Lexical cues to foreignness in Japanese*

Ben Gelbart and Shigeto Kawahara
University of Massachusetts, Amherst

This article addresses two questions regarding lexical stratification in Japanese: Are lexical strata psychologically real to native speakers of Japanese? And, if real, do they include words that lack phonotactic cues to foreignness? The results of a perceptual experiment suggest affirmative answers to both questions. Japanese speakers identified acoustic continua in pairs of lexical items that differed in little else other than historic origin. The perceptual boundaries shifted when the carrier words were of foreign origin. These results indicate that native-foreign distinctions are psychologically real even for words that lack phonotactic cues to stratal affiliation.

1. Introduction

Lexical stratification has been a topic of recent debate in Japanese phonology, inspired by the pioneering works by Itô and Mester (1995, 1999, 2001). A concrete example of lexical stratification in Japanese is the treatment of voiced obstruent geminates. Not only are voiced obstruent geminates absent from native lexical items, they are actively avoided by an alternation. The suffix *-ri* induces gemination of root-final consonants, as in (1a), but when the target of gemination is a voiced obstruent the result is instead a homorganic nasal + voiced obstruent sequence, as shown in (1b).

(1)	a.	u**k**a(-uka)	u**kk**a-ri		‘absentmindedly’
		hi**s**o(-ka)	hi**ss**o-ri		‘secretly’
		ho**n**o(-ka)	ho**nn**o-ri		‘dimly, faintly’
	b.	za**b**u(-zabu)	za**mb**u-ri	*za**bb**u-ri	‘jumping into water’
		u**z**a(-i)	u**nz**a-ri	*u**zz**a-ri	‘annoyed’
		ko**g**a(-su)	ko**ng**a-ri	*ko**gg**a-ri	‘toasted, roasted’

(Itô and Mester 1999)

However, voiced geminates are allowed in loanwords: *we**bb**u* ‘website,’ *be**dd**o* ‘bed,’ *ba**gg**u* ‘bag,’ etc. Itô and Mester propose that foreign words are indexed as such in the lexicon, and that the indexes correspond to independently

* Thanks to John Kingston, Kazu Kurisu, Joe Pater, Caren Rotello, Kazuko Shinohara, Takahito Shinya, Lisa Shiozaki, and the audience at Tokyo Circle of Phonologists (Jan, 2006) and FAJL 4 (Aug, 2006) for their comments on this project. Also thanks to Hiroko Yagi for setting up the occasion to run the experiment at Chuo University. All remaining errors are ours.

rankable stratum-specific faithfulness constraints. The stratified lexicon results when these faithfulness constraints are ranked differently with respect to some markedness constraints, e.g. a constraint against voiced obstruent geminates.

Itô and Mester's lexical stratification analysis, however, has met with criticism. For example, from a representational perspective, Inkelas, Orgun, and Zoll (1997) attempt to reduce the exceptional properties and behaviors of the foreign stratum to differences in underlying forms, obviating the need for separate strata. They posit a unique phonological representation for each unique phonological pattern. This approach can be described as a *representational hypothesis*, in contrast to the *lexical stratification hypothesis* proposed by Itô and Mester. The plausibility of lexical strata has also been questioned by Rice (1997), largely on the grounds that when a contrast is prohibited only in some strata, learners should conclude that in the overall grammar, that contrast is contrastive. Thus, when learners see voiced geminates in loanwords, they are likely to consider that voicing is contrastive in geminates in Japanese, rather than positing two different lexical strata (see Itô, Mester, and Padgett 1999 for a reply). Finally, Maye, Suzuki, and Ohno (2000) argue against the psychological reality of lexical strata in Japanese based on their experiments on Rendaku voicing, a process primarily restricted to native words. Maye et al. found that Rendaku voicing applies no more frequently for nonce words that conformed to native phonotactic restrictions than for those that did not.

In addition to these opposing positions, a third position has emerged that concedes that stratification exists but insists that it does not extend over the entire lexicon. For example, Ota argues that "even if we can justify the existence of phonological sublexica, we cannot determine the classhood of all lexica based solely on distributional evidence" (2004: 23). In this view, morphemes that show active alternations are parsed out from the rest of the lexicon, while words that do not display active alternations are treated alike by the phonology; i.e., the lexicon is stratified into words that show alternations and those that do not (Fukazawa, Kitahara, and Ota 2002a, b; Ota 2004; Tateishi 2003). Following Ota (2004), we refer to this approach as the *weak lexical stratification hypothesis*, and Itô and Mester's original position as the *strong lexical stratification hypothesis*.

The results of our experiment reported below support the strong lexical stratification hypothesis, as originally envisioned by Itô and Mester. In the remainder of this article, we provide experimental evidence for it, but first we address the pros and cons of the representational hypothesis and the weak stratification hypothesis.

The representational hypothesis has the advantage of using only independently motivated phonological devices. Its possible drawback is that it has nothing to say about why certain representations such as voiced obstruent geminates are considerably rarer than others or about why these rarer structures only occur in loanwords. However, proponents of the representational hypothesis could plausibly argue that these are results of historical factors and need not be addressed by the synchronic grammar (see also Itô and Mester 2001 for further critical discussion on the representational hypothesis).

Stronger evidence against the representational hypothesis comes from

speech perception experiments which demonstrate speakers' recourse to lexical strata. It has been well established in the speech perception literature that listeners are biased against perceiving grammatically illicit structure (e.g. Massaro and Cohen 1983). Building on this observation, Moreton and Amano (1999) showed that stratum specific phonotactic constraints can affect speech perception. In Japanese, only the foreign stratum contrasts short [a] and long [a:] word-finally. Other strata can contain short [a], but not long [a:] in word-final position. Moreton and Amano used carrier words of the shape C_1oC_2a, where the final [a] varied gradiently in duration. The C_1 and C_2 independently cued stratal affiliations of the stimuli: [p] and [f] cued foreignness, and [r^y] and [h^y] cued the Sino-Japanese (SJ) (Moreton, Amano, and Kondo 1998), and [r] and [t] were neutral i.e. the stimuli could be taken to belong to any strata. The [a]-[a:] boundaries—points in the continuum where the long [a:] percept becomes dominant—for each type of stimulus are shown in Figure 1.

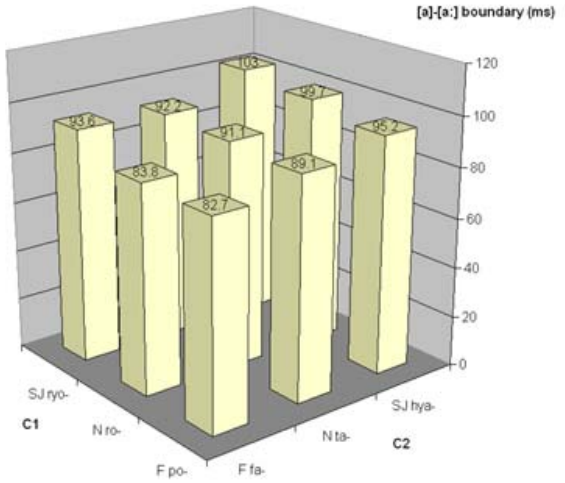

Figure 1. The results of Moreton and Amano (1999) (their Figure 1).

The result in Figure 1 suggests that given ambiguous stimuli with the same duration, listeners are most likely to judge [a] as long [a:] once the stimuli are perceived as belonging to the foreign stratum in which word-final long [a:] is permitted. On the other hand, listeners are biased against judging [a] as long once the stimuli are perceived as belonging to the SJ stratum in which word-final long [a:] is prohibited. These results suggest two conclusions: (i) the sounds that cue stratal affiliation ([p,f,r^y,h^y]) cause listeners to associate the stimuli with a particular stratum, and (ii) a stratum-specific phonotactic constraint (*[a:] /_ $]_{wd}$ in SJ) can affect speech perception. In addition, the effects of consonants that cue lexical affiliation of stimuli are cumulative; e.g. two consonants that cue foreignness have a stronger effect than just one.

The result of Moreton and Amano cannot be explained from the representational perspective, as the representational model does not encode

lexical affiliation in the lexicon. The result is also problematic for the weak lexical stratification hypothesis, which asserts that lexical stratification is restricted to cues that participate in alternations. The weak stratification hypothesis maintains that the phonotactics cannot trigger phonological subclasses: "the existence of a particular sound only gives us a positive piece of evidence that it exists, not that all the other sounds do not exist, so the existence of [p] in Japanese, for example, simply adds another phoneme to phonological inventories of Japanese." (Tateishi 2003:13). Yet in Moreton and Amano's experiment, [p,f,r^y,h^y] all caused listeners to associate the stimuli with a particular stratum, as evidenced by the perceptual boundary shifts of the word-final [a]-[a:] continuum, even though none of these consonants participate in any overt alternations.[1]

However, the results of Moreton and Amano (1999) do not fully address the criticism posed by the weak lexical stratification hypothesis concerning the question of lexical items that are phonotactically ambiguous as to their stratal affiliation. As an example, Ota (2004) uses the pair of words *tombo* "dragonfly" (native) vs. *kombo* "combo" (foreign). Neither word contains marked structure that might cue their stratal affiliations, so there is no reason, at least on phonological grounds, for native speakers to treat them as belonging to different phonological lexical strata. Our experiment nonetheless shows that Japanese phonology treats pairs like *tombo* and *kombo* differently.

2. Experiment

Our experiment builds on Moreton and Amano (1999) in attempt to provide additional evidence for the psychological reality of stratification. The novel aspect of our experiment is that it addresses the issue of whether lexical stratification exists even when there are no phonological cues at all (i.e. for pairs like *tombo* and *kombo*). In short: in the absence of any phonotactic cues, does Japanese phonology treat such pairs of lexical items differently? We answer positively to this question.

2.1 Design

Our design resembles that of Moreton and Amano (1999) in that each item contains both an unambiguous cue to trigger stratal identity and a continuum to measure a perceptual boundary. The main difference between our experiment and theirs is that instead of using phonotactic cues as triggers, we used real lexical items that contained no phonotactic cues. For us, the lexical item itself was the cue. Given that our goal is to show the perceptual effects of stratal affiliation as a robust phenomenon, several contrasts were used as measures: the voiced stop geminacy contrasts [b]~[bb], [d]~[dd], and [g]~[gg], as well as the vowel length contrast [a]~[a:] in word-final position. In all four contrasts, the

[1] Underlying /p/ is arguably mapped to [h] in the native phonology. However, this alternation is necessary only when we posit /p/ for the underlying form of [h]. We could posit /h/ for [h], in which case no overt active alternations exist for the mapping (Fukazawa, Kitahara, and Ota 2002b).

long variant is legal only in the foreign stratum while the short variant is legal in all strata. Participants are asked to identify the geminacy or length of each stimulus in a forced-choice task. The same physical continuum is presented in different carrier words. Thus, any difference in the location of the perceptual boundary is interpreted to be the result of differences between carrier words.

Eight word pairs were selected. Each pair consisted of two words, one native and one foreign, which were otherwise matched for prosodic similarity. Neither word contained a non-native phone, but each contained a shared phone ([b],[d],[g],[a]) that, when made long, would violate the phonotactics of the native phonology. Prolonging these phones makes both words into non-words, so the lexical bias is not greater in either the foreign stimuli or the native stimuli. The only difference is that the non-words are phonotactically possible for the foreign items but not for the native items.

There were two pairs for each of the four manipulated segment types ([b],[d],[g],[a]), as listed in (2). For each word, the manipulated segment is shown in boldface and underscore.

(2)

Native word		Foreign word	
tá**b**i	'Japanese socks'	ná**b**i	'navigation'
nara**b**u	'line up'	kura**b**u	'dance club'
ku**d**aru	'to descend'	me**d**aru	'medal'
hára**d**a	Last name	sára**d**a	'salad'
tó**g**u	'sharpen'	má**g**u	'mug'
sá**g**a	Place name, 'nature'	né**g**a	'negative'
mós**a**	'tough guy'	nás**a**	'NASA'
nóhar**a**	'field'	sáhar**a**	'Sahara'

The pairs differed in the number of syllables (two or three) and the location of the manipulated segment: second consonant (=C_2), third consonant (=C_3), or final vowel. This variation was included in order to test how robust the predicted perceptual bias is: a robust effect would not be expected to be limited to a particular segment, syllable or position. Unfortunately, it was not possible to fully cross the three factors of segment type, segment position, and number of syllables. The word-final [a] must be word-final and could not be varied for segment position. Additionally, no suitable three-syllable word pairs could be found that contained [g] and no suitable two-syllable word pairs could be found that contained [d]. The two [d]-pairs differed by segment position, though, and the two [a]-pairs differed by number of syllables.

We chose pairs of words that have the same accent locations. We also attempted to control for frequencies of paired items based on Amano and Kondo's (2000) database, although not all of the paired words were found in it. Additionally, a questionnaire was given to all participants at the start of the experiment in which they were asked to rate how familiar the words seemed to them on a scale of 1 to 7. The results of the questionnaire suggested that the word pairs varied significantly in terms of how well they matched for frequency, although familiarity ratings do not appear to have affected the results of the experiment in any consistent way.

All 16 words were manipulated to form continua from a word to a non-word, where the endpoint containing the long segment ([bb], [dd], [gg], and word-final [aa]) corresponds to the non-word. The prediction of strong lexical stratification hypothesis is that listeners will cross over from word to non-word judgments at shorter durations for the foreign words than for their native counterparts i.e. the sound with the same duration will be more likely to be heard as long in the foreign stimuli than in the native stimuli, because long segments are legal only in the foreign stratum.

2.2 Method

2.2.1 Stimuli

In order to create the stimuli, a male native speaker of Japanese was recorded. He pronounced the stimuli in a carrier phrase in a sound-attenuated booth at the University of Massachusetts, Amherst. The speaker was asked to pronounce the words in two types of repetitions. In the first type of repetition, the word was pronounced as it normally is. In the second type, it was pronounced as if the manipulated segment ([b],[d],[g],[a]) was unnaturally long (longer than a naturally produced token in a similar context). This extra-long segment was then used to create a continuum by repeatedly cutting out several glottal pulses until the last step in the continua was as short or shorter than a naturally produced singleton. The [b], [d], and [g] continua had 10 steps each. The [a] continuum had 9 steps.

The normally pronounced stimulus words from the first type of repetition were then used to make the carrier words. First, the manipulated segment was removed from the carrier word. It was then replaced by all possible steps of the relevant continuum. Since 12 words had manipulated consonants with 10 steps each and the 4 remaining words had manipulated vowels with 9 steps each, the result was 156 different versions of the 16 words.

2.2.2 Procedure

In each trial, two options appeared on a screen while the audio stimulus was played over headphones. The options were written in katakana script for the foreign words and hiragana for the native words. The option on the left was always the standard form of the word (e.g. *togu*), and the option on the right was the form corresponding to the non-word end of the continuum (e.g. *toggu*). The participants were told to press the button on the side that corresponded to the form on the screen that more closely represented the word they had heard. Participants had 2.5 seconds to respond before the trial timed out. The inter-stimulus interval was .25 seconds.

Participants first worked through a training block to familiarize themselves with the task and with the endpoint stimuli. Feedback was provided for each trial for which a response was recorded. The feedback was an audio recording of a female native speaker of Japanese saying either *yoku deki-mashita* “well done” or *zannen* “sorry” as might be said by a schoolteacher.

Following the training, participants were tested on a full block of all 156 stimuli without feedback. Then they were tested on a series of 6 blocks. Each of these blocks presented the two endpoint stimuli once, and the remaining intermediate stimuli twice. In total, therefore, each listener responded to the endpoint stimuli 7 times, and to the intermediate stimuli 13 times. The order of the stimuli was randomized within each block.

2.2.3 Participants

The experiment was first run with 16 native speakers of Japanese recruited at the University of Massachusetts, Amherst. However, it turned out after the experiment that the sizes of response prompts were not uniform. Therefore, a follow-up study was carried out with 26 native speakers of Japanese at Chuo University (Tokyo, Japan). No participants reported a history of hearing problems and all received payment for their time. The results from the two studies are compatible with each other. Due to space limitation, we report here the results from the second running only. See Gelbart (2005, chapter 2) for the complete description of both the first and second running.

2.3 Results and discussion

Overall, almost all the pairs of native-foreign items in (2) show boundary shifts, and hence our results support the strong lexical stratification hypothesis. All participants converged towards 0% and 100% "long" identifications for the endpoint stimuli. For each listener, we calculated the total percentage of tokens identified as "long" in all steps of the foreign word minus the total percentage of tokens identified as "long" in all steps of the native word. Positive values for the percentage difference indicate that listeners identified more of the continuum as long in the foreign item than in the native item. Negative values indicate that listeners identified more of the continuum as long in the native item than in the foreign item. This percentage was thus predicted to be positive in all cases.

The two pairs that had the final vowel continuum and the two pairs that had the C_3 continuum showed positive shifts, as shown in Figure 2 and Figure 3, respectively. In these figures, the error bars represent 95% confidence intervals.

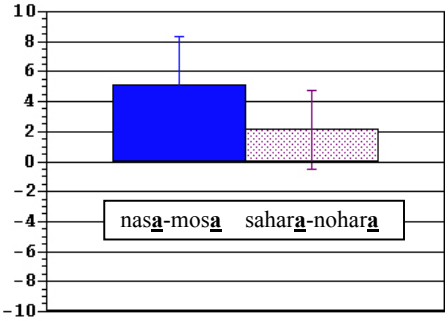

Figure 2. The difference in the total "long" responses (total "long" responses for the foreign item minus the total "long" responses for the native item). The items with the word-final [a] continuum.

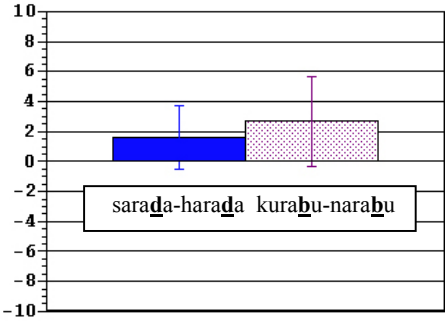

Figure 3. The difference in the total "long" responses. The C_3-continua.

These results directly support the hypotheses that (i) Japanese speakers store foreign items as foreign and native items as native in the lexicon, as predicted by the strong lexical stratification hypothesis, even in the absence of phonological evidence for their stratal affiliations and (ii) speech perception is influenced by stratum-specific phonotactic constraints. Listeners are biased against judging the continua as long for the native items, because of the general ban against long [a:]s and voiced geminates in the native phonology. It is worth pointing out that the bias against the voiced geminates in C_3 positions in Figure 3 does not seem to be more robust than the effect of final [a:] in Figure 2. This result is interesting because the ban on voiced geminates involves active alternations (see (1)), while the ban on final [a:] is merely a static phonotactic cue.

An interesting reversal was observed when the manipulated segment was the second consonant; the shift in "long" identifications was zero or negative, as shown in Figure 4.

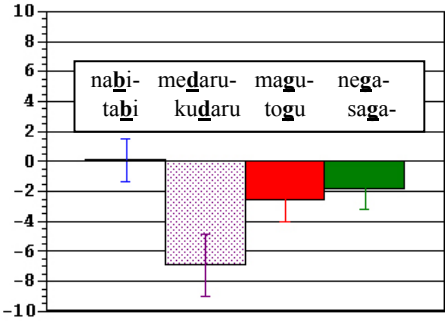

Figure 4. The difference in the total "long" responses. The C_2-continua.

With a single exception (*tabi~nabi*), perceptual boundaries shifted between the foreign items and native items. The unexpected finding was that this bias runs in the opposite direction, as listeners judged more of the continuum as long in the native items than in the foreign items. Although the direction of the boundary

shifts goes the opposite direction as predicted, the perceptual biases were observed in both the first running (not documented in this article; see Gelbart 2005) and the second running. This consistency indicates that there is something principled involved in the perception of geminacy in voiced C_2.

Our explanation for the negative bias in Figure 4 is as follows. While no native Japanese (or Sino-Japanese) words contain voiced obstruent geminates lexically, there is a productive process of emphasis that geminates the onset of the second syllable in native words. This process is not blocked when the consonant in question is a voiced obstruent (Kawahara 2001), as in (3).

(3)				
	ya**b**ai	'dangerous'	ya**bb**ai	'very dangerous'
	ku**d**aranai	'trivial'	ku**dd**aranai	'very trivial'
	su**g**oi	'very'	su**gg**oi	'very (emphatic)'

What may have happened in the experiment is that listeners assumed that native words like *kudaru* were being pronounced emphatically. The emphatic pronunciation very rarely applies to nouns, presumably for semantic reasons, and since foreign items are predominantly nouns, foreign items that contain emphatically geminated voiced geminates are rare at best. As a result, listeners may fail to associate emphatic voiced geminates with foreign items, which could yield a bias *against* judging the stimuli as geminates in the foreign words. Since emphatic gemination primarily targets the C_2 consonant (Kawahara 2006), this bias against voiced geminates in foreign items was not present when the segment to be identified occurred later in the word. If this hypothesis is right, the negative shift is another piece of evidence for the stratified lexicon.

Emphatic gemination in (3) may also explain why the negative shift for *kudaru~medaru* is so much larger than the other pairs in Figure 4. The word *kudaranai* 'trivial, useless', an adjectival derivative of *kudaru*, is frequently emphatically geminated as *ku**dd**aranai*. The other native words in the pairs in (2) may be emphatically geminated, but none of the geminated form has the near-lexical status of *kuddaranai*. Another related issue is that word pairs whose C_2 was manipulated show a negative shift of a highly variable size, while word pairs manipulated in other positions show a positive shift of a less variable size. The difference in variability suggests that the two effects may be of a different nature. The positive effect is due to the difference in strata between the two words in each pair. The negative effect is due to listeners having heard a particular native word pronounced with an emphatic geminate. It would thus depend on how often they have heard that particular word pronounced with an emphatic geminate: the word *kudaru*—or its corresponding adjective *kudaranai*—is frequently emphatically geminated as *kuddaranai* "very trivial," while *tabi* may not have been heard often enough with an emphatic geminate to override the positive shift.

Since all the perceptual boundary shifts were identified in the first running of the experiment, we used one-tailed *t*-tests to test the significance of each word pair in the second running, as shown in (4). The familywise error was not adjusted since all the tests were all pre-planned. From the results of the first running, the first four word pairs were predicted to have a negative mean % difference, while the last four were predicted to have a positive difference.

(4)	pair	mean % diff.	s.d.	*t*(25)	one-tailed *p*
a.	me**d**aru-ku**d**aru	-6.93	5.11	6.91	< 0.001
b.	ma**g**u-to**g**u	-2.58	7.18	1.83	= 0.04
c.	ne**g**a-sa**g**a	-1.82	5.04	1.84	= 0.04
d.	na**b**i-ta**b**i	0.10	3.47	-0.15	= 0.56
e.	sara**d**a-hara**d**a	1.60	5.18	1.58	= 0.06
f.	kura**b**u-nara**b**u	2.69	7.48	1.83	= 0.04
g.	sahar**a**-nohar**a**	2.13	6.59	1.65	= 0.06
h.	nas**a**-mos**a**	5.09	7.94	3.27	< 0.01

Of the four word pairs whose difference was predicted to be positive, all four were positive. Two were significant (f, h), and the other two were marginally significant (e, g). Of the four word pairs whose difference was predicted to be negative, three were indeed negative (a-c) and the fourth was very close to zero (d). All three negative differences were significant.

The positive shifts for the stimuli that had manipulated C_3 consonant and word-final [a] (4e-h) directly support the hypothesis that foreign items are stored in the lexicon as foreign, even in the absence of phonotactic cues: listeners are biased against perceiving phonotactically illegal segments in native items. The negative bias for the stimuli that had C_2-continuum (4a-d) may likewise support this hypothesis through the morphological process of emphatic gemination.

Finally, overall there seems to be little evidence that word length or the position of the manipulated segment in the word influences the observed perceptual bias caused by lexical strata. Stratal bias proved to be a robust effect observed across these conditions.

3. Concluding remarks

The results of the experiment, along with Moreton and Amano (1999), support the strong stratification hypothesis by Itô and Mester over the representational hypothesis and the weak stratification hypothesis. Our experiment identified two instances of stratal bias even in the absence of phonological cues to stratal affiliations. First, listeners are less likely to identify voiced geminates and final long [a:] in native words than in foreign words. Second, in identifying the geminacy status of C_2, listeners are more likely to identify voiced geminates in native words than in foreign words. The observed biases are incompatible with the representational approach, because in the absence of different phonological behavior, items would not be distinguished representationally in the lexicon. The results are also at odds with the weak stratification hypothesis, which asserts that without alternation evidence, lexical items would be treated alike. In conclusion, our experiment shows the psychological reality of the stratified lexicon, adding substantially to the strength and plausibility of the strong lexical stratification hypothesis.

The last remaining question is how to reconcile the assertion that the grammar is unable to learn the stratal affiliation of words which lack alternation evidence (see especially Ota 2004), with our results that native speakers

perceive words as belonging to certain strata even when they lack phonological cues to their stratal affiliation. First, there is a recent proposal by Pater (2005) suggesting that lexical stratification may be learnable from static phonotactics. Shaw (2006) likewise proposes a learning algorithm which learns stratal stratification without evidence from alternations. Furthermore, we suspect that speakers may learn the stratal affiliation of unmarked words through extragrammatical means as well.

For example, Japanese speakers may learn the stratal affiliation of each item through explicit education, especially when they learn Japanese orthography. Different lexical classes are written with different sets of letters and therefore, it is necessary for Japanese students to learn stratal affiliations of each lexical item. One question that arises now is whether perceptual biases in our experiment could have arisen due to the orthographic differences we used in the experiment. We believe that even if orthography cued the affiliation, the listeners were still actively using the knowledge of stratal affiliations—and different phonological properties of different strata—when categorizing the stimuli. Furthermore, Gelbart (2005) obtained similar bias effects due to lexical strata in languages that do not use different orthographic systems for different strata, which suggests that orthography cannot be the whole story.

Another kind of information which may aid the learning of stratal affiliation is morphological: native verbs and foreign verbs inflect differently (foreign verbs in general inflect with the light verb *-suru*). This sort of morphological information might help Japanese speakers to notice that, for example, *togu* is a native word, as it inflects as a native item. A homonym pair like *aisu* ('to love', native verb) vs. *aisu* ('ice cream', foreign noun) can be distinguished in the same way: only the former item inflects like a native verb. However, this sort of morphology-based learning of stratification is of no help for nouns, because they do not inflect at all. For example, morphology does not provide any evidence that *aisu* 'ice cream' belongs to the foreign stratum, rather than the native stratum; nor does it distinguish the stratal affiliations of pairs like *mosa-nasa*, as neither shows inflection.

We conclude our overall discussion with a thought experiment. Imagine a speaker of a language that inherited the same surface forms as Japanese, but the speaker did not inherit the particular set of prejudices that Japanese speakers impose on their surface forms. Such a speaker would most probably not exhibit the kind of perceptual bias that Japanese speakers exhibited in our experiment, because the bias depends on knowing that words like *kurabu* are foreign and a belief that there is something special about such words. Yet such knowledge and prejudices are not only part of the Japanese language, but are actually used by native speakers to resolve ambiguities in the sound signal.

References

Amano, Shigeaki and Tadahisa Kondo (2000). *Nihongo no Goitokusei*. Tokyo: Sanseido.

Fukazawa, Haruka, Mafuyu Kitahara, and Mitsuhiko Ota (2002a). Constraint-based modeling of split phonological systems. *Phonological Studies* 5: 115-120.

Fukazawa, Haruka, Mafuyu Kitahara, and Mitsuhiko Ota (2002b). Acquisition of

phonological sublexica in Japanese. *Proceedings of Tokyo Conference on Psycholinguistics* 3: 97-113.
Gelbart, Ben (2005). *Perception of Foreignness*. Doctoral dissertation, University of Massachusetts, Amherst.
Inkelas, Sharon, Orhan Orgun, and Cheryl Zoll (1997). The implications of lexical exceptions for the nature of grammar. In Iggy Roca (ed.), *Derivations and Constraints in Phonology*. Oxford: Oxford University Press, pp. 393-418.
Itô, Junko and Armin Mester (1995). Japanese phonology. In John Goldsmith (ed.), *The Handbook of Phonology*. Oxford: Blackwell, pp 817-847.
Itô, Junko and Armin Mester (1999). The phonological lexicon. In Natsuko Tsujimura (ed.), *The Handbook of Japanese Linguistics*. Oxford: Blackwell, pp 62-100.
Itô, Junko and Armin Mester (2001). Covert generalizations in Optimality Theory: The role of stratal faithfulness constraints. *Studies in Phonetics, Phonology and Morphology* 7: 3-33.
Itô, Junko, Armin Mester, and Jaye Padgett (1999). Lexical classes in Japanese: A reply to Rice. *Phonology at Santa Cruz* 6: 39-46.
Kawahara, Shigeto (2001). *Similarity among Variants: Output-Variant Correspondence*. BA Thesis, International Christian University.
Kawahara, Shigeto (2006). Mimetic gemination in Japanese: A challenge for Evolutionary Phonology. *Theoretical Linguistics* 32: 411-424.
Massaro, Dominic and Michael Cohen (1983). Phonological context in speech perception. *Perception and Psychophysics* 34: 338-348.
Maye, Jessica, Keiichrio Suzuki, and Kazutoshi Ohno (2000). The psychological reality of lexical stratification in Japanese. Ms., University of Rochester.
Moreton, Elliot and Shigeaki Amano (1999). The effect of lexical stratum phonotactics on the perception of Japanese vowel length. *Proceedings of the 6th European Conference on Speech Communication and Technology*.
Moreton, Elliot, Shigeaki Amano, and Tadahisa Kondo (1998). Statistical phonotactics in Japanese. *Transactions of Technical Committee on Psychological Acoustics,* H-98-120.
Ota, Mitsuhiko (2004). The learnability of the stratified phonological lexicon. *Journal of Japanese Linguistics* 20: 19-40.
Pater, Joe (2005). Learning a stratified grammar. In *Proceedings of the 29th Boston University Conference on Language Development*. Somerville: Cascadilla Press, pp. 482-492.
Rice, Keren (1997). Japanese NC clusters and the redundancy of postnasal voicing. *Linguistic Inquiry* 28: 541-551.
Shaw, Jason (2006). Learning stratified lexicon. *Proceedings of NELS 36*.
Tateishi, Koichi (2003). Phonological patterns and lexical strata. *Proceedings of the 7th International Congress of Linguistics* (CD-ROM). Prague: Matfyz Press.

Department of Linguistics
South College 226, University of Massachusetts, Amherst
Amherst MA 01003, U.S.A.

gelbartb@yahoo.com & *kawahara@linguist.umass.edu*