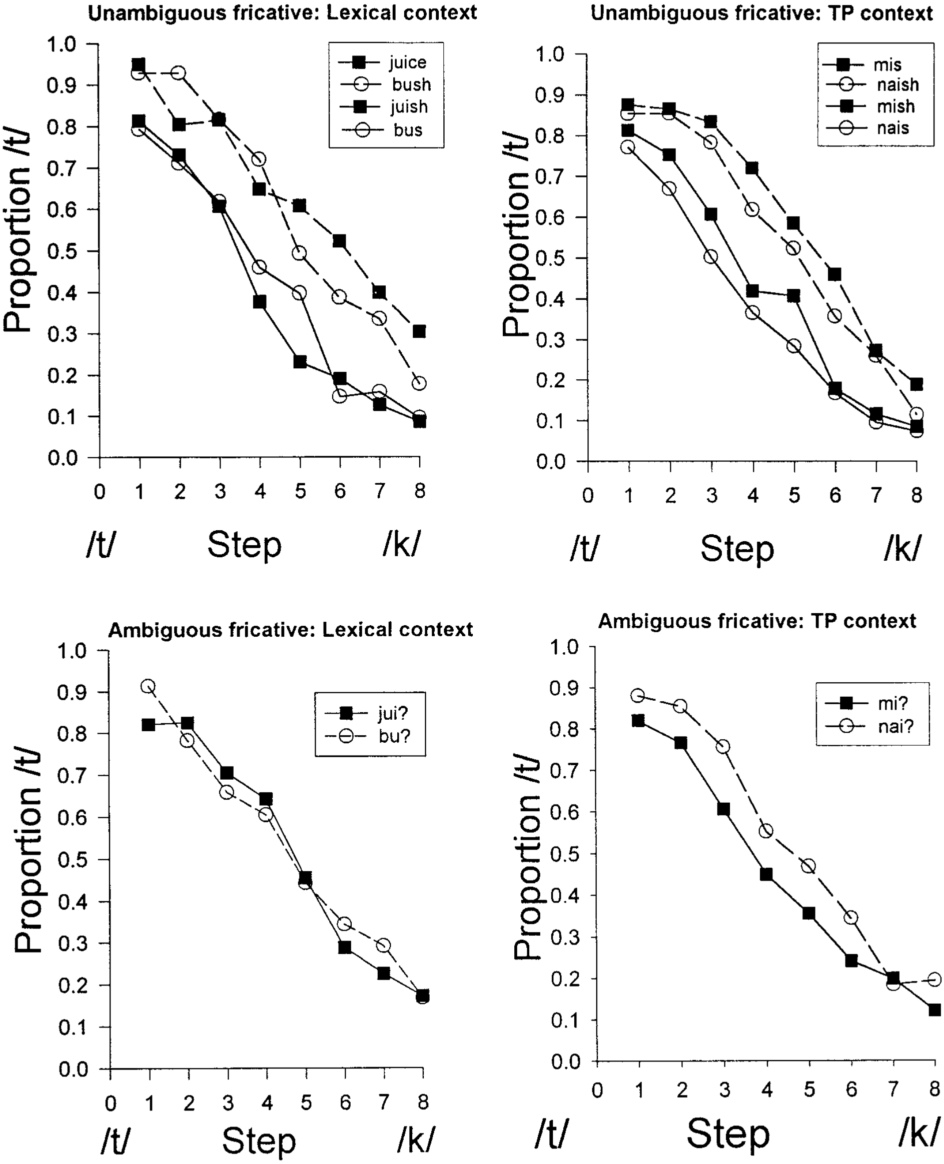

FIG. 3. Experiment 3. Overall proportion of /t/ responses to each step of the eight-step /t/–/k/ (*tapes–capes*) continuum. The upper panel shows categorization following unambiguous fricatives in both lexical- and TP-bias contexts. The dashed lines show labeling after /ʃ/, and the solid lines show labeling after /s/. Filled squares show contexts with an /s/ bias (*jui* and *mee*); open circles show contexts with an /ʃ/ bias (*bu* and *nai*). The lower panel shows categorization following ambiguous fricatives (/?/) in both lexical- and TP-bias contexts. Filled squares and solid lines show contexts with an /s/ bias (*jui* and *mee*); open circles and dashed lines show contexts with an /ʃ/ bias (*bu* and *nai*).

(.07) was found in the lexical context. In the ambiguous condition, context effects were identical in magnitude, .17. Statistical analyses showed that the difference between the /s/ and /ʃ/ biasing contexts was reliable, $F(1,11) = 7.57$, $p < .02$.

The same pattern of results was therefore obtained when the confound of vowel rounding was removed in the TP contexts. The present data affirm the results of Experiments 1 and 2 in suggesting that transitional probabilities between phonemes can cause compensation for coarticulation. The two contexts (lexical and TP) affected fricative labeling, but only in the TP context was any compensation in stop labeling observed. No compensation was found in the lexical context.

EXPERIMENT 4

In the first three experiments, we have argued that the bias in fricative labeling found in the TP context (Table 3) is due to the transitional probability between the vowel and fricative. In this final experiment, we undertook a more thorough test of this claim. We did not examine compensation per se, but rather assessed the reliability and generality of the novel finding that transitional probability between utterance-final VCs can bias listeners' labeling of the consonants. Instead of using an /s/–/ʃ/ continuum, as in Experiments 1–3, we assessed the generalizability of the TP effect to other consonants in the context of a standard phonetic categorization experiment. Listeners were asked to categorize the stops /p/ and /t/ in the continua *chaip–chait* and *yeep–yeet.* /p/ is more frequent after *yee,* and /t/ is more frequent after *chai.* If listeners are indeed sensitive to the frequency with which phonemes co-occur in English, then just as is found in lexical contexts (e.g., *kiss–kish* and *fis–fish,* McQueen, 1991a), the labeling functions should split apart in the ambiguous region and converge at the endpoints. Specifically, /p/ responses should be more frequent on the *yeep–yeet* function than on the *chaip–chait* function.

Method

Participants. Fourteen students in an introductory psychology course participated.

Stimuli. In order to select the contexts for the /p/–/t/ continuum, the transitional probabilities between vowels and these stop consonants were again computed using the on-line dictionary. When conditional on the stop, /t/ was more likely than /p/ after /eɪ/, and /p/ was more likely than /t/ after /i/. When conditional on the vowel, /t/ was also more likely than /p/ after /eɪ/, but /t/ was also slightly more likely than /p/ after /i/. Although this small reversal, which was not found in a larger dictionary (CELEX, Burnage, 1991), has the potential of diluting any TP effect, it also makes the experiment a relatively conservative test of the hypothesis that listeners use TP in phonetic categorization.

The stop–final continua were created by appending synthetic stops onto natural tokens of /ʧeɪ/ and /ji/. Recordings of the two CV context items in isolation (i.e., without a final consonant) were made. *Chai* and *yee* were, respectively, 327 and 281 ms in duration. The stop consonant was an aspirated burst. Steps along the continuum varied in aspiration (longer for /t/) and the center frequency and amplitude of Formants 1–3 (higher on both dimensions for /t/). The continuum was synthesized following the procedure used to make the fricative continuum in Experiment 1. /p/ and /t/ endpoint tokens were made, and the six intervening steps were created by interpolating between the parameter values of the endpoint tokens in approximately equal-size steps. The Klatt (1980) parameter values for the endpoint tokens are listed in the Appendix. A stop closure interval (70 ms of silence) was inserted between the CV context and stop when making the CVC stimuli.

Procedure. Because many of the procedural details were the same as those in the preceding experiments, only differences will be mentioned. Listeners were tested in a single session, in which there were two blocks of 160 trials, with each step in each continuum being presented ten times, for a total of 20 presentations of each step. Sixteen practice trials (one presentation of each step) were presented at the beginning of the session.

Results and Discussion

The categorization data (collapsed across subjects) are plotted in Fig. 4. Endpoint labeling is close to the extremes. The transition in labeling from one response category to the other is abrupt, indicating that no more than a few steps were perceived as partially ambiguous. In spite of this, one can see that the TP context biased labeling in the expected direction. /p/ responses were more frequent in the *yee* context than *chai*

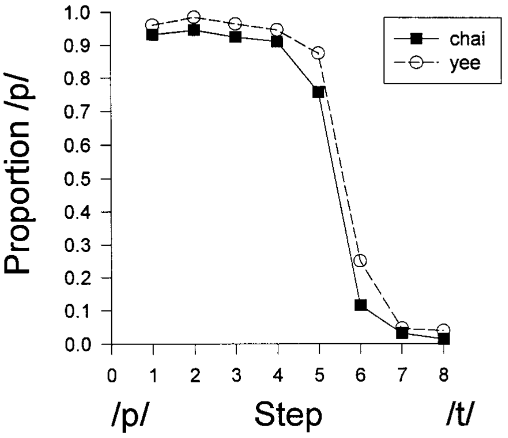

FIG. 4. Experiment 4. Overall proportion of /p/ responses to the eight-step /p/–/t/ continuum (filled squares and solid lines for the context with a /t/ bias, *chaip–chait,* and open circles and dashed lines for the context with a /p/ bias, *yeep–yeet*).

context. The size of the TP effect was .09 units (using the area-between-the-curves measure), which was reliable, $F(1,13) = 16.67, p < .001$.

The present results suggest that the transitional probability between the vowel and following stops biased listeners' perception of those stops. The TP effect observed with fricatives in the earlier experiments was generalized to a stop continuum. Furthermore, this effect was found using the simpler and widely used phonetic categorization paradigm. The data are qualitatively similar to those found with this paradigm when a lexical context is used. These findings provide another piece of evidence to the growing literature that suggests listeners are sensitive to the statistical regularities of the language (e.g., McQueen & Pitt, 1996; Saffran et al., 1996). In the present case, transitional probabilities between phonemes appear to be an additional source of information that listeners can use to process spoken language.

GENERAL DISCUSSION

The pattern of results across the four experiments is clear. In Experiment 1, the TPs of /s/ and /ʃ/ at the end of nonwords influenced both the way listeners labeled an ambiguous fricative at the end of the nonwords and the way they labeled the following word-initial stops. TP appears to influence the compensation for coarticulation process, so that the stops after an ambiguous fricative tend to be labeled as if the listeners had heard the more probable unambiguous fricative. The TP effects in fricative and stop labeling were replicated in Experiment 2, where there were no cues to fricative identity in the initial CV contexts, and in Experiment 3, where in addition to the absence of formant transition cues to fricative identity in the contexts, there was also no confound due to vowel rounding in the TP contexts. TP effects in categorization were also observed in the labeling of nonword-final stops (Experiment 4). The influence of TP on phonetic categorization is robust and replicable, as is the indirect influence that fricative TP has on compensation for coarticulation.

In contrast, Experiment 1 also showed that when TP on word–final fricatives was controlled, there was no compensation for coarticulation on stops following word–final ambiguous fricatives. If the compensation process were influenced by the lexicon, there should have been a tendency to label ambiguous stops as /t/ after *bu?,* as if the listeners had heard *bush,* and as /k/ after *jui?,* as if the listeners had heard *juice.* The failure to observe a lexically mediated compensation for coarticulation effect was not a consequence of a poor manipulation of lexical context. Listeners showed strong lexical effects in their categorization of the word–final fricatives; they did tend to hear *bu?* as *bush* and *jui?* as *juice.* It would therefore appear that although listeners were using lexical knowledge in making their perceptual decisions to the fricatives, this knowledge was unable to influence their decisions to the stops. The fact that compensation for coarticulation was found after the TP-biased nonwords (even though the TP effect on the fricatives was smaller than the lexical effect) shows that there was sufficient power in Experiment 1 to observe a lexically mediated compensation effect had there been one. Experiments 2 and 3 reinforce these findings. In both experiments, the lexical and TP effects on fricative identification were equivalent, but although the TP bias influenced compensation for coarticulation, there was still no evidence of lexical involvement in the compensation process.

Taken together, these results suggest that Elman and McClelland's (1988) demonstrations of apparent lexical involvement in compensation for coarticulation may have been demonstrations of a TP effect. The materials in their Experiment 1 were confounded with TP, and, as the present results suggest, it is this aspect of the materials, not their lexical status, which produced compensation for coarticulation.

Readers familiar with Elman and McClelland (1988) may recall that in their Experiment 2, listeners were presented with only the final vowel and the ambiguous fricative from *ridiculou?* and *Spani?*. If compensation were due to the TPs of /s/ and /ʃ/ after the vowels alone, as in the present experiments, compensation effects should have been observed in Elman and McClelland's Experiment 2. As shown in their Fig. 6b, there was such an effect, albeit an unreliable one. This nonsignificant trend is consistent with the present results. Inspection of their Fig. 6b shows that all of the data-points in the ambiguous region of the continuum, bar one, were in the predicted direction (more velar /g/ responses in the /ə?/ context than in the /ɪ?/ context). Other than this one data-point, the shift appears to be as large as that observed in the present Experiment 1. Elman and McClelland's (1988) ANOVA, on the proportion of responses across the entire continuum, assigns equal weight to the endpoint and ambiguous responses in its test of a main effect of context. Since the compensation effect was limited to responses to ambiguous stops, an analysis procedure focusing only on those responses would have been more likely to have yielded a significant effect. The outcome of this experiment is therefore ambiguous. Whereas Elman and McClelland (1988) chose to dismiss the trend, we believe it is suggestive of a TP effect and certainly in line with the present findings.

Another concern with the present results is that they could be due to the acoustic information in the vowels, rather than to TP. Experiments 2 and 3 address the most serious potential acoustic confounds in the materials of Experiment 1. In Experiment 2, the contexts were constructed from CVs spoken in isolation and thus could not contain any spectral cues to the following fricatives. In Experiment 3, the contexts were again constructed from CVs spoken in isolation, and, furthermore, the rounded–unrounded status of the vowels, which had previously been shown to influence fricative identification (Mann & Repp, 1980; Whalen, 1981, 1989; Soli & Mann, 1982), was controlled. Both vowels in the lexical context were rounded, and both vowels in the TP context were unrounded.

Any remaining acoustic confound, acting to bias fricative and/or stop identification, would probably have to be due to information specifying vowel identity. This seems rather implausible, because the explanation would have to apply across a wide range of different vowels. In some cases, the acoustic biases in the vowels would need to produce a shift in identification (for the fricatives and stops following the vowels used in the TP contexts in Experiments 1–3, for the stops following the vowels used in Experiment 4, and for the stops following the vowels used by Elman & McClelland, 1988). In other cases, the acoustic biases would need to produce no shift in identification (for the stops following the vowels used in the lexical contexts in Experiments 1–3, and for the stops following two further pairs of vowels—/ɛ/ paired with /ɪ/ and /ɒ/ paired with /ɑ/—in lexical contexts tested by McQueen, 1991b, which also failed to show any compensation for coarticulation effects). What is common across these different pairs of vowels are their TPs, not their acoustic structure. It therefore seems much more plausible that the effects observed are due to TPs rather than to acoustic structure.

A further control experiment from Elman and McClelland (1988) raises a different concern. In their Experiment 3, an ambiguous final syllable, /lV?/, where the vowel /V/ was ambiguous between /ə/ and /ɪ/, and /?/ was an ambiguous fricative between /s/ and /ʃ/, based on *ridiculous* and *foolish*, was placed after *''ridicu''* and *''foo.''* There was still reliable compensation for coarticulation on subsequent stops, consistent with the final fricatives being lexically restored. Elman and McClelland argued that because, in addition to this final syllable, the vowels before /l/ in each context were the same, the phonological contexts were identical for three phonemes before the final fricative, and hence that the effect they observed here could not be due to different sequential constraints in the two contexts. But this is an empirical question. Given the present demonstrations of local TP effects (effects due to the transition into the ambiguous fricative), it is possible that listeners are also sensitive to longer-range sequential constraints spanning several segments.

It is unlikely that the vowels before /l/ in *ridiculous* and *foolish* were in fact the same. In its normal pronunciation, the third vowel in *ridiculous* is the back rounded vowel /U/, as in *bush.* In contrast, the first vowel in *foolish* is the higher back rounded vowel /u/, as in *juice.* If these vowels were indeed different, the TPs in these two contexts would have been very biased.[2] If the vowels before /l/ were /U/ and /u/, and listeners are sensitive to longer-range dependencies, the effect observed in Experiment 3 in Elman and McClelland (1988) could have been due to TP biases.

It is important to note that the present experiments tested only local TP effects. The above account depends on longer-range TP sensitivities, for which there is currently no direct evidence. It therefore remains possible that the effects after /lV?/ do indeed indicate lexical involvement in compensation for coarticulation. The present results, however, argue against this view, or at least constrain it to longer words such as those used by Elman and McClelland. It is clear that further experiments which test listeners' sensitivity to longer-range TPs are required to address these issues.

For example, it will be necessary to distinguish between TP effects due to local diphone transitions, as studied here, and possibly independent TP effects due to larger units, such as syllables and words. In particular, a clear distinction between TP effects due to several phoneme transitions and effects due to specific lexical entries will need to be established. At some point, particularly for longer words, the transitional likelihood of the final phoneme may be determined only by the transitions between all the phonemes in a single lexical entry. It will thus be important to ascertain whether contextual effects are due to knowledge about specific words and/or to knowledge generalized over the vocabulary as a whole.

Finally, TP effects need to be distinguished not only from effects due to specific words but also from the size and composition of the lexical neighborhood (i.e., effects due to the combined influence of the lexical neighbors of a word or nonword; Luce, Pisoni, & Goldinger, 1990). Although lexical neighborhoods and TPs are highly correlated (words in dense lexical neighborhoods tend to consist of common phoneme sequences), it appears that they have independent effects: words with dense neighborhoods and high TPs are named more slowly than words with sparse neighborhoods and low TPs, while the reverse is true for nonwords (Vitevitch & Luce, in press). Such results suggest that sensitivity to transitional probabilities is not due solely to lexical neighborhoods. Vitevitch and Luce (in press) argue that the slowdown in naming that was observed with words reflects a neighborhood effect at the lexical level, while the speedup observed with nonwords reflects a TP effect due to prelexical processes.

Theoretical Implications

What implications do the present findings have for interactive and autonomous models of speech perception? When evaluating these models, it is necessary to distinguish between effects due to a specific lexical entry and effects due to the simultaneous influence of multiple entries (i.e., knowledge of TPs extracted from the complete vocabulary). In interactive models like Trace (McClelland & Elman, 1986), both lexical and TP effects on phonetic categoriza-

[2] /ulɪʃ/ occurs in words like coolish, foolish, and ghoulishness. CELEX shows that this string occurs about 26 times per million words. /ulɪs/ and /uləs/ occur less than once per million words, and /uləʃ/ does not occur. So /ʃ/ is much more likely given /ulV?/. The opposite bias operates after /U/. The string /Uləs/ is quite common, in words like incredulously, ridiculous, and stimulus. The CELEX estimate is 86 times per million words. /Ulɪs/ also occurs (in words like oculist and somnambulist, 9 times per million), but /Uləʃ/ and /Ulɪʃ/ never occur. So /s/ is much more likely after /UlV?/.

tion originate in the lexicon (the word–node level). All phonetic decisions are based on the activation levels of phoneme nodes. Lexical influences on phonetic decisions arise when top-down facilitation from word nodes boosts the activation of lexically consistent phoneme nodes. TP influences operate in the same way in Trace. Words sharing the same phonemic sequences (i.e., those in the lexical neighborhood) feed back small amounts of activation individually, but when combined over many words, activation can be sufficiently large to bias perception. For example, phonotactic context effects in phonetic categorization (Massaro and Cohen, 1983; Pitt, in press), which can be described as extreme TP effects that involve sequences of phonemes with TPs of zero, have been explained in Trace as being due to a ‘‘conspiracy’’ of lexical entries, acting top-down to bias phoneme–node activation (McClelland & Elman, 1986; McClelland, 1991).

There are two important consequences of lexical and TP effects originating from the same source. First, it is difficult to distinguish between them, since they are quantitatively similar outcomes rather than qualitatively different phenomena. Second, these lexical and TP effects must come and go together. Since both are due to top-down processing, lexical effects should be present when there are TP effects and vice versa. If one effect is absent, so should the other one be. Of course, TP sensitivities could be built into the prelexical level of an interactive model. Such a design modification seems unnecessary because it is redundant. The top-down connections already provide a mechanism by which knowledge about the vocabulary as a whole can influence prelexical processing.

Autonomous models account for lexical and TP effects in a different manner. Because by definition there are no top-down connections, there can be no effects on prelexical processing due to specific lexical entries. Prelexical processes act to map spoken input onto the lexicon, but are unaffected by knowledge stored in the lexicon. There can however be lexical effects in perceptual decisions. In autonomous models like Merge (McQueen et al., submitted; Norris et al., submitted) and FLMP (Oden & Massaro, 1978), knowledge about specific words is output from the lexicon and can then influence categorization decisions. In principle, therefore, knowledge about the vocabulary as a whole could also be output from the lexicon and influence categorization decisions. TP effects on perceptual decisions could thus emerge in autonomous models, as in interactive models, as a result of biases due to gangs of activated words in lexical neighborhoods. Whether or not this could occur depends on another important assumption about the nature of the output from the lexicon in an autonomous model. If lexical output to the decision process is all or none, TP effects due to output from the lexicon could only influence categorization decisions to words, not to nonwords. But if lexical output were continuous and was fed partial information to the decision process (as in Merge), TP effects due to output from the lexicon could influence perceptual decisions in both words and nonwords.

But TP sensitivity could also arise at a prelexical level in autonomous models. Indeed, for knowledge about the vocabulary as a whole to influence prelexical processing, and not only perceptual decision-making, that knowledge must be stored at the prelexical level in an autonomous model. Recurrent networks (Chater et al., submitted; Norris, 1993) can learn sequential dependencies between speech sounds in the absence of learning any specific words. They show that TPs can be stored prelexically. Any autonomous model with the ability to store knowledge generalized over the vocabulary at a prelexical level could thus account for effects of TP on prelexical processing.

This discussion raises a distinction which should be made fully explicit regarding autonomous models. This is the distinction between effects of context on perceptual decisions and effects on prelexical processing. These are not the same in an autonomous model. The defining feature of these models, the lack of top-down connections, requires that knowledge about specific words cannot influence prelexical processing. This knowledge can nonetheless influence perceptual decisions on the basis of lexical output. On the other hand, the effects of TP depend on where that knowledge is stored. If TP knowledge is built into the prelexical level, this knowledge can influence both prelexical processing and perceptual decision-making. But if it is coded at the lexical level, it can of course only influence perceptual decision-making. The distinction between prelexical processing and phonetic decision-making is not so clear-cut in interactive models. Prelexical processing devices (e.g., the phoneme nodes in Trace) provide the basis for perceptual decisions. Context effects on phonetic decisions should thus go hand-in-hand with effects on prelexical processes.

How then do the present results constrain these theoretical alternatives? Viewed alone, the effects of TP-bias on fricative and stop categorization are consistent with both classes of model, but they do constrain the autonomous account. Since it is reasonable to assume, as Elman and McClelland (1988) did, that the compensation for coarticulation process is prelexical, the TP effect on stop categorization suggests that on an autonomous account, sequence constraints have their effect prelexically and are not due to a neighborhood effect at the output from the lexicon. In isolation, this result does not constrain the interactive account: The TP effects on fricative decisions and on the compensation process could be due either to top-down processing or to the prelexical storage of sequence constraints.

The effects from the word contexts, however, challenge interactive models. There was no modulation of the compensation process when lexical items were balanced in TP (e.g., *juice*, *bush*). It is important to emphasize that this was not simply a failure to obtain a lexical context effect. There were strong lexical biases in listeners' fricative identifications in all three experiments. It is this result, lexical effects on fricative identification with no consequent effect on stop identification, that is problematic for interactive models such as Trace. Its explanation for the increased proportion of /s/ responses in the *jui?* context, for example, is that lexical feedback increases the activation of the /s/ node. Fricative node activation in turn modulates the strengths of the connections between feature nodes and stop nodes, causing perceptual compensation (see Elman & McClelland 1986, 1988 for further details). If lexically consistent fricative nodes have elevated activation, as the lexical bias in fricative categorization implies, there should have been a compensation effect.

The current instantiation of Trace has difficulty accounting for this dissociation between fricative and stop categorization. This is because interactive models like Trace predict that effects on perceptual decisions and prelexical processing should go hand-in-hand. We obtained an effect on fricative decisions but no corresponding effect on the prelexical compensation for coarticulation process (i.e., on the stop decisions). The TP-bias nonwords add a second compelling dissociation. If the TP-bias effect were due to top-down processing, the results from the TP-bias conditions would suggest that the lexicon influenced both fricative identification and compensation for coarticulation. But this pattern was not observed in the lexical-bias conditions. Since top-down TP and lexical effects should come and go together, the different pattern of results across the two context conditions is problematic for Trace.

What instead is required to account for both dissociations is a model in which transitional probabilities are represented separately from lexical information (i.e., prelexically), such as in an autonomous model. In a model such as Merge, even when listeners are using knowledge about specific lexical entries in their fricative decisions, these entries cannot influence the compensation for coarticulation process, and identification of stops will therefore not be affected by lexical bias. At the same time, however, knowledge about the vocabulary as a whole, coded at the prelexical level, can influence both fricative labeling and (via the compensation process) stop labeling in the TP-bias contexts.

We believe the present results provide an informative constraint on models of speech perception. They support models in which knowledge generalized over the vocabulary is coded in the system. They suggest that information about local TPs, but not lexical knowledge (at least for short words), can affect compensation, which leads one to believe that TPs are represented prelexically and not at the level of the lexicon. The dissociation between TP and lexical neighborhood effects observed in naming words and nonwords (Vitevitch & Luce, in press) also suggests that TPs are coded independently of the lexicon. Further research should establish the extent of listeners' sensitivity to transitional probabilities (how long-range these sensitivities are) and seek to establish the type of mechanism responsible for this sensitivity. Such an endeavor might also provide insight into the characteristics of the representation on which TPs operate. For example, are effects of TPs contingent on the representation's size (e.g., diphone, syllable), or is such information stored with individual phonemes? It will also be necessary to determine the nature of apparent lexical effects in tasks requiring phonetic judgments. That is, are such effects in fact due to knowledge about specific lexical entries in a listener's language, or are they due to knowledge of the sequence constraints of that language?

APPENDIX

Synthesis Parameters for the Endpoints of the /ʃ/-/s/ and /p/-/t/ Continua

Parameter	Time	/ʃ/ value	/s/ value	Parameter	Time	/p/ value	/t/ value
AF	0	0	00	AF	0	0	0
AF	30	15	15	AF	5	65	52
AF	50	54	54	AF	40	50	30
AF	190	54	54	AF	45	0	27
AF	240	40	40	AF	65	0	20
F0	0	120	120	AF	75	0	0
F0	240	120	120	AH	0	0	0
F1	0	460	460	AH	5	70	52
F1	240	460	460	AH	55	58	30
F2	0	2300	2300	AH	65	29	20
F2	240	2300	2300	AH	75	0	0
F3	0	2800	2800	F0	0	132	132
F3	240	2800	2800	F0	75	120	120
F4	0	3400	3400	F1	0	300	460
F4	240	3400	3400	F1	75	300	360
A3	0	5	0	F2	0	1300	2000
A3	90	38	0	F2	75	1300	1600
A3	190	40	0	F3	0	2000	2600
A3	240	35	0	F3	75	2000	2300
A4	0	5	0	AB	0	0	0
A4	90	47	0	AB	10	60	0
A4	190	47	0	AB	75	20	0
A4	240	40	0	A3	0	0	0
A5	0	5	10	A3	5	0	30
A5	30	20	25	A3	75	0	10
A5	50	30	30	A4	0	0	0
A5	90	41	30	A4	5	0	42
A5	190	41	30	A4	75	0	10
A5	240	30	10	A5	0	0	0
A6	0	5	10	A5	5	0	52
A6	30	19	30	A5	75	0	39
A6	50	24	48	A6	0	0	0
A6	90	38	48	A6	5	0	50
A6	190	38	48	A6	75	0	39
A6	240	30	20				
B1	0	100	100	B1	0	300	300
B1	50	200	200	B1	75	300	300
B1	240	200	200	B2	0	200	240
B2	0	300	300	B2	75	200	240
B2	50	200	200	B3	0	150	200
B2	240	120	120	B3	75	150	200
B3	0	240	240				
B3	50	260	160				
B3	240	240	140				

REFERENCES

Burnage, G. (1990). *CELEX: A guide for users.* Nijmegen, The Netherlands: CELEX.

Burton, M. W., Baum, S. R., & Blumstein, S. E. (1989). Lexical effects on the phonetic categorization of speech: the role of acoustic structure. *Journal of Experimental Psychology: Human Perception and Performance,* **15,** 567–575.

Cairns, P., Shillcock, R., Chater, N., & Levy, J. (1995). Bottom-up connectionist modelling of speech. In J. P. Levy, D. Bairaktaris, J. A. Bullinaria, & P. Cairns (Eds.), *Connectionist models of memory and language* (pp. 289–310). London: UCL Press.

Chater, N., Shillcock, R., Cairns, P., & Levy, J. (submitted). Bottom-up explanation of phoneme restoration: Comment on Elman and McClelland (1988).

Connine, C. M., & Clifton, C. (1987). Interactive use of lexical information in speech perception. *Journal of Experimental Psychology: Human Perception and Performance,* **13,** 291–299.

Cutler, A., Mehler, J., Norris, D., & Segui, J. (1987). Phoneme identification and the lexicon. *Cognitive Psychology,* **19,** 141–177.

Cutler, A., & Norris, D. (1979). Monitoring sentence comprehension. In W. E. Cooper & E. C. T. Walker (Eds.), *Sentence processing: Psycholinguistic studies presented to Merrill Garrett* (pp. 113–134). Hillsdale, NJ: Erlbaum.

Elman, J. L., & McClelland, J. L. (1986). Exploiting lawful variability in the speech wave. In J. S. Perkell & D. H. Klatt (Eds.), *Invariance and variability in speech processes* (pp. 360–385). Hillsdale, NJ: Erlbaum.

Elman, J. L., & McClelland, J. L. (1988). Cognitive penetration of the mechanisms of perception: Compensation for coarticulation of lexically restored phonemes. *Journal of Memory and Language,* **27,** 143–165.

Fox, R. A. (1984). Effect of lexical status on phonetic categorization. *Journal of Experimental Psychology: Human Perception and Performance,* **10,** 526–540.

Ganong, W. F. (1980). Phonetic categorization in auditory word perception. *Journal of Experimental Psychology: Human Perception and Performance,* **6,** 110–125.

Jusczyk, P. W., Luce, P. A., & Charles-Luce, J. (1994). Infants' sensitivity to phonotactic patterns in the native language. *Journal of Memory and Language,* **33,** 630–645.

Klatt, D. H. (1980). Software for a cascade/parallel formant synthesizer. *Journal of the Acoustical Society of America,* **67,** 971–975.

Luce, P. A., Pisoni, D. B., & Goldinger, S. D. (1990). Similarity neighborhoods of spoken words. In G. T. M. Altmann (Ed.), *Cognitive models of speech processing: Psycholinguistic and computational perspectives* (pp. 122–147). Cambridge, MA: MIT Press.

Mann, V. A., & Repp, B. H. (1980). Influence of vocalic context on perception of the [ʃ]-[s] distinction. *Perception & Psychophysics,* **28,** 213–228.

Mann, V. A., & Repp, B. H. (1981). Influence of preceding fricative on stop consonant perception. *Journal of the Acoustical Society of America,* **69,** 548–558.

Massaro, D. W., & Cohen, M. M. (1983). Phonological constraints in speech perception. *Perception & Psychophysics,* **34,** 338–348.

Massaro, D. W., & Oden, G. C. (1995). Independence of lexical context and phonological information in speech perception. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* **21,** 1053–1064.

McClelland, J. L. (1991). Stochastic interactive processes and the effect of context on perception. *Cognitive Psychology,* **23,** 1–44.

McClelland, J. L., & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology,* **18,** 1–86.

McQueen, J. M. (1991a). The influence of the lexicon on phonetic categorization: Stimulus quality in word-final ambiguity. *Journal of Experimental Psychology: Human Perception and Performance,* **17,** 433–443.

McQueen, J. M. (1991b). *Phonetic decisions and their relationship to the lexicon.* Ph.D. dissertation, University of Cambridge.

McQueen, J. M. (1998). Segmentation of continuous speech using phonotactics. *Journal of Memory and Language,* **39,** 21–46.

McQueen, J. M., & Pitt, M. A. (1996). Transitional probability and phoneme monitoring. *Proceedings of ICSLP 96,* Vol. 4, pp. 2502–2505. Delaware: Alfred I. du Pont Institute.

McQueen, J. M., Norris, D., & Cutler, A. *Lexical influence in phonetic decision making: Evidence from subcategorical mismatches.* Manuscript submitted for publication.

Norris, D. (1993). Bottom-up connectionist models of ''interaction.'' In G. Altmann & R. Shillcock (Eds.), *Cognitive models of speech processing: The second Sperlonga meeting* (pp. 211–234). Hillsdale, NJ: Erlbaum.

Norris, D. (1994). SHORTLIST: A connectionist model of continuous speech recognition. *Cognition,* **52,** 189–234.

Norris, D., McQueen, J. M., & Cutler, A. (submitted). Merging phonetic and lexical information in phonetic decision-making.

Oden, G. C., & Massaro, D. W. (1978). Integration of featural information in speech perception. *Psychological Review,* **85,** 172–191.

Pitt, M. A. (1995). The locus of the lexical shift in phoneme identification. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* **21,** 1037–1052.

Pitt, M. A. (in press). Phonological processes and the perception of phonotactically illegal consonant clusters. *Perception & Psychophysics.*

Pitt, M. A., & Samuel, A. G. (1993). An empirical and meta-analytic evaluation of the phoneme identification task. *Journal of Experimental Psychology: Human Perception and Performance,* **19,** 699–725.

Pitt, M. A., & Samuel, A. G. (1995). Lexical and sublexical feedback in auditory word recognition. *Cognitive Psychology,* **29,** 149–188.

Repp, B. H. (1981). Perceptual equivalence of two kinds of ambiguous speech stimuli. *Bulletin of the Psychonomic Society,* **18,** 12–14.

Repp, B. H., & Mann, V. A. (1981). Perceptual assessment of fricative-stop coarticulation. *Journal of the Acoustical Society of America,* **69,** 1154–1163.

Repp, B. H., & Mann, V. A. (1982). Fricative-stop coarticulation: Acoustic and perceptual evidence. *Journal of the Acoustical Society of America,* **71,** 1562–1567.

Saffran, J. R., Newport, E. L., & Aslin, R. N. (1996). Word segmentation: the role of distributional cues. *Journal of Memory and Language,* **35,** 606–621.

Samuel, A. G. (1981). Phonemic restoration: Insights from a new methodology. *Journal of Experimental Psychology: General,* **110,** 474–494.

Samuel, A. G. (1996). Does lexical information influence the perceptual restoration of phonemes? *Journal of Experimental Psychology: General,* **125,** 28–51.

Shillcock, R., Lindsey, G., Levy, J., & Chater, N. (1992). A phonologically motivated input representation for the modelling of auditory word perception in continuous speech. In *Proceedings of the 14th Annual Cognitive Science Society Conference,* Bloomington, pp. 408–413.

Soli, S. D., & Mann, V. A. (1982). The influence of vocalic context on the /s/ - /ʃ/ distinction in initial and final position. Paper presented to the 103rd meeting of the Acoustical Society of America. *Journal of the Acoustical Society of America,* **71,** S75.

Svartvik, J., & Quirk, R. (1980). *A corpus of English conversation.* Lund: Gleerup.

Vitevitch, M. S., & Luce, P. A. (in press). When words compete: Levels of processing in spoken word perception. *Psychological Science.*

Whalen, D. H. (1981). Effects of vocalic formant transitions and vowel quality on the English [s]-[š] boundary. *Journal of the Acoustical Society of America,* **69,** 275–282.

Whalen, D. H. (1989). Vowel and consonant judgments are not independent when cued by the same information. *Perception & Psychophysics,* **46,** 284–292.

(Received March 20, 1997)
(Revision received January 30, 1998)