Available online at www.sciencedirect.com
ScienceDirect
Cognition 104 (2007) 591–630
COGNITION
www.elsevier.com/locate/COGNIT

What we know about what we have never heard: Evidence from perceptual illusions ☆

Iris Berent [a,*], Donca Steriade [b], Tracy Lennertz [a], Vered Vaknin [c]

[a] *Department of Psychology, Florida Atlantic University, 777 Glades Road, Boca Raton, FL 33431-0991, United States*
[b] *Department of Linguistics and Philosophy, Massachusetts Institute of Technology, Cambridge, MA, USA*
[c] *Faculty of Education, University of Haifa, Haifa, Israel*

Received 3 June 2005; revised 24 May 2006; accepted 28 May 2006

Abstract

Are speakers equipped with preferences concerning grammatical structures that are absent in their language? We examine this question by investigating the sensitivity of English speakers to the sonority of onset clusters. Linguistic research suggests that certain onset clusters are universally preferred (e.g., *bd* > *lb*). We demonstrate that such preferences modulate the perception of unattested onsets by English speakers: Monosyllabic auditory nonwords with onsets that are universally dispreferred (e.g., *lbif*) are more likely to be classified as disyllabic and misperceived as identical to their disyllabic counterparts (e.g., *lebif*) compared to onsets that are relatively preferred across languages (e.g., *bdif*). Consequently, dispreferred onsets benefit from priming by their epenthetic counterpart (e.g., *lebif–lbif*) as much as they benefit from identity priming (e.g., *lbif–lbif*). A similar pattern of misperception (e.g., *lbif* → *lebif*) was observed among speakers of Russian, where clusters of this type occur. But unlike English speakers, Russian speakers perceived these clusters accurately on most trials, suggesting that the perceptual illusions of English speakers are partly due to their linguistic experience, rather than phonetic confusion alone. Further evidence against a purely phonetic explanation for our results is offered by the capacity of English speakers to perceive such onsets accurately under

---

☆ This manuscript was accepted under the editorship of Jacques Mehler.
* Corresponding author. Tel.: +1 561 297 2905; fax: +1 561 297 2160.
E-mail address: iberent@fau.edu (I. Berent).

doi:10.1016/j.cognition.2006.05.015

conditions that encourage precise phonetic encoding. The perceptual illusions of English speakers are also irreducible to several statistical properties of the English lexicon. The systematic misperception of universally dispreferred onsets might reflect their ill-formedness in the grammars of all speakers, irrespective of linguistic experience. Such universal grammatical preferences implicate constraints on language learning.

Keywords: Phonology; Markedness; Sonority; Typology; Optimality theory; Speech perception

---

1. Introduction

1.1. Markedness in typology and individual grammars

There is ample evidence that speakers (including young infants) are sensitive to the statistical structure of linguistic input (e.g., Dell, Reed, Adams, & Meyer, 2000; Saffran, Aslin, & Newport, 1996). Whether statistical learning is *sufficient* to account for phonological competence is far less clear: Is linguistic knowledge fully explicable by the properties of linguistic tokens, or are speakers equipped with preferences concerning the representation and processing of language–preferences that shape learning, but are irreducible to their experience with specific linguistic expressions? Do speakers possess knowledge concerning linguistic structures that they have never encountered?

These questions can be addressed by comparing the linguistic preferences of individual speakers to typological surveys. Typology, the comparative study of linguistic systems, often reveals universal laws. A typological survey of languages along any given structural dimension shows that certain structural variants (e.g., structure A) are more frequent than others (e.g., structure B). Moreover, the existence of infrequent structures implicates common ones: If a language tolerates the relatively infrequent structure B, it is likely to allow the more frequent structure A. For example, any language that permits a complex onset (e.g., *drug*) also permits simple onsets (e.g., *rug*, see Blevins, 1995; Greenberg, 1978; Prince & Smolensky, 1993/2004). The rare variants (e.g., structure B) are called *marked* whereas the more frequent ones are *unmarked* (e.g., structure A). Such laws reflect universal regularities in the distribution of linguistic structures. Of interest is whether these typological universals reflect the inherent linguistic preferences of *individual* speakers.

A long line of linguistic research (Chomsky & Halle, 1968; Jakobson, 1941; Trubetzkoy, 1938/1958) links typology to individual linguistic competence. Building on that tradition, Prince and Smolensky (1993/2004), propose that structures that are typologically marked (i.e., dispreferred across languages) are also grammatically marked–such structures violate a set of grammatical restrictions, called markedness constraints. These two meanings of "markedness" are causally linked: Typological markedness is a consequence of grammatical markedness in the linguistic competence of individual speakers. Grammars of distinct languages differ only on whether they tolerate marked structures. However, *all* grammars express preferences for unmarked over marked structures. For example, the typological preference for simple onsets (e.g., *rug*) over complex onsets (e.g., *drug*) is attributed to a principle that marks complex onsets as dispreferred in all grammars, including the grammars of English (a language that tolerates complex onsets such as *drug*) and Japanese (a language that allows only simple onsets like *rug*). To generalize, if B is marked and A is unmarked (in the typological, distributional sense) then, on this view, all grammars contain statements that identify A as preferred to B, irrespective of whether either or both structures are attested in the relevant language.

The following research takes some initial steps to investigate this proposal (for related work, see Davidson, 2000; Moreton, 2002; Wilson, 2003, in press; Zuraw, 2005). To this end, we explore speakers' sensitivity to the markedness of structures that are absent in their lexicon. If speakers are equipped with markedness preferences, such preferences should generalize to unattested structures. To illustrate our approach, let us assume that a phonological structure A is unmarked (i.e., preferred) relative to structure B across languages, but neither A or B are attested in the English lexicon. Of interest is whether English speakers nonetheless prefer A to B. To the extent that such preferences are found, and they are inexplicable by the statistical characteristics of attested exemplars and the acoustic properties of the input, they could potentially reflect on inherent preferences of the language system. Our goal here is to probe for such preferences.

We chose to assess the scope of markedness constraints using the restrictions on onset clusters as a case study. Cross-linguistic surveys (e.g., Greenberg, 1978) suggest that certain onset clusters as in *blif* are preferred to certain others, as in *bnif*, which, in turn, are preferred to onsets such as in *bdif*. At the bottom of the hierarchy is the type of onsets found in *lbif*. English tolerates onsets like *blif*, but it offers speakers little evidence as to the remainder of the hierarchy. Of interest is whether speakers of English are sensitive to the hierarchy of onsets that are absent in their lexicon. We begin by documenting some cross-linguistic preferences concerning the co-occurrence of consonants in the onset. We next examine whether English speakers generalize such preferences to onsets that are unattested in their language.

1.2. The markedness of onset clusters in typology and grammar

Languages constrain the co-occurrence of segments in the syllable. For instance, English speakers accept *plin*, but not *lpin* as a possible word. Such preferences have been analyzed as bearing on the sonority profile of the syllable (Blevins, 1995; Clements, 1990). Sonority is a scalar property of segments correlated with acoustic intensity (i.e., loudness, Ladefoged, 1975; Ohala, 1990, for critical discussion; Parker, 2002, for details). Louder segments (e.g., *l*) are more sonorous than quieter segments (e.g., *p*, *t*). Segments are arrayed on the sonority scale as follows:

(1) A sonority scale

Vowels/glides (5)	Liquids (4)	Nasals (3)	Fricatives (2)	Stops (1)
a, e/y, w	*l, r*	*n, m*	*s, sh, z, f, v, th*	*p, b, t, d, k, g*
More sonorous				Less sonorous

One can formulate general conditions on the sonority profile of the syllable relative to this scale. In an English syllable like *plank*, the sonority level rises abruptly from [p] to [l], it continues to rise to [a], and then steadily declines. We are concerned here with the large initial rise in the onset [pl]: this sonority profile is typical of English onsets. English biconsonantal onsets manifest systematic combinations of obstruents (i.e., stops and fricatives) with either liquids (e.g., ***pl**ay*, ***dr**ive*) or glides (e.g., ***cu**te*, ***sw**eet*)– both manifesting an abrupt sonority rise (Hammond, 1999, pp. 51–56). A third type of onsets, including s-stop combinations (i.e., onsets of falling sonority), presents a systematic exception, since *s* is the only consonant that forms a falling sonority onset in English and many other languages (for different views on the status of sC onsets, see Blevins, 1995; Kiparsky, 1979; Selkirk, 1982; Wright, 2004).[1] Leaving the *s*C onsets aside, there is a consensus that English requires a large sonority rise in its onset (e.g., Blevins, 1995; Clements, 1990; Kenstowicz, 1994; Selkirk, 1982). Although onsets with a large sonority rise are most common across languages, other sonority profiles are tolerated (for illustration, see Table 1). For example, Ancient Greek manifests onsets with small rise (obstruent–nasal onsets; e.g., *pneuma*, "breath"); Hebrew tolerates sonority plateaus (e.g., *ptil*, "wick"), whereas Russian even allows sonority falls (e.g., *rzhan*, "zealous; Halle, 1971). Interestingly, however, languages that tolerate rare profiles tend to tolerate more frequent ones as well. These regularities can be described in the following three implicational statements about onsets:

(2) Implicational universals regarding sonority profiles in *typology*. In any given language:
(a) The presence of a small sonority rise in the onset implies that of a large one.
(b) The presence of a sonority plateau in the onset implies that of some sonority rise.
(c) The presence of a sonority fall in the onset implies that of a plateau.

The statements in (2a–c) are hypotheses regarding the typological contingencies among sonority-profiles. Statement (2a) maintains that languages allowing small rises like [pn], should also allow large rises, as in [pl]; languages that allow small, but not large rises should be rare. Likewise, the contingency between sonority plateau and sonority rise in (2b) states that if [pt] – a plateau – is a possible onset, so is [pl] or [pn]–so in principle, counter examples (e.g., languages that allow [pt] to the exclusion of [pl] or [pn]) should not occur. Finally, according to (2c), languages that

---

[1] Because s-*stop* sequences violate the sonority restrictions of numerous languages, we will not consider them further.

Table 1
The sonority profile of various onset clusters

Onset cluster	Sonority level of C1	Sonority level of C2	Sonority profile of onset
blif	1	4	Large rise
bnif	1	3	Small rise
bdif	1	1	Plateau
lbif	4	1	Fall

allow a falling-sonority onset like [lp] should allow a plateau onset like [pt]. English disallows both, Russian allows both, while Greek allows only plateaus and rises, but not falls in word-initial onsets. Together, (2a–c) predict that smaller sonority rises (plateaus and falls) imply larger rises.

1.2.1. The typological markedness of sonority profiles: Evidence from a survey of onset clusters

Typological hypotheses such as (2) are of interest to us in so far as they reflect on preferences in the grammar of all individual speakers. Before we proceed to define and test speakers' preferences, it is first necessary to verify that the typological generalizations in (2a–c) are, in fact, true. Although some version of (2a–c) has been assumed by many linguists (cf. Blevins, 1995; Smolensky, 2006; Steriade, 2003), no survey has been carried out specifically to test these generalizations. As an initial test of (2), we turned to an existing, widely cited survey of onset clusters by Greenberg (Greenberg, 1978). Because he was not testing any explicit hypothesis about sonority profiles, Greenberg's conclusions do not address (2a–c) directly. However, the data described in his survey allows one to reconstruct the implicational relations among onsets with a sonority rise (either small or large), plateau and fall, along the lines of (2a–c).[2] This reconstruction is presented in the corresponding Table 2a–c below.

An inspection of Table 2 suggests two general conclusions. First, abrupt sonority rises (e.g., obstruent–liquid, 83% of the sample) are more frequent as word onsets than small rises (64% of the sample), smaller rises are more frequent than sonority plateaus (49% of the sample), and sonority plateaus are more frequent than sonority falls (13% of the sample). Second, if a language tolerates onsets with a rare sonority profile, it is likely to tolerate onsets with more common profiles. For example, consider the contingency between small and large sonority rises. Not only are word onsets with small rises less frequent across languages, but the presence of an onset with a smaller rise in any one language implies the presence of onsets with larger rises. There are no counter-examples to this generalization in Greenberg's survey: Only one language with a small sonority rise lacks a large rise (i.e., it lacks obstruent–liquid onsets), but because that language (Santee Dakota) lacks liquids altogether, it is not a true counter example to 2a. Similar

---

[2] In Greenberg's (1978) sample, onsets with large sonority rises are obstruent–liquid clusters; onsets with smaller rises are nasal–liquids and obstruent–nasals; onsets with sonority plateaus are stop–stop and fricative–fricative onsets, and onsets with sonority falls are liquid–obstruent and liquid nasal clusters.

Table 2
The contingency between small sonority rise and larger sonority rise (a); sonority plateau and sonority rise (b); and sonority fall and sonority rise (c)

a.		Large rise	
		+	-
Small rise	+	57	1
	-	18	14

b.		Rise	
		+	-
Plateau	+	41	3
	-	35	11

c.		Plateau	
		+	-
Fall	+	11	1
	-	33	45

The presence of a cluster is indicated by +, whereas its absence is indicated by −. Data from Greenberg (1978).

implicational links are present between sonority falls and plateaus and between sonority plateaus and rises. The frequency of large sonority rises and their implicational dependence on profiles of smaller (or no) rise is consistent with their hypothesized status as relatively unmarked.

To assure that the contingency of frequent (i.e., unmarked) sonority profiles on the presence of infrequent (i.e., marked) ones is not merely due to the overall preponderance of frequent profiles, we evaluated the statistical significance of these contingencies. Consider first the contingency of large sonority rises on the presence of small rises in (2a). We first assessed the statistical significance of this contingency by means of a contingency table. The overall contingency was significant ($\chi^2(1) = 23.28$, $p < .0001$, with continuity correction). This result suggests a dependency between these two profiles, but it does not specify its precise nature. In particular, this result cannot distinguish the possibility that common (i.e., unmarked) cases depend on the presence of rare (i.e., marked) ones – the principal prediction of the markedness account – from alternative contingencies that are not predicted (e.g., that unmarked cases imply marked ones). We next tested the specific hypothesis that the presence of a marked, small sonority rise implies the presence of an unmarked, larger one by means of a binomial coefficient. As predicted in (2a), the probability that most languages with a small sonority rise also manifest a large rise (.98 of the cases) was inexplicable by the share of large rises in the sample (.83, $p = .01$, testing for the binomial co-efficient).

Similar results were obtained for the contingency between sonority rises and sonority plateaus (2b). The contingency between these two profiles was marginally significant ($\chi^2(1) = 3.79$, $p < .06$, with continuity correction). The probability that most languages with sonority plateaus have sonority rises (.93 of the cases) was inexplicable by the share of large rises in the sample (.84, $p = .049$, testing for the binomial co-efficient). Finally, we tested the contingency between sonority falls and sonority plateaus (2c). The overall contingency between sonority falls and sonority plateaus was significant ($\chi^2(1) = 8.26$, $p < .005$, with continuity correction). With the exception of a single language (Chatino), all languages with sonority falls manifested plateaus. A test of the binomial co-efficient specifically showed that the probability that most languages with sonority falls have sonority plateaus (.92) is inexplicable by the overall probability of sonority plateaus in the sample (.49, $p = .002$).

1.2.2. The markedness of sonority profiles: from typology to grammar

The typological evidence derived from Greenberg's (1978) sample supports the contingency between sonority falls, plateaus and rises described in (2a–c). Not only is the statistical contingency between marked and unmarked profiles significant, but the number of counter examples is remarkably low (virtually no counter examples were found to 2a, one counter example to 2c, and three to 2b). The observation of such reliable contingencies in the typology sets the stage for examining whether they reflect the presence of universal preferences encoded in the grammars of individual speakers. We ask whether the generalizations in (2a–c) have a counterpart in the grammar of individual speakers, as in (3)[3]:

(3) The markedness of sonority profiles in the grammars of *individual* speakers.
  a. Small sonority rises in the onset are more marked than large rises.
  b. Sonority plateaus in the onset are more marked than rises.
  c. Sonority falls in the onset are more marked than plateaus.

Suppose further that these markedness differences among onsets are potentially present in every grammar, regardless of what types of onset clusters occur overtly. This could be either because the preferences in (3) are innate (Prince & Smolensky, 1993/2004) or because speakers are equipped with mechanisms that allow them to induce (3) from the information available to them (e.g., information regarding the perception and production of such sound sequences, see Hayes & Steriade, 2004). On either of these interpretations, these preferences among onset types will be potentially available to all speakers, even if the relevant onsets are absent in their language. Under this scenario, even if speakers are exposed overtly only to [pl]-onsets, they should be sensitive to the markedness difference between [pn]–[pt] and [pt]–[lp] onsets. This is the possibility investigated here.

---

[3] The statements in (3) express grammatical preferences – they do not constitute a formal account of how such preferences are encoded by the grammar. For different views regarding the mechanisms that underlie these differences in marked status, cf. Blevins (1995), Gouskova (2002), Hammond (1999), Smolensky (2006), and Steriade (1982).

1.3. Are speakers sensitive to the markedness of onset clusters?

To probe for such universal grammatical preferences, it is necessary to determine whether speakers generalize them to onset clusters that are unattested in their language. Although much research examined sonority-related preferences for attested clusters (e.g., Gierut, 1999; Ohala, 1999; Romani & Calabrese, 1998; Treiman, 1984; Treiman, Bowey, & Bourassa, 2002; Treiman & Danis, 1988; Treiman, Straub, & Lavery, 1994), the evidence with respect to preferences for unattested clusters is limited. A few studies have examined the effect of sonority on the production of unattested onsets, but their results are unclear. Pertz & Bever (1975) compared the responses of English speakers to unattested onsets which participants were instructed to articulate. They found that marked onsets were judged as less likely to occur than unmarked onsets. Likewise, Broselow & Finer (1991) reported that Korean and Japanese speakers (languages manifesting only obstruent-*y* onsets and lacking the phoneme *f*) produced English onsets with large sonority rises (e.g., [fj] in *fuse*) more accurately than onsets with smaller rises (e.g., [fr] in *frugal*). Using different sequences, however, Davidson (2000) found no systematic effects of sonority profile in the production of novel onset clusters by English speakers. Thus, the few existing results regarding the production of unattested onsets are inconsistent. Moreover, the preferences inferred from tasks that require the production of unattested clusters might be due to articulatory limitations, rather than grammatical markedness. It is thus desirable to complement these findings by assessing the effect of sonority on perception.

How should the markedness of an onset cluster affect its perception? Prior research suggests that, in the case of extremely marked structures, the representation of the input may be systematically distorted. Pitt (1998) showed that English speakers confuse marked illicit clusters (e.g., *tla*) with their epenthetic counterparts (e.g., təla). Additional perceptual difficulties with illicit syllables are observed with speakers of English (Massaro & Cohen, 1983; Moreton, 2002) French (Hallé, Segui, Frauenfelder, & Meunier, 1998) and Japanese (Dupoux, Kakehi, Hirose, Pallier, & Mehler, 1999; Dupoux, Pallier, Kakehi, & Mehler, 2001). Because the perceptual difficulties with marked, unattested syllables are typically assessed relative to syllables that are both unmarked and attested, it is conceivable that they might be due to the unfamiliarity with marked structures. However, linguistic analysis suggests a different possibility. In this view, perceptual illusions might reflect the organization of the grammar: Highly marked inputs are repaired in perception to abide by the grammatical restrictions of the language.

In the framework of Optimality theory (Prince & Smolensky, 1993/2004), the grammar includes two types of constraints: markedness constraints (expressing structural restrictions on outputs, e.g., “avoid the output *tl*”) and faithfulness constraints (e.g., which require identity between outputs and inputs, e.g., map the input *tl* to the output *tl*). The linguistic representation inferred by listeners is determined by the ranking of these two sets of constraints (Smolensky, 1996). Given relatively unmarked structures (e.g., *dr*), the faithfulness to the input is likely to be ranked above the avoidance of marked outputs, hence inputs are encoded accurately. In contrast, with marked structures, the avoidance of marked inputs might outweigh their faithful encoding. Accordingly, such inputs are less likely to be represented in a way that faithfully matches the acoustic properties of the stimulus. Instead, they would be systematically recoded as less marked structures. The grammatical recoding of marked onsets should trigger systematic illusions in their perception.

The present research exploits such perceptual illusions to investigate speakers' sensitivity to the markedness of onsets that are unattested in English. Previous research (e.g., Dupoux et al., 1999; Pitt, 1998) shows that syllables that are both marked and unattested in one's language are more likely to trigger perceptual illusions compared to syllables that are both unmarked and attested. Our investigation probes for such illusions among syllables that are *all* unattested (e.g., syllables with the onsets *bn*, *bd*, *lb*). Moreover, these unattested syllables instantiate a hierarchy of structural types (i.e., sonority profiles) that are not systematically allowed in English (i.e., small sonority rises, plateaus and falls). Of interest is whether English speakers are sensitive to the markedness of linguistic structures that they have never heard before. If speakers are sensitive to the markedness of unattested sonority profiles, then marked structures should be more likely to undergo repair compared with unmarked ones. In production, English speakers are known to repair highly marked clusters by means of vowel epenthesis – a process that inserts a vowel between adjacent consonants (e.g., *tla* → *tela*). Such repairs are utilized for this purpose in inflection (e.g., /buʃz/ → [buʃəz]; Anderson, 1974, 54ff) and on-line loan adaptation (Davidson & Stone, 2004). Assuming a single grammar for perception and production (Smolensky, 1996), we expect similar repair strategy in the *perception* of marked onsets. Consequently, English speakers should fail to distinguish marked onsets from their epenthetic counterparts.[4]

Participants in our experiments are presented with monosyllabic nonwords with three types of onset clusters that are unattested in English: Onsets with small sonority rises, plateaus or falls (e.g., *bnif*, *bdif*, *lbif*). In Experiments 1–2, participants judge the number of syllables in the input, whereas in Experiments 3–4, they determine whether monosyllabic inputs are identical to their disyllabic counterparts (e.g., is *lbif* identical to *lebif*?). Experiments 5–6 compare the potential of marked and unmarked structures to exert identity priming. If marked unattested clusters are repaired in perception (e.g., *lbif* → *lebif*), then as the markedness of the onset cluster increases, speakers should be more likely to consider the cluster as disyllabic (in Experiments 1–2) and misperceive it as identical to its disyllabic counterpart (in Experiments 3–4). Consequently, marked onsets should benefit from priming by their epenthetic counterparts (e.g., *lebif–lbif*) as much as they benefit from identity priming (e.g., *lbif–lbif*, in Experiment 5).

Our investigation also seeks to determine the source of such perceptual illusions. We consider the possibility that the difficulties with marked onsets might

---

[4] In what follows, we use an orthographic transcription such as *tela* to indicate the phonetic representation [təla].

be due to the statistical properties of the English lexicon. We also assess whether such errors are due to the acoustic properties of the different clusters, rather than grammatical repair – the active recoding of an intact phonetic percept by the grammar (e.g., [lbɪf] → /ləbɪf/). We address this possibility in two ways. First, we examine whether marked clusters can be accurately perceived by speakers of Russian – a language that tolerates all cluster-types used in our experiments. If the tendency of English speakers to misperceive marked monosyllabic nonwords is due to their linguistic knowledge (both universal and language-specific), then the performance of English and Russian speakers should diverge: Although both groups might be sensitive to markedness, Russian speakers should perceive most marked clusters accurately (as a result of their specific experience with these types of clusters). A second test of the phonetic explanation evaluates the ability of English speakers to encode marked clusters. Here we examine whether the perceptual illusions with marked onsets can be eliminated under conditions that encourage precise phonetic encoding of the input. If English speakers can encode the acoustic properties of marked onsets accurately, then marked and unmarked onsets might yield similar performance when attention to their phonetic form is enhanced. The demonstration of systematic markedness preferences that are inexplicable by the statistical properties of the lexicon or phonetic confusions implicates inherent constraints on the organization of the language system.

2. The effect of markedness on syllable judgment

2.1. Experiment 1

In Experiments 1–2, participants are asked to quickly indicate whether the auditory stimulus has one or two syllables. If English speakers are sensitive to the markedness of unattested clusters, and if marked clusters are repaired by epenthesis (e.g., *lbif* → *lebif*), then accuracy should be inversely related to markedness: As the markedness of a monosyllabic cluster increases, speakers should be more likely to (incorrectly) classify it as disyllabic. The dispreference for marked sonority profiles should exert the opposite effect on the perception of the disyllabic counterparts. In our forced choice task between *lbif* and *lebif*, aversion to the marked monosyllabic stimulus (e.g., *lbif*) might increase the likelihood of interpreting the acoustic input as consistent with the disyllabic counterpart (e.g., *lebif*). Accordingly, the perception of disyllabic forms should become more accurate as the markedness of their monosyllabic counterpart increases.

2.2. Method

2.2.1. Participants

Sixteen native English speakers, students at Florida Atlantic University took part in the experiment in partial fulfillment of a course requirement.

2.2.2. Materials

The experimental materials consisted of 90 pairs of CCVC monosyllabic and CəCVC disyllabic nonwords, presented aurally (see Appendix A). The monosyllabic nonwords had an unattested onset cluster, either sonority rises (mostly obstruent–nasal or obstruent–liquid combinations), plateaus (obstruent–obstruent combinations) or falls (liquid–obstruent or nasal–obstruent combinations). The three types of onsets were arranged in triplets, matched for the rhyme (e.g., *bnif*, *bdif*, *lbif*). These nonwords were matched to disyllabic counterparts that differed only on the presence of a schwa between the onset consonants (e.g., *bənif*, *bədif*, *ləbif*). To encourage participants to treat the experimental materials as English, they were presented mixed with 30 pairs of fillers with attested English onsets and their disyllabic counterparts (e.g., *blif*, *belif*). Because responses to items with attested onsets might be due to familiarity with their onsets, we did not include these items in the analyses.

To estimate the phonotactic properties of our nonwords with unattested onsets, we calculated the position-specific probability of each of their four segments (i.e., the probability that a given phoneme occurs at a specific position along the array of an item's four phonemes in all English words, regardless of word-length), as well as their position-specific bi-phone probability in the word (i.e., the probability that a given two-phone combination occurs at a given word position in all English words, regardless of word-length Vitevitch & Luce, 2004; http://www.people.ku.edu/~mvitevit/PhonoProbHome.html). The means provided in Table 3 were computed by summing the probability of each segment and each bi-phone across the four segments in the word. We also estimated the number of phonological neighbors (the number of words generated by substituting, deleting or adding a single phoneme) and their summed frequency using an on-line database prepared by Mitch Sommers (http://128.252.27.56/Neighborhood/Home.asp). The phonotactic and neighborhood properties were calculated based on an electronic version of the Merriam-Webster Pocket Dictionary of 1964, whereas the frequencies were based on Kucera and Francis (1967). An ANOVA suggested that the three types of nonwords with unattested onsets differed only on the position-specific probability of their phonemes ($F(2, 58) = 6.43$, $p < .004$). Planned comparisons of sonority rises and plateaus, and sonority plateaus and falls demonstrated that the phoneme probability of onsets with sonority rises was significantly higher than that of the

Table 3
The statistical properties of the monosyllabic nonwords used in Experiments 1–4

Sonority profile	Phoneme probability		Biphone probability		Number of neighbors		Neighbors' frequency	
	M	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*
Rise (e.g., *bnif*)	.1426	.0454	.0094	.0398	2.03	3.95	33.8	55.21
Plateau (e.g., *bdif*)	.1265	.0328	.0016	.0013	1.50	1.91	36.00	49.83
Fall (e.g., *lbif*)	.117	.117	.0018	.0016	1.17	1.05	50.77	126.72

plateaus ($F(58) = 2.26$, $p < .04$). No other planned contrasts were significant (all $p > .05$).

The monosyllabic and disyllabic materials were recorded by a native Russian speaker: because Russian manifests such cluster-types, they can be naturally produced by Russian speakers. To assure that the speaker considered these nonwords as Russian-like stimuli, they were presented in a Cyrillic alphabet embedded in a context ("X-raz", "once X").

2.2.3. Procedure

Each trial began with a fixation (*) and a message indicating the trial number. Participants initiated the trial by pressing the space bar, which resulted in the presentation of an auditory stimulus. Participants indicated whether the stimulus had one syllable or two using the 1 and 2 keys, respectively. To familiarize participants with the task, they were first given practice categorizing existing English words (e.g., *polite*, *plight*). In this and subsequent experiments, trial order was randomized.

2.3. Results

2.3.1. Response accuracy

Mean response accuracy is provided in Fig. 1. An inspection of the means suggests that, as the markedness of the monosyllabic form increased, monosyllabic inputs were categorized less accurately, whereas their disyllabic counterparts yielded more accurate responses. An ANOVA (2 syllables × 3 cluster type) yielded a significant interaction ($Fs(2, 30) = 32.82$, $p < .0001$; $Fi(2, 58) = 59.15$, $p < .0001$).

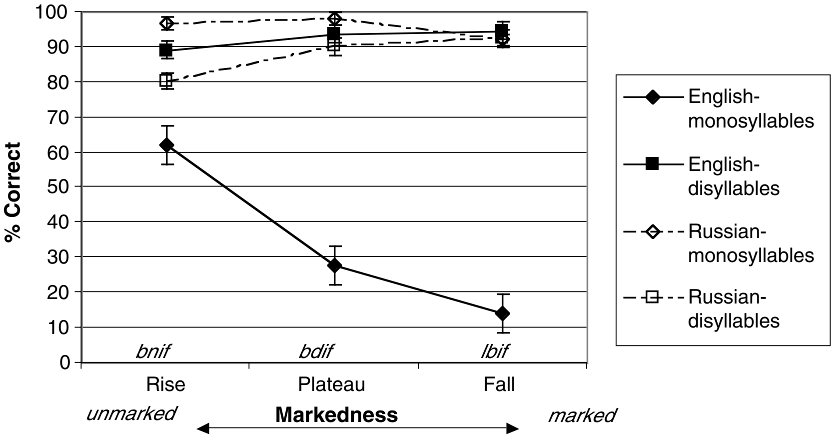

Fig. 1. Mean response accuracy of English and Russian speakers to monosyllabic nonwords and their disyllabic counterparts in Experiments 1–2 as a function of the number of syllables and the markedness of the monosyllabic counterpart. Error bars represent the confidence interval constructed for the difference among the means.

The simple effect of cluster type was significant for monosyllabic inputs (Fs(2, 30) = 42.25, $p < .0002$; Fi(2, 58) = 46.73, $p < .0002$) and marginally significant for disyllabic ones (Fs(2, 30) = 2.93, $p < .07$; Fi(2, 58) = 6.89, $p < .003$).

The effect of cluster type was further investigated separately for monosyllabic and disyllabic inputs by means of planned contrasts. With monosyllabic inputs, onsets with sonority rises elicited more accurate responses than sonority plateaus (ts(30) = 6.38, $p < .0001$; ti(58) = 6.72, $p < .0001$), which, in turn, yielded more accurate responses than sonority falls (ts(30) = 2.53, $p < .02$; ti(58) = 2.65, $p < .02$). The disyllabic counterparts yielded the opposite pattern: disyllabic counterparts of clusters with sonority rises elicited more errors than the counterparts of plateaus (ts(30) = 1.92, $p < .07$; ti(58) = 2.96, $p < .005$). Response accuracy for the disyllabic counterparts of sonority plateaus and falls did not differ (both $t < 1$).

2.3.2. Response time

Correct responses falling 2.5 *SD* above or below the mean (3.3% of the correct responses) were eliminated from the analysis of response latency. The means are presented in Table 4. An ANOVA (2 syllable × 3 cluster type) did not yield a significant interaction (Fs(2, 24) = 1.08, $p < .35$, n.s.; Fi(2, 46) = 1.17, $p < .32$ n.s.) nor were any of the planned comparisons significant.

2.4. Discussion

The findings of Experiment 1 suggest that onsets that are unattested in English are not all perceived alike: Highly marked clusters with sonority falls (e.g., *lbif*) triggered more errors than sonority plateaus (e.g., *bdif*), which, in turn, triggered more errors than sonority rises (e.g., *bnif*). Conversely, the disyllabic counterparts of marked clusters (e.g., *bedif*) were perceived more accurately than the counterparts of less-marked onsets (e.g., *benif*). As markedness increases, monosyllabic onsets appear to be confused with their disyllabic counterparts. In fact, onsets with sonority plateaus (ts(15) = 3.85, $p < .002$, ti(29) = 7.19, $p < .002$) and falls (ts(15) = 8.42, $p < .0002$, ti(29) = 13.18, $p < .0002$) were categorized as disyllabic on most trials.

Why are English speakers subject to such perceptual illusions? An Optimality-theoretic account attributes these errors to universal constraints on the organization of the grammar: Because the constraints banning highly marked structures (e.g., avoid

Table 4
Mean correct response time (in ms) in Experiment 1 for monosyllabic and disyllabic inputs as a function of the number of syllables and the sonority profile of the monosyllabic counterpart

Cluster type	Monosyllabic		Disyllabic	
	M	*SD*	*M*	*SD*
Rise (e.g., *bnif*)	1581	145	1476	78
Plateau (e.g., *bdif*)	1676	190	1510	96
Fall (e.g., *lbif*)	1702	291	1503	58

lbif) outrank the constraints that ensure their faithful representation (e.g., encode the input *lbif* as *lbif*), people do not represent highly marked structures accurately. Instead, such structures are systematically recoded as structures that are less marked. Since English is known to repair illicit clusters by means of vowel epenthesis, we expected unattested marked onsets to be repaired in the same fashion (e.g., *lbif* → *lebif*). The perception of marked onsets as disyllabic is consistent with this prediction.[5]

An alternative explanation attributes the superiority of unmarked clusters to their statistical lexical properties and to grammatical restrictions that are unrelated to sonority. Greenberg (1978) notes the tendency of onset consonants to avoid sharing

---

[5] Although our results clearly suggest that marked onsets are perceived as disyllabic, previous research indicates that unattested onsets are also confusable with attested clusters (e.g., *tl* → *tr*, e.g., Massaro & Cohen, 1983; Pitt, 1998). To determine the prevalence of epenthetic repair, we asked a group of 12 native English speakers, students at Florida Atlantic University to transcribe our materials using English letters. Speakers rarely produced outputs with an erroneous onset cluster (attested or otherwise) – such responses amounted to only 5.18% of the total observations ($M = 4.26\%$, $M = 0.83\%$, $M = 0.09\%$ for sonority rises, plateaus, and falls, respectively). In contrast, epenthetic responses (e.g., *ptik* → *petik*) were the single most frequent response category in the experiment ($M = 37.4\%$) as well as the most frequent response type within each of the three sonority profiles. An ANOVA suggested that the rate of epenthetic responses was modulated by cluster type in the analysis by participants (Fs(2, 22) = 4.13, $p < .04$; Fi(2, 58) = 2.44, $p < .10$). Epenthetic responses were more frequent for inputs with sonority plateaus ($M = 15\%$) than with sonority rises ($M = 10.5\%$, $p < .05$, Fisher PSLD), but, unexpectedly, they were numerically less frequent with sonority falls ($M = 11\%$). This finding might be due to the prevalence of epenthetic responses that altered the onset consonants (e.g., *ptik* → *pelik*). We therefore collapsed all erroneous disyllabic responses into a single category, including medial epenthesis (*ptik* → *petik*, 37.4% of the responses), initial epenthesis (*ptik* → *eptik*, 1.11% of the responses) and medial epenthesis that altered the onset consonants (e.g., *ptik* → *pelik*, 34.4% of the total responses). We likewise formed a second broad category representing all erroneous monosyllabic responses, including replacement of one of the consonant (e.g., *ptik* → *ktik*, 5.2% of the total responses), consonant deletion (e.g., *ptik* → *pik*, 7.7% of the total responses), and other erroneous monosyllabic responses (2.78% of the total responses).We next considered the effect of onset types on the prevalence of three broad categories of response: correct responses (2.9% of the total responses), erroneous disyllabic responses (73% of the total responses) and erroneous monosyllabic responses (15.5% of the total responses). The remaining responses were lexicalizations and omissions (8.4% of the total responses). The effect of sonority profile was examined separately for each response type using one-way ANOVA's. The effect of onset type was significant for correct responses (Fs(2, 22) = 8.48, $p < .002$; Fi(2, 58) = 7.93, $p < .001$), monosyllabic responses (Fs(2, 22) = 22.93, $p < .0002$; Fi(2, 58) = 11.69, $p < .0002$), and disyllabic responses (Fs(2, 22) = 58.72, $p < .0002$; Fi(2, 58) = 14.007, $p < .0002$). The rate of disyllabic responses was significantly lower in onsets of rising sonority ($M = 18\%$) compared to those with sonority falls and plateaus ($M = 26\%$; $M = 28\%$, respectively, $p < .05$, Fisher PLSD by participants and items). Conversely, onsets with sonority rises ($M = 1.94\%$) yielded significantly more accurate responses compared to plateaus and falls (0.8% and 0.18%, respectively, Fisher PLSD by participants and items), as well as a higher rate of monosyllabic responses ($M = 8.7\%$) compared to onsets with sonority plateaus and falls (4.8% and 2.2%, respectively, Fisher PLSD by participants and items). Onsets with sonority plateaus and falls did not differ significantly on the rate of correct, disyllabic or monosyllabic responses (in all cases, Fisher PLSD did not reach significance by both participants and items). The lack of difference between sonority plateaus and falls might well reflect the insensitivity of the transcription task, as such distinctions are observed under implicit testing in Experiments 1–6. The demand to explicitly state the representation of an unattested onset and the memory and attention load associated with its transcription might render the transcription of an unattested onset a rather indirect reflection of its implicit representation.

Table 5
The contribution of various predictors of response accuracy in Experiment 1 (a linear step-wise regression analysis using forced entries of predictors)

	Step	Predictor	R^2 change	F change	*df*	P value
Rises vs. plateaus	1	Statistical properties	.096	1.143	5, 54	.349
	2	Homorganicity	.519	71.28	1, 53	.000
	3	Markedness	.114	21.77	1, 52	.000
Plateaus vs. falls	1	Statistical properties	.084	1.26	4, 55	.29
	2	Homorganicity	.027	1.62	1, 54	.21
	3	Markedness	.133	9.35	1, 53	.003

the same place of articulation (i.e., homorganic consonants), and this constraint also appears to affect the perception of unattested clusters in laboratory tasks (Hallé et al., 1998). A subset of our materials violates this constraint (e.g., *ltap*). The dislike of marked onsets could thus reflect their infrequent statistical properties and homorganicity, rather than their marked sonority profile. To evaluate this explanation, we submitted participants' response accuracy to two linear regression analyses; one comparing onsets with sonority rises vs. plateaus, and another comparing plateaus vs. falls. We first forced into the model the statistical characteristics of our materials (given in Table 3); in the second step we forced in the homorganicity factor, whereas markedness was entered last. Markedness accounted for significant unique variance in the comparison of sonority rises and plateaus as well as in the comparison of plateaus and falls (see Table 5).[6] Thus, the misperception of marked clusters is inexplicable by either segment co-occurrence or homorganicity alone.

2.5. Experiment 2

The findings of Experiment 1 suggest that marked monosyllabic inputs are recoded as disyllabic. The markedness hypothesis attributes the finding to speakers' linguistic knowledge. Because extremely marked clusters are not tolerated in English, they are recoded as their unmarked, disyllabic counterparts. Alternatively, the misperception of marked clusters could reflect their phonetic properties: clusters such as *lbif* might be confusable with their disyllabic counterparts, *lebif*, because the spectral properties of sonorants and vowels are similar. The erroneous responses to marked onsets might thus reflect their imperceptibility (specifically, the indistinct quality of the contrast between C_1C_2VC and $C_1əC_2VC$) rather than the organization of the English grammar.

To begin evaluating this possibility, we first sought to determine whether marked onsets are perceptible for some speakers. To this end, we replicated Experiment 1 with speakers of Russian – a language that tolerates all cluster types used in our experiment. If the perception of marked monosyllabic inputs as disyllabic reflects

---

[6] Similar analyses performed on response time did not yield a unique effect of markedness in either analysis (both $F < 1$).

only differences in acoustic properties (e.g., differences between the acoustic properties of *bdif* and *lbif*), then Russian speakers should exhibit similar patterns of confusion. Conversely, if such perceptual errors are at least partly due to the preference for sonority rise in all grammars, then (a) unlike English speakers, Russian speakers should perceive most monosyllabic inputs accurately (since their grammar tolerates all cluster types examined in our experiments); and (b) like English speakers, Russian speakers' performance should be modulated by markedness: As markedness increases, accuracy should decrease.

2.6. Method

2.6.1. Participants

Sixteen native Russian speakers participated in the experiment. These participants were students at the University of Haifa, Israel. They emigrated to Israel between the ages 11–18 years, and they all reported speaking Russian at home. Participants were paid $5 for taking part in the experiment.

Materials and procedure were the same as in Experiment 1, except for the use of Russian words for practice (e.g., *drov*, 'log'; *darov*, '*present*').

2.7. Results

We first compared the performance of Russian and English speakers (in Experiments 2 and 1, respectively) by means of a 2 Language (English × Russian) × 2 Syllable (one vs. two) × 3 Cluster types (sonority rises, plateaus, and falls) ANOVA's. The significance of the three-way interaction in response accuracy (Fs(2, 60) = 13.89, MSE = .011, $p < .0001$), Fi(2, 58) = 22.80, MSE = .013, $p < .0001$; In response time: Fs(2, 46) < 1, MSE = 18,482, Fi(2, 54) = 1.58, MSE = 8727, $p < .22$, n.s. indicates that the discrimination of the various types of clusters from their disyllabic counterparts is modulated by linguistic experience. We thus turned to examine the responses of Russian speakers separately.

2.7.1. Response accuracy

Mean response accuracy of the Russian speakers is presented in Fig. 1. An inspection of the means suggests that Russian speakers perceived most monosyllabic inputs accurately. However, their accuracy was modulated by the markedness of the monosyllabic form. An ANOVA (2 syllables × 3 cluster types) yielded a significant interaction (Fs(2, 30) = 16.51, $p < .0001$); (Fi(2, 58) = 13.08, $p < .0001$). The simple effect of cluster type was significant for both monosyllabic inputs (Fs(2, 30) = 7.02, $p < .004$; Fi(2, 58) = 9.62, $p < .0003$) and their disyllabic counterparts (Fs(2, 30) = 14.97, $p < .0002$); (Fi(2, 58) = 10.19, $p < .0003$). Planned contrasts showed that, with monosyllabic inputs, participants responded more accurately to sonority plateaus than to falls (ts(30) = 3.54, $p < .002$); (ti(58) = 4.15, $p < .0002$). With disyllabic inputs, clusters with sonority rises elicited more errors than plateaus (ts(30) = 4.17, $p < .0003$; ti(58) = 3.44, $p < .002$). No other planned contrasts were significant.

Table 6
Mean correct response time (in ms) in Experiment 2 for monosyllabic and disyllabic inputs as a function of the number of syllables and the sonority profile of the monosyllabic counterpart

Cluster type	Monosyllabic		Disyllabic	
	M	*SD*	*M*	*SD*
Rise (e.g., *bnif*)	1293	133	1378	180
Plateau (e.g., *bdif*)	1308	128	1340	111
Fall (e.g., *lbif*)	1372	124	1311	101

2.7.2. Response time

Correct responses falling 2.5 SD above or below the mean (2.86% of the correct responses) were eliminated from the analysis of response latency. Mean correct response time is presented in Table 6. In accord with the accuracy data, as the markedness of the monosyllabic form increased, Russian speakers were slower to classify it as monosyllabic, and faster to respond to its disyllabic counterpart. The ANOVA (2 syllable × 3 type) on correct response time yielded a significant interaction (Fs(2, 30) = 5.99, $p < .007$; Fi(2, 56) = 5.09, $p < .01$). The simple effect of sonority type was marginally significant with monosyllabic inputs (Fs(2, 30) = 6.72, $p < .004$; Fi(2, 58) = 3.14, $p < .06$), but not with their disyllabic counterparts (Fs(2, 30) = 1.11, $p < .35$, n.s.; Fi(2, 58) = 2.22, $p < .12$). Planned contrasts showed that monosyllabic inputs with sonority plateaus elicited faster responses than sonority falls (ts(30) = 2.63, $p < .02$; ti(58) = 1.92, $p < .06$). No other planned contrasts were significant.

The main finding of Experiment 2 is that, unlike English speakers, Russian speakers perceived monosyllabic inputs accurately (all $p < .0002$ compared to 50% chance). In fact, Russian speakers were more accurate in perceiving monosyllabic inputs than their disyllabic counterparts, a fact that might be due to phonotactics constraints on CəCVC words in Russian. Like English speakers, however, the performance of Russian speakers was modulated by the markedness of the input. Although the convergent effects of markedness with speakers of English and Russian are consistent with the hypothesis that both grammars favor unmarked to marked structures, we currently cannot rule out the possibility that the performance of Russian speakers might reflect the statistical properties of their lexicon. Nonetheless, the ability of Russian speakers to accurately perceive onset clusters suggests that these items are not invariably imperceptible. Accordingly, the misperceptions of English speakers are at least partly due to their linguistic experience, rather than to stimuli properties alone.

3. The effect of markedness preferences on identity judgment

3.1. Experiment 3

The findings of Experiments 1–2 suggest that the linguistic knowledge of English speakers triggers perceptual illusions: English speakers misperceive marked clusters (e.g., *lbif*) as their disyllabic counterparts (e.g., *lebif*). Experiments 3–4 directly test this hypothesis. Participants are presented with two auditory stimuli, sampled from the nonwords used in previous experiments. The stimuli are either identical (e.g., *lbif–lbif*) or epenthetically related (e.g., *lbif–lebif*). Participants are simply asked to determine whether the two stimuli are identical. If English speakers recode marked clusters as their epenthetic counterparts (e.g., *lbif* → *lebif*), then as markedness increases, they should experience difficulty (i.e., an increase in errors and correct response time) in determining that epenthetically related items are distinct.

3.2. Method

3.2.1. Participants

Thirty native English speakers, students at Florida Atlantic University took part in the experiment in partial fulfillment of course requirements.

3.2.2. Method

The materials consisted of the 90 pairs of monosyllabic and disyllabic nonwords from Experiments 1–2. These materials were arranged in pairs. In half of the trials, the pair members were identical (either monosyllabic, *lbif–lbif*, or disyllabic, *lebif–lebif*), whereas in the other half, they were epenthetically related (e.g., *lbif–lebif*, *lebif–lbif*). The materials were next arranged in two lists, matched for the number of stimuli per condition (target type × identity × order) and counterbalanced, such that, within a list, each item appeared in either the identity or the nonidentity condition. As in previous experiments, these items were mixed with fillers including onsets that are attested in English.

3.2.3. Procedure

Each trial began with a fixation point (*) and a message indicating the trial number. Participants initiated the trial by pressing the space bar, triggering a change in the color of the computer screen and the presentation of the first auditory stimulus, followed by the second one (with an onset asynchrony of 1500 ms). Participants indicated whether the stimuli were identical by pressing the 1 or 2 keys, for "identical" and "non-identical" responses, respectively. Slow responses ($rt > 2500$ ms) received a computerized warning signal. The experiment was preceded by a short practice using English words (e.g., *plight–plight* vs. *polite–plight*).

3.3. Results

Trials with identical pairs elicited accurate ($M = 96\%$) and fast ($M = 940$ ms) responses. Our interest, however, is in responses to trials with nonidentical items. Correct responses falling 2.5 *SD* above or below the mean (2.7% of the correct responses) were eliminated from the analysis of response latency. Mean response accuracy and reaction time for nonidentical trials (e.g., *lebif–lbif*) is provided in Figs. 2 and 3, respectively.

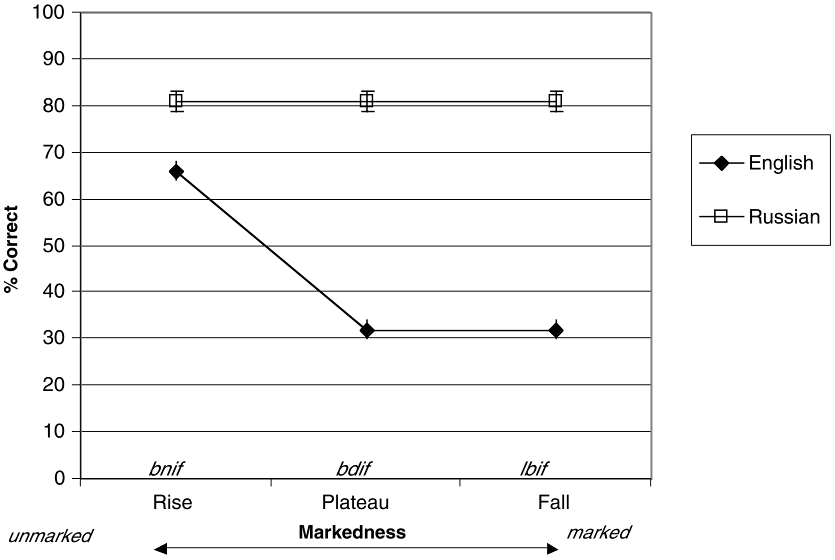

Fig. 2. Mean response accuracy of English and Russian speakers to non-identical trials in Experiments 3–4 as a function of the markedness of the monosyllabic input. Error bars represent the confidence interval constructed for the difference among the means.

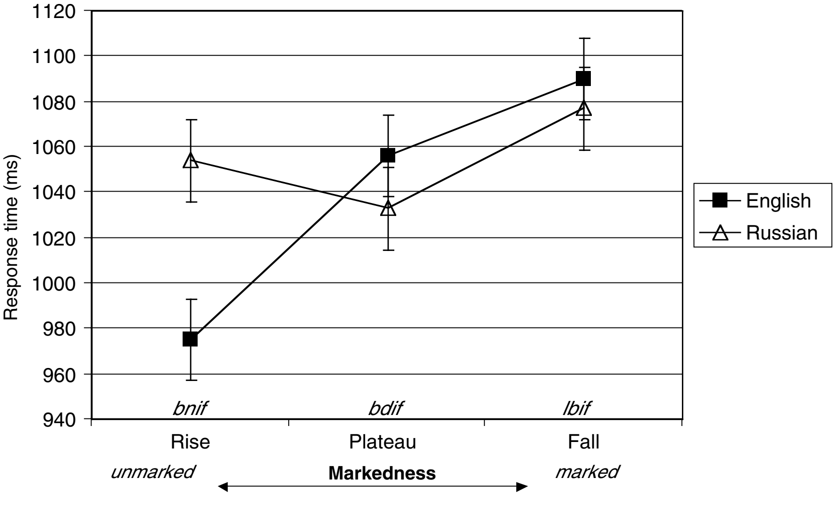

Fig. 3. Mean correct response time of English and Russian speakers to non-identical trials in Experiments 3–4 as a function of the markedness of the monosyllabic input. Error bars represent the confidence interval constructed for the difference among the means.

A one-way ANOVA yielded a significant effect of cluster type in response accuracy (Fs(2, 58) = 178.82, $p < .0002$; Fi(2, 58) = 27.66, $p < .0002$) and response time (Fs(2, 58) = 22.54, $p < .0002$; Fi(2, 56) = 15.65, $p < .0002$). Planned contrasts

Table 7
The contribution of various predictors of response accuracy in Experiment 3 (a linear step-wise regression analysis using forced entries of predictors)

	Step	Predictor	R^2 change	F change	*df*	P value
Rises vs. plateaus	1	Statistical properties	.08	1.25	4, 55	.30
	2	Homorganicity	.39	39.93	1, 54	.001
	3	Markedness	.22	38.70	1, 53	.001
Plateaus vs. falls	1	Statistical properties	.059	<1	4, 55	n.s.
	2	Homorganicity	.034	2.06	1, 54	.15
	3	Markedness	.003	<1	1, 53	n.s.

showed onsets of rising sonority elicited significantly faster (ts(58) = 4.61, $p < .0001$; ti(56) = 3.36, $p < .002$) and more accurate (ts(58) = 16.46, $p < .0001$; Fi(58) = 6.48, $p < .0001$) responses compared to onsets with sonority plateaus. Responses to onsets with sonority plateaus were faster compared to falls (ts(58) = 1.92, $p < .06$; ti(56) = 2.19, $p < .04$; albeit no more accurate, both $t < 1$).

3.4. Discussion

The findings of Experiment 3 demonstrate that, as the markedness of the onset increases, people judge monosyllabic nonwords as identical to their disyllabic counterparts. On most trials, people misjudge non-identical pair-members with sonority plateaus ($M = 68\%$, ts(29) = 6.74, $p < .0002$; ti(29) = 5.63, $p < .0002$) and falls ($M = 65\%$, ts(29) = 6.04, $p < .0002$; ti(29) = 4.62, $p < .0002$) as identical.

As in Experiment 1, we examined whether the effect of markedness is due to the statistical properties of the English lexicon and the homorganicity of the onset consonants. To this end, we submitted participants' response accuracy data to a linear regression analysis in which we first forced in the statistical properties of the items, followed by homorganicity. Markedness was entered in the third and last step. We conducted two such analyses, one comparing sonority rises and plateaus, and one contrasting plateaus vs. falls. The comparison of plateaus and falls (see Table 7) did not yield a reliable unique effect of markedness. In contrast, the comparison of sonority rises and plateaus yielded a unique significant effect of markedness after controlling for the statistical properties of the items and homorganicity.[7]

The misperception of marked onsets as identical to their epenthetic counterparts agrees with the results form Experiment 1, in which the same items were misjudged as disyllabic. To determine whether this effect is at least partly due to linguistic experience (rather than across the board imperceptibility), we next compared the results of English speakers with Russian speakers.

---

[7] Similar outcomes were obtained when the same analyses were conducted on response time. Markedness accounted for significant unique variance in comparing onsets with sonority rises vs. plateaus ($R^2 = .099$, $F(1, 52) = 6.59$, $p < .02$), but not in the comparison of sonority plateaus vs. falls ($R^2 = .016$, $F(1, 52) < 1$).

3.5. Experiment 4

3.5.1. Method

Participants. Thirty Russian speakers, students at the University of Haifa, Israel, took part in the experiment. The characteristics of this population are described in Experiment 2. The materials and procedure are as in Experiment 3, except the use of Russian words for practice (e.g., *drov–drov, darov–drov*).

3.5.2. Results

Mean response accuracy and response time to identical trials were 96% and 1011 ms. Of interest are the responses to trials with nonidentical items (e.g., *lbif–lebif*), presented in Figs. 2 and 3. Correct responses falling 2.5 *SD* above or below the mean (2.62% of the correct responses) were eliminated from the analysis of response latency.

We first compared the responses of Russian and English speakers (in Experiments 4 and 3, respectively) to non-identical items by means of a 2 Language (English × Russian) × 3 Cluster type (rises, plateaus, falls) ANOVA's. The significance of the two-way interaction (In response accuracy: Fs(2, 116) = 99.46, MSE = .006, $p < .001$; Fi(2, 58) = 37.83, MSE = .016, $p < .0001$), In response time: (Fs(2, 116) = 10.24, MSE = 4695, $p < .0001$; Fi(2, 56) = 6.67, MSE = 4732, $p < .0025$) indicates that the effect of cluster type is modulated by linguistic experience. We thus turned to examine the effect of cluster type on the responses of Russian speakers separately.

The effect of cluster type approached significance in the analysis of response time (Fs(2, 58) = 3.01, $p < .06$; Fi(2, 58) = 3.24, $p < .05$), but not in response accuracy (both $F < 1$). Planned contrasts showed that participants were slower to respond to highly marked clusters of falling sonority compared to lesser marked clusters with sonority plateaus (ts(58) = 2.45, $p < .02$; ti(58) = 2.48, $p < .02$). However, performance with the most marked onsets of falling sonority did not differ from the least marked onsets with sonority rise (ts(58) = 1.2, $p < .24$; ti(58) = 1.72, $p < .10$). As indicated previously (see Experiment 2), our analysis cannot distinguish the effect of markedness from the statistical properties of the Russian lexicon. Whether Russian speakers are sensitive to markedness remains to be seen. Nonetheless, the performance of the Russian speakers offers a valuable control in interpreting the results from English speakers. Unlike English speakers, Russian participants were clearly able to distinguish monosyllabic inputs from their disyllabic counterparts ($M = 81\%$, all $p < .0002$, testing against 50% chance). These findings demonstrate that the tendency of English speakers to consider marked onsets as identical to their disyllabic counterpart must be partly shaped by their linguistic experience.

4. Why *lbif* is perceived as *lebif*: phonetic confusion vs. phonological repair?

Experiments 1–4 have established that English speakers represent marked clusters as identical to their epenthetic counterparts. These results are consistent with the proposal that marked onsets are actively recoded by the grammar. However, an alternative phonetic explanation attributes the findings to an inability to represent the acoustic properties of these auditory inputs (hereafter, *phonetic confusion*). In this view, English speakers misperceive marked clusters because the acoustic properties of these stimuli are more similar to their disyllabic counterparts. For example, the misperception of *lbif* as *lebif* might be due to the similar spectral properties of vowels and sonorants. Because the acoustic boundaries between sonorants and vowels are harder to identify, the initial sonorant in *lbif* might be more confusable with a sonorant-vowel sequence (e.g., *lebif*) compared to the initial obstruent in *bdif* and its epenthetic counterpart (e.g., *bedif*). The finding that Russian speakers can perceive such onsets accurately (in Experiments 2 and 4) demonstrates that the difficulties in encoding marked onsets can be partly overcome by linguistic experience, but it does not rule out the possibility that those difficulties are phonetic, rather than grammatical in nature.

Although markedness preferences may have their roots in the avoidance of perceptual confusion (cf, contributions in Hayes, Kirchner, & Steriade, 2004; Hume & Johnson, 2001), we do not think that this fact explains our findings. A careful inspection of the results challenges the attribution of our findings to phonetic confusion alone, specifically, the confusion between CC and CəC.[8] If the perceptual errors with marked *lb* onsets were solely due to the difficulty of discriminating the acoustic properties of *lbif* from *lebif*, then, to the extent they had any effect on the disyllabic counterpart, this effect should have been deleterious (i.e., since the input *lebif* is indistinguishable from *lbif*). But our findings show that disyllabic counterparts of marked onsets are perceived *more* accurately than counterparts of unmarked onsets. For instance, people were more accurate to classify *bedif* (the counterpart of *bdif*) compared to *benif* (the counterpart of *bnif*). This effect is inexplicable by the phonetic properties of disyllabic forms, such as the length of the pre-tonic vowel (which contrasts them with their monosyllabic counterparts). Disyllabic forms like *bedif* did not have a longer pre-tonic vowel than disyllabic forms like *benif* ($M = 0.077$, 0.068, and 0.086 s, for the counterparts of sonority rises, plateaus, and falls, respectively). In contrast, this effect has a simple grammatical explanation: Because marked monosyllabic forms are dispreferred, people avoid interpreting the acoustic input as a marked monosyllable (e.g., interpreting *lebif* as an intended *lbif*), and opt, instead, for a disyllabic interpretation. To further test this explanation, we submitted the response accuracy for disyllabic forms in Experiment 1 to a linear stepwise regression analysis, forcing the length of initial vowel in the disyllabic forms, and the statistical properties and homorganicity of the monosyllabic counterparts as the initial third predictors. The markedness of the monosyllabic counterpart, entered as the last predictor,

---

[8] The phonetic confusion account also makes incorrect predictions regarding the perception of medial clusters. If the misperception of *lbif* as *lebif* were solely due to the acoustic properties of these auditory inputs, then similar confusion should have occurred in word medial positions (cf. *e**lb**ow*). Although our experiments do not test this prediction, it is nonetheless significant that the literature on English contains no reports of confusion between falling sonority clusters like /lb/ and their epenthetic counterparts.

Table 8
The contribution of various predictors of the response accuracy of English speakers to disyllabic nonwords in Experiment 1 (a linear step-wise regression analysis using forced entries)

	Step	Predictor	R^2 change	*F* change	*df*	*P* value
Cluster type	1	Vowel length	.026	2.33	1, 88	.13
	2	Statistical properties	.458	18.60	4, 84	.001
	3	Homorganicity	.008	1.38	1, 83	.24
	4	Markedness	.033	5.61	1, 82	.02

accounted for unique variance even after controlling for vowel length, statistical properties and homorganicity (see Table 8).

In the following experiments, we further test whether the perceptual errors of English speakers with marked monosyllabic clusters are due to grammatical markedness or to phonetic confusions alone. To this end, we compare the perception of marked onsets under conditions that are either neutral (e.g., Experiment 5) or encourage participants to attend to phonetic detail (in Experiment 6). Our experiments specifically focus on the markedness of onsets with sonority plateaus and falls. The phonetic confusability account predicts that marked onsets should always be harder to perceive, hence, they should always be more confusable with their disyllabic counterparts relative to less marked onsets with sonority plateaus (e.g., *bdif*). Conversely, on the scenario that involves active phonological repair (illustrated in Fig. 4), grammatical repair selectively modifies the phonological form of marked onsets, a representation that is distinct from the output of the initial stage in speech perception (hereafter, we refer to that initial output as the *phonetic form*). This allows for the possibility that the surface form of marked onsets is not more confusable to their epenthetic counterparts compared to the less marked onsets with sonority plateaus. If so, the effect of phonological markedness might be dissociable from difficulties related to phonetic encoding. Previous research suggests that speakers store detailed phonetic episodes that encode indexical information (Dupoux & Green,

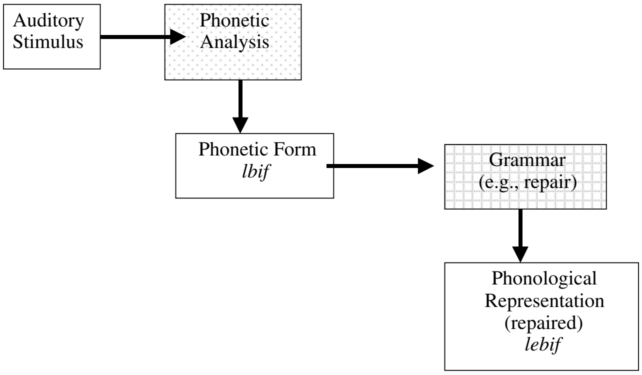

Fig. 4. The phonological repair hypothesis.

1997; Goldinger, 1998; McLennan & Luce, 2005) and allophonic variations (Gow, 2001; Maye, Werker, & Gerken, 2002; McLennan, Luce, & Charles-Luce, 2003; McMurray & Aslin, 2005), especially when such information is relevant to task demands (e.g., Goldinger, 1996). If speakers could access this (un-repaired) surface form under some special circumstances, then conditions that encourage the recovery of phonetic detail (e.g., in Experiment 6) should eliminate the greater confusability of marked onsets.

4.1. Experiment 5

Experiment 5 examines whether highly marked onsets of falling sonority are more likely to be repaired than the less marked onsets with sonority plateaus given conditions that do not specifically encourage accurate phonetic encoding. To this end, we compare the potential of marked and unmarked onsets to exert identity priming using a variant of the lexical decision procedure. Participants in our experiment are presented with two auditory stimuli, and they are asked to determine whether the stimuli are both existing English words (a double lexical-decision task, Meyer & Schvaneveldt, 1971). Although participants are instructed to respond to both stimuli in the same fashion, we expect the processing of the first stimulus to affect the second, hence, we refer to them as prime and target, respectively. In the critical trials, participants are presented with foils displaying sonority falls (e.g., *lbif*) or plateaus (e.g., *bdif*). Each foil is preceded either by an identity prime (e.g., *lbif–lbif*; *bdif–bdif*) or its epenthetic counterpart (e.g., *lebif–lbif*; *bedif–bdif*). Identity priming is assessed by comparing responses with the identity prime and the epenthetic control. Given the findings from Experiments 1–4, we expect the representation of novel onsets to be modulated by their markedness: Highly marked onsets with sonority falls (e.g., *lbif*) should be more likely to be perceived as identical with their epenthetic counterparts (e.g., *lebif*) than the less marked onsets with sonority plateaus (e.g., *bdif*). Accordingly, the magnitude of identity priming (reflecting the disadvantage of epenthetic primes relative to identity primes) should be smaller for the highly marked onsets with sonority falls relative to those with sonority plateaus.

Table 9
The structure of Experiments 5–6

	Prime	Target	Example	Experiment 5		Experiment 6	
"No" trials	Identity	Plateau	*bdif–bdif*	15	30	15	30
		Fall	*lbif–lbif*	15		15	
	Epenthesis	Plateau	*bedif–bdif*	15	30	15	30
		Fall	*lebif–lbif*	15		15	
	Dissimilar	Fillers	*train-bdif* *parade-mcug*		60		60
"Yes" trials	Identity	Fillers	*plight–plight*		30		60
	Epenthesis		*polite–plight*		**30**		**60**
	Dissimilar		*sustain–from* *crest-brow*		60		0

To interpret the source of these effects, it is useful to consider whether speakers' sensitivity to the relationships between targets and primes might reflect task-specific strategies. To this end, we provide the structure of our experimental materials in Table 9. Note that, in addition to the critical trials with sonority plateaus and falls (both items are nonwords, hence, they require a "no" response), the experiment also includes trials comprised of two words (e.g., *plight-polite*, requiring a "yes" response). Some of the "yes" and "no" trials include identical items, some are epenthetically related, whereas others are dissimilar. However, the distribution of epenthetic, identical and dissimilar trials is the same for "yes" and "no" trials. Because the detection of an epenthetic relationship between the target and prime is not beneficial for performing the experimental task, these conditions do not encourage attention to acoustic properties differentiating CCVC from CəCVC. Experiment 6 will next examine whether the potential of marked onsets to yield identity priming can be enhanced when attention to the phonetic distinction between CCVC and CəCVC is encouraged.

4.2. Method

4.2.1. Participants

Thirty-four native English speakers, students at Florida Atlantic University took part in the experiment in partial fulfillment of a course requirement.

4.2.2. Materials

The experiment included 240 trials (see Table 9). Each trial presented a pair of auditory stimuli. Half of the trials consisted of two English words (e.g., "yes" trials) whereas in the other half, at least one of the pair members was not a real English word ("no" trials). The second pair member (hereafter, the target) had a complex onset, whereas the first member either had a complex or a simple onset. Of interest are the trials in which the target contained an unattested onset. These targets included 30 pairs of CCVC targets whose complex onsets manifested either sonority plateaus or falls (see Appendix B). The statistical properties of these items (the position-specific probability of the four segments, bi-phone probability, neighborhood count and neighborhood frequency, calculated using the same databases described in Experiment 1) are presented in Table 10. An ANOVA suggested that items with sonority plateaus had significantly higher phoneme probability (Fi(1, 29) = 5.07, $p < .04$) and tended to have more neighbors than those with sonority falls

Table 10
The statistical properties of the nonwords with unattested onsets used in Experiments 5–6

Cluster type	Phoneme probability		Biphone probability		Number of neighbors		Neighbors' frequency	
	M	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*
Sonority plateau (e.g., *bdif)*	.124	.0298	.0018	.001	1.033	.999	21.80	34.80
Sonority fall (e.g., *lbif*)	.112	.0144	.0018	.001	.733	.785	20.0	42.79

(Fi(1, 29) = 3.85, $p < .06$), but the two types of items did not differ significantly on their bi-phone probability or the frequency of their neighbors (both $F < 1$).

Each participant was presented with 60 trials with unattested targets. In half of these trials, the target had an onset containing a sonority plateau (e.g., *bdif*), whereas in the other half it had a sonority fall (e.g., *lbif*). Each such target was preceded by either an identical nonword (e.g., *bdif–bdif*; *lbif–lbif*) or its epenthetic counterpart (e.g., *bedif–bdif*; *lebif–lbif*). To avoid repetition of the target, these materials were arranged in two sub-lists using a Latin-square design that balanced the prime × target combinations across participants and items (a total of 60 such trials per list).

To assure that participants attend to both stimuli, the remaining 60 "no" trials consisted of word–nonword combinations in a counterbalanced order. The structure of the nonword–word fillers was similar to that of the experimental trials: The nonwords had either sonority falls or sonority plateaus. In half of the filler trials, the prime had an onset cluster, whereas in the other half it was disyllabic.

The structure of "yes" trials was similar to the "no" trials. In half of the trials, the prime was monosyllabic, either identical to the target (e.g., *brick–brick*) or dissimilar (e.g., *speck–sport*), whereas in other half, it was disyllabic, either similar to the target (e.g., *support–sport*) or dissimilar (e.g., *sustain–from*).

4.2.3. The preparation of the materials

Because this experiment compared nonwords with English words, it was necessary to record both types of materials by an English speaker. Words and nonwords were recorded in separate lists by a female voice. Words were recorded as natural speech. Nonwords were generated by splicing out the vowels from the recordings of their disyllabic counterparts (e.g., excising [bdIf] from [bədIf]) using visual and auditory scrutiny of their waveforms. The initial syllable of the disyllabic source had either a secondary stress or no stress, and the precise quality and duration of its nucleus (schwa or vowel) were determined by the speaker. An inspection of the disyllabic waveforms of syllables beginning with a sonorant (e.g., *lebif*) revealed a silence between the initial vowel and the following consonant. To prevent the perception of initial sonorant consonants as syllabic, it was sometimes necessary to maintain that silence or, at times, lengthen it. The mean length of clusters with sonority falls (see Table 11) was shorter than those with sonority plateau ($F(1, 29) = 9.07$, $p < .006$).

The spliced materials were inspected by a linguist (D. Steriade) to assure that the splicing procedure effectively eliminated the traces of the initial vowel in the disyllabic source. As further assurance, we obtained rating of these materials by a group of

Table 11
The measurements (in ms) of the nonwords with unattested clusters used in Experiments 5–6

Cluster type	Consonant 1		Silence		Consonant 2		Vowel		Consonant 3		Total length	
	Mean	*SD*	Mean	*SD*	Mean	*SD*	Mean	*SD*	Mean	*SD*	Mean	*SD*
Plateau	118	28	0	0	101	34	262	91	330	77	812	99
Fall	135	29	16	14	88	27	246	61	260	85	745	111

seven speakers of Russian. The mean response accuracy for the disyllabic source words was 92%. Response accuracy for the monosyllabic words with sonority plateaus and falls (89%) did not differ significantly (both $F < 1$).

4.2.4. Procedure

Participants were seated in front of a computer screen wearing headphones. Each trial began with a fixation point in the center of the computer screen and participants initiated the start of the trial by pressing the space-bar key. Upon pressing the space-bar, two spoken stimuli were presented consecutively through the headphones. The onset-asynchrony between the two auditory stimuli was 1500 ms. Participants were instructed to determine if both of the spoken stimuli were real English words, and to indicate their response as fast and as accurately as possible by pressing the corresponding number on the numeric key pad (1 –"yes", indicating that both stimuli are words, 2 –"no", indicating that at least one stimulus is a nonword). Slow (slower than 3000 ms) and inaccurate responses triggered negative feedback from the computer (all response times are measured from the onset of the second word). The trials were presented in random order. Prior to the experimental phase, participants completed a practice session consisting of eight trials. Participants were tested in groups of up to three people.

4.3. Results and discussion

Correct responses falling 2.5 *SD* above or below the mean (2.5% of the correct responses) were eliminated from the analysis of response latency. Mean response latency for targets with sonority plateaus vs. falls preceded by identity prime and epenthetic controls are presented in Fig. 5. An inspection of the means suggests that

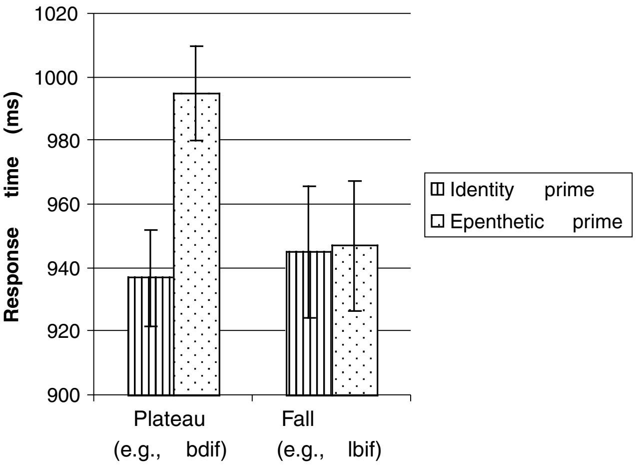

Fig. 5. Mean correct response time as a function of prime and target type in Experiment 5. Error bars represent the confidence interval constructed for the difference among the means.

identity priming is found for targets with sonority plateaus, but not for targets with sonority falls. This pattern is confirmed by the significance of the 2 (target type) × 2 (prime type) interaction in the ANOVA's by participants and items (Fs(1, 33) = 6.75, $p < .02$; Fi(1, 29) = 8.32, $p < .008$). The simple main effect of prime type was significant for onsets with sonority plateaus (Fs(1, 33) = 15.17, $p < .001$; Fi(1, 29) = 11.36, $p < .003$). In contrast, for targets with sonority falls, responses with the identity prime did not differ relative to controls (both F's < 1). Mean response accuracy was 97%, and it was unaffected by either prime or target type.

The finding that targets of falling sonority benefit from an epenthetic prime (e.g., *lebif–lbif*) as much as they do from an identity prime (e.g., *lbif–lbif*) suggests that the representation of such onset clusters is isomorphic to their epenthetic controls. Experiment 6 examines the source of this perceptual illusion.

4.4. Experiment 6

The results of Experiments 1–5 suggest that onsets of falling sonority, typologically the most marked class of onsets, are more likely to be perceived as identical to their epenthetic counterparts relative to less marked onsets with sonority plateaus. The findings are compatible with either a phonetic explanation (i.e., the inability to perceive the phonetic form of *lbif*) or a phonological explanation (i.e., the active recoding of the intact phonetic form *lbif* as *lebif*). On either account, identity primes are represented in a manner that is indistinguishable from epenthetic controls, hence, they do not differ on their priming potential.

To adjudicate between these explanations, it is necessary to determine whether the effects of markedness can be dissociated from the difficulties to represent the acoustic properties of such onsets. If the perceived isomorphism of marked onsets with their disyllabic counterparts is due to phonetic confusions, then such confusions should be always maintained, regardless of whether participants attend to phonetic cues in the input. In contrast, if the difficulties with marked onsets are due to an active grammatical repair, a process that modifies their (possibly intact) surface phonetic forms, then conditions that encourage attention to phonetic detail might eliminate the confusion of marked onsets with their disyllabic counterparts.

To encourage phonetic discrimination, Experiment 6 introduces a contingency between the response and the presence of epenthesis. Recall that Experiment 5 included a group of trials in which the prime and target were related epenthetically. However, the epenthetic relationship among primes and targets was equally likely in "yes" (e.g., *polite–plight*) and "no" trials (e.g., *lebif–lbif*, see Table 9). Because the presence of epenthesis was not predictive of the response, this manipulation did not encourage participants to attend to epenthesis. To promote attention to epenthesis, Experiment 6 creates a contingency between the presence of epenthesis and response. In this experiment, an epenthetic relationship between target and prime is twice as likely in trials consisting of two words (50% of the "yes" response) compared to trials consisting of at least one nonword (e.g., 25% of "no" response). Because epenthesis is a good predictor of response, this condition tacitly encourages participants to attend to the presence of acoustic cues for a vowel in the auditory stimulus. To the extent that such phonetic cues are securely encoded, they might allow participants to distinguish stimuli that carry acoustic cues for an extra vowel (e.g., *lebif*) from those in which epenthesis is generated by the grammar (e.g., *lbif*).

The two hypotheses of phonetic confusability and grammatical repair make distinct predictions with respect to the effect of this manipulation on identity priming. If the failure of targets of falling sonority onsets to produce identity priming (in Experiment 5) is due to their phonetic confusion with their epenthetic counterparts, then the same finding should emerge in the present experiment: Although attention to phonetic detail might improve the accurate perception of these onsets overall, it cannot eliminate their greater susceptibility to confusion with their disyllabic counterparts (e.g., of *lbif* with *lebif*) compared to onsets with sonority plateaus. Conversely, if the lack of priming was due to active phonological repair, then the heightened attention to phonetic detail might allow speakers to discriminate between marked clusters and their disyllabic counterparts, thereby enhancing the capacity of onsets with sonority falls to produce identity priming.

4.5. Method

4.5.1. Participants

Thirty-four native English speakers, students at Florida Atlantic University took part in this experiment in partial fulfillment of a course requirement.

4.5.2. Materials

The materials were identical to Experiment 5 with the exception that the group of 60 dissimilar word-word trials was eliminated (see Table 9). These trials were replaced by increasing the number of epenthetic word-word trials (e.g., *polite–plight*) from 30 to 60, and increasing the number of identical word-word trials from 30 to 60. Consequently, the proportion of epenthetic trials among the "yes" responses was 50%, which is twice as frequent as the proportion of epenthetic trials among "no" responses (25%).

4.5.3. Procedure

The procedure is the same as in Experiment 5.

4.6. Results

Correct responses falling 2.5 *SD* above or below the mean (2.4% of the correct responses) were eliminated from the analysis of response latency. Mean response latency as a function of prime type (identity vs. epenthetic control) and target (targets with sonority plateau vs. sonority fall) is depicted in Fig. 6.

We first examined whether our attempt to induce phonetic discrimination modulated the pattern of priming across Experiments 5–6 by means of a 2 Strategy (phonetic discrimination is either encouraged or not, Experiments 6 vs. 5) × 2 Target (sonority plateaus vs. falls) × 2 Prime (identity or epenthetic). The three-way interaction was

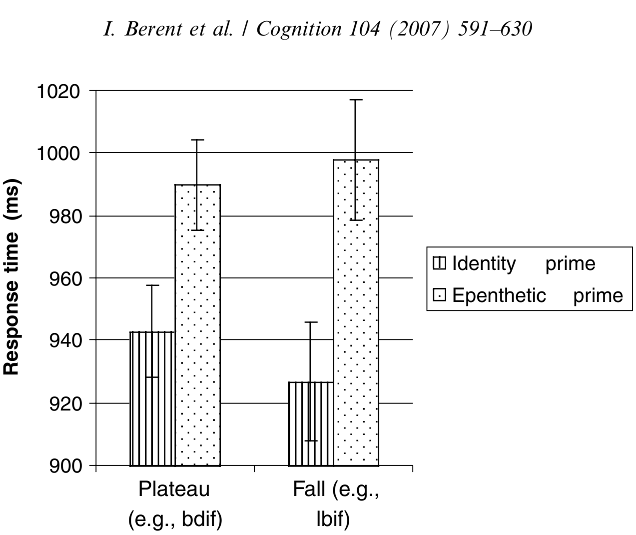

Fig. 6. Mean correct response time as a function of prime and target type in Experiment 6. Error bars represent the confidence interval constructed for the difference among the means.

significant in the analysis of response time (Fs(1, 66) = 7.23, MSE = 3716, $p < .01$, Fi(1, 29) = 5.41, MSE = 3901, $p < .03$; for accuracy: Fs(1, 66) = 1.0, MSE = .001, Fi(1, 29) = 1.45, MSE = .001, $p < .24$, n.s.), suggesting that the strategy manipulation indeed modulated the effect of priming for targets with sonority plateaus and falls.

We next turned to examine the pattern of priming when attention to phonetic discrimination is specifically encouraged (in Experiment 6). An inspection of the means suggests that participants responded more quickly in the presence of an identity prime compared to the epenthetic control for both types of targets. This conclusion is confirmed by the ANOVA (2 target × 2 prime) on response latency. There was a significant main effect of prime type (Fs(1, 33) = 19.38, $p < .0002$; Fi(1, 29) = 12.75, $p < .002$), and no evidence of an interaction (Fs(1, 33) = 1.37, $p < .26$, n.s.; Fi(1, 29) = 1.03, $p < .32$, n.s.). The simple effect of prime type was significant for targets with sonority falls (Fs(1, 33) = 13.93, $p < .002$; Fi(1, 29) = 14.58, $p < .002$) as well as for those with sonority plateaus (Fs(1, 33) = 10.53, $p < .004$; Fi(1, 29) = 4.47, $p < .05$). Mean response accuracy was 97%, and it was unaffected by either prime or target type.

4.7. Discussion

The results of Experiment 5 demonstrated that nonwords with sonority falls do not produce identity priming relative to epenthetic controls. This finding suggests that people do not discriminate the representation of onset clusters with sonority falls (e.g., *lbif*) from their epenthetic counterparts (e.g., *lebif*). However, these results could not determine the source of this effect: Is it due to the active phonological repair of such targets, or does it reflect their phonetic confusability with sequences that lack clusters? To adjudicate between these explanations, Experiment 6 promoted attention to the number of syllables in the stimulus. We reasoned that if the previous failure of onsets with sonority falls to produce identity priming were due to phonetic confusion, then attention to the acoustic properties that distinguish CCVC from CəCVC might improve their overall encoding, but not their inferiority relative to onsets with sonority plateaus. Conversely, if marked onsets are subject to grammatical repair, then the heightened attention to such acoustic properties might improve their discrimination from their epenthetic counterparts, resulting in a benefit to the identity prime relative to its epenthetic control.

The findings of Experiment 6 are consistent with the phonological repair hypothesis: Not only did onsets with sonority falls benefit from identity priming relative to controls, but the magnitude of the priming ($\Delta = 71$ ms) was numerically (albeit not significantly) larger than the priming of onsets with sonority plateaus ($\Delta = 47$ ms). These results suggest that when attention to the number of syllables in the stimulus is encouraged, people can encode highly marked onsets just as accurately as they encode onsets of lesser markedness. Note that, according to the phonological repair account, the null priming effect for targets of falling sonority in Experiment 5 is due to the capacity of the epenthetic counterpart to prime the target fully, whereas the resurrection of the priming effect, in Experiment 6, actually reflects a *decrease* in its priming potential due to reliance on a phonetically precise representation for the target. A comparison of the two experiments is consistent with this interpretation (see Fig. 7). When attention to the number of syllables in the stimulus was not emphasized (in Experiment 5), the effect of the epenthetic prime on targets of falling sonority was comparable to the identity prime, suggesting that the epenthetic counterpart primed targets of falling sonority as well as the identity prime. The enhanced attention to phonetic detail (in Experiment 6) slowed down responses by 51 ms to the

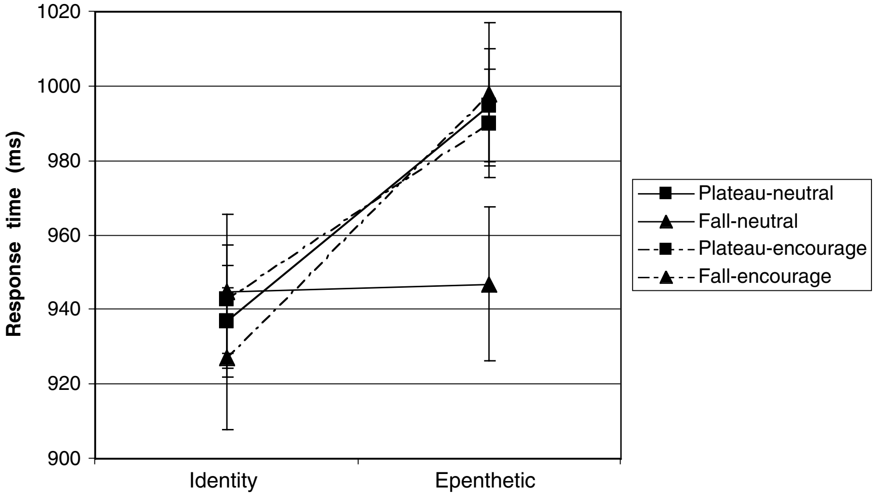

Fig. 7. Mean correct response time as a function of prime and target type under neutral conditions and conditions that encourage attention to phonetic detail (Experiments 5–6, respectively). Error bars represent the confidence interval constructed for the difference among the means.

lebif–lbif condition (Fs(1, 66) = 2.33, $p < .14$, n.s., Fi(1, 29) = 17.02, $p < .0004$; for comparison, targets with sonority plateaus exhibited a benefit of 5 ms), suggesting that the change in task demands affected participants' reliance on an (intact) phonetic representation of such onsets.

The nature of this task shift is not entirely clear from our results. One possibility is that the emphasis on accurate phonetic encoding (in Experiment 6) encouraged participants to abandon the phonological repair of the falling sonority onsets. This account, however, cannot explain why repair is used in the first place: If onsets with sonority falls are admissible by the grammar in their un-repaired form, then why are they repaired in Experiments 1–5? An alternative explanation is that the change in task demands only affected the reliance on surface phonetic cues. On this view, spoken language is encoded in (minimally) two representational formats (see Fig. 4): One corresponds to the surface phonetic form, whereas another captures the output of the phonological grammar. Although listeners normally access just the phonological form, special circumstances might encourage participants to access the unrepaired phonetic form. It is possible that some of the phonetic cues used by participants in our experiments are specifically due to the preparation of our materials by splicing. Nonetheless, the results of Experiment 6 make it clear that onsets with sonority falls are not invariably more confusable with their epenthetic counterparts than sonority plateaus. Accordingly, the repeated failure of participants to distinguish onsets like *lbif* from their epenthetic counterpart *lebif* (given the same materials and experimental task) implicates a process of phonological repair, triggered by their markedness.

5. General discussion

This research has investigated whether speakers are equipped with universal preferences regarding the structure of language. As a case study, we have examined restrictions on the sonority profile of onsets. Typological research suggests that onsets with large sonority rises (e.g., *blif*) are preferred to onsets with smaller rises (e.g., *bnif*), which, in turn are preferred to sonority plateaus (e.g., *bdif*); the plateaus, in turn, are preferred to onsets with sonority falls (e.g., *lbif*). The structure of the English lexicon offers English speakers little evidence for the sonority hierarchy: English manifests onsets with large sonority rises (e.g., *blif*), and it systematically lacks onsets with smaller rises or lacking rises (e.g., *bnif*, *bdif*, *lbif*). Of interest is whether English speakers are nonetheless sensitive to the markedness of onset clusters that are unattested in their lexicon.

Our experiments examined such sonority-related preferences by investigating the potential of unattested onset clusters to elicit perceptual illusions. Previous linguistic and psycholinguistic research suggests that English speakers repair illicit clusters by the epenthesis of a vowel between the onset consonants (e.g., *tla* → *tela*, Anderson, 1974; Davidson & Stone, 2004). Such repairs reflect the recoding of marked inputs to abide by the phonotactics of the language. Our experiments exploit such perceptual illusions to gauge the organization of the grammar. We reasoned that, if speakers are equipped with grammatical knowledge regarding the markedness of all sonority profiles, attested or unattested, then marked onsets should be more likely to elicit repair than less marked ones. Our experimental results are consistent with this prediction. As the markedness of the onset increases, speakers are more likely to consider the onset as disyllabic, and they tend to misperceive it as identical to its disyllabic counterpart. For this reason, highly marked onsets also benefit from priming by their epenthetic counterpart (e.g., *lebif–lbif*) as much as they benefit from identity priming (e.g., *lbif–lbif*).

It is unlikely that the perceptual distortions of marked onsets are solely due to their phonetic properties. First, unlike English speakers, speakers of Russian (a language tolerating all onset types examined in our experiments) do perceive the same items accurately on most trials, suggesting that these clusters are not invariably imperceptible. Second, the markedness of onset clusters affects the perception of their disyllabic counterparts: As the markedness of the monosyllabic form increases, people are more likely to perceive its disyllabic counterpart accurately. For example, people perceive *bedif* (the counterpart of the marked *bdif*) more accurately than *benif* (the counterpart of the less marked *bnif*). These findings have a simple grammatical explanation (i.e., the dispreference for marked monosyllabic onsets prevents the representation of their disyllabic counterparts as monosyllabic), but we are unaware of a phonetic explanation for this result. Finally, English speakers are capable of perceiving marked onsets accurately when attention to phonetic detail is emphasized. Specifically, under conditions that encourage attention to the presence of epenthesis, marked onsets with sonority falls benefit from identity priming to the same extent as less marked onsets with sonority plateaus, suggesting that participants are capable of distinguishing highly marked onset clusters from their epenthetic counterparts.[9]

The systematic misperception of marked sonority profiles (e.g., misperception of $lb > bd > bn$) suggest that marked onsets are dispreferred relative to less marked ones (e.g., preference of $bn > bd > lb$). These preferences are found despite the absence of such onsets in the English lexicon, and even when several statistical properties of the lexicon (the position-specific frequency of phonemes, bi-phone combinations, the number of neighbors and their frequency) are controlled by means of a regression analysis. These findings challenge the wide spread belief that phonological knowledge can be reduced to statistical knowledge of patterns instantiated in the lexicon. Dell et al. (2000) articulate this hypothesis quite clearly in discussing the effect of phonological knowledge on speech errors:

> "You may say "blug" instead of "bug" because you have said words beginning with [bl] in your recent past. But you would never say "lbug". Perhaps that is

---

[9] Our proposal that English speakers can recover an accurate phonetic representation of marked clusters does not imply that phonetic encoding is independent of linguistic experience. The well-established effect of linguistic experience on categorical perception (e.g., Kuhl, Williams, Lacerda, Stevens, & Lindblom, 1992; Werker & Tees, 1984) speaks against this possibility.

> because your phonology does not allow syllable initial [lb]. But perhaps it is instead because you do not have any recent experience saying this cluster. Some modern phonological theories tend to see the phonology of a language as being projected from its lexicon (see Broe and Pierrehumbert, 1999). That is, *there is no independent abstract representation of phonological patterns aside from what is stored in the lexicon*. Or, if there is some abstraction, it is computed from the contents of the lexicon. We would only add to this notion that the phonology is projected preferentially from those parts of the lexicon that are most accessible, such as recently experienced sound forms.'' (Dell et al., 2000, p. 1365, the italics are added).

Our findings obviously do not negate speakers' demonstrable sensitivity to statistical lexical information (e.g., Dell et al., 2000; Jusczyk, Luce, & Luce, 1994; Mattys, Jusczyk, Luce, & Morgan, 1999; Vitevitch, Armbruster, & Chu, 2004). Likewise, we cannot exclude all statistical accounts of our findings. For example, we do not evaluate the possibility that the preferences for novel onsets might reflect statistical knowledge of feature co-occurrence (a similar possibility is considered by Warker & Dell, 2006). In particular, participants might discriminate among unattested clusters by tracking the frequency of obstruents and sonorants at C1 and C2 positions. Unlike the hypothesis that speakers are equipped with knowledge about the hierarchy of sonority profiles (e.g., 3a–c), such a modified statistical explanation might fail to exhibit the preference for unmarked sonority profiles given onsets with unattested feature combinations (e.g., a preference for the sonority rise in *mlif* to the plateau in *mnif*–both manifesting an unattested combination of two sonorants) or unattested features (cf., Berent, Marcus, Shimron, & Gafos, 2002). Whether speakers manifest such preferences remains to be seen.

Our findings suggest that a full account of phonotactics must incorporate a productive grammatical component that encodes sonority-related knowledge. The present results nonetheless leave several unanswered questions regarding the nature of this knowledge and its origins. Our findings clearly implicate knowledge that is related to sonority. But whether this knowledge in fact concerns sonority per se, or other factors that correlate with sonority (e.g., amplitude modulation) remains unknown. The attempt to account for the preferences among different syllable structures in terms of sonority has been faced with numerous challenges, including the failure to clearly define the phonetic basis for sonority (Kawasaki-Fukumori, 1992; Ohala, 1990; but see Parker, 2002), the difficulties in devising a general scale of sonority (cf., Clements, 1990 vs. Steriade, 1982), and the presence of counterexamples to the sonority generalization (Levin, 1985; Selkirk, 1982; Steriade, 1982; Wright, 2004). Indeed, our present results do not necessarily demonstrate that the grammar directly constrains the markedness of sonority profiles. The preference for onsets such as *bdif* over *lbif* could reflect a narrow preference for obstruent–obstruent over liquid–obstruent onsets – a preference that does not extend to all profiles of sonority plateaus vs. falls.

Likewise, important questions remain regarding the source of sonority-related knowledge. Our findings demonstrate that English speakers manifest sonority-related preferences despite the lack of lexical evidence, either direct (i.e., the existence of the relevant onsets in the English lexicon) or indirect (the statistical co-occurrence of segments in English words). The lack of lexical evidence, however, does not amount to the absence of linguistic evidence altogether. English speakers might encounter consonant combinations such as *bn* due to schwa reduction in fast speech (e.g., *beneath* → ***bn****eath*), and the familiarity with such reduced forms could lead to the sonority-related preferences implicated by our results. Although the effect of sonority on schwa reduction has not been fully explored, the existing evidence does not support the attribution of our findings to the distribution of those clusters in reduced forms. First, speakers are unlikely to encounter clusters in reduced forms, since the rate of schwa reduction in conversational American English is low (Patterson, LoCasto, & Connine, 2003). Second, English listeners distinguish the outputs of schwa-reduction from underlying clusters (e.g., *support* vs. *sport*, LoCasto & Connine, 2002; Manuel et al., 1992). Thus, the encounter with a reduced form (e.g., the reduction of *beneath)* does not provide evidence for the existence of a CC cluster (e.g., *bn*). Finally, the distribution of reduced forms fails to account for the preference for sonority plateaus over falls. The distributional account specifically assumes that speakers are more likely to encounter rises or plateaus than falls in fast speech. But a careful phonetic analysis of the reduced obstruent–vowel–obstruent sequences in novel forms (e.g., ***pat****een* → ***pt****een*) indicates cues (e.g., the release phase between consonants and the aspiration of the second obstruent) that distinguish such forms from the underlying clusters (Walter, 2006). These findings suggest that sonority plateaus do not occur at the beginning of English words (including fast speech).

Although it is unlikely that participants encounter the onset clusters used in our experiments, their preferences for such onsets might be informed by indirect sources, including their experience with the phonetic properties of such consonant sequences in word-medial positions (e.g., *o****bn****oxious*, *a****bd****omen*, *A****lb****any*), as well as their experience with the onset clusters attested in the English language. Whether such phonetic experience is necessary for the induction of markedness remains to be seen. But the possibility that the sonority markedness hierarchy might be induced from phonetic experience is perfectly compatible with the existence of innate constraints on the organization of the grammar (Hayes & Steriade, 2004; Wright, 2004). Such constraints on language acquisition would further account for the repeated emergence of the sonority hierarchy across languages and its convergence with the preferences of individual speakers.

Acknowledgements

This research was supported by NIDCD grant DC003277. We are indebt to Paul Smolensky, who urged us to pursue this project and laid out its theoretical foundation. We also thank Mike Sagristano, for statistical advice, Gary Marcus, for comments on an earlier version of this manuscript, and Adam Albright, for discussions of this project. Finally, we are grateful to Samantha Murphy, for her continuous assistance in the recording and splicing of the auditory materials, and to Sasha Makarova, for recording the Russian materials.

Appendix A

The nonwords used in Experiments 1–4

Sonority rises	Sonority plateaus	Sonority falls
bwif	bdif	lbif
bnɑp	bdɑp	ɾgɑp
knim	kpim	lpim
knɛk	ktɛk	ɾtɛk
dlif	dbif	ɾdif
dlɑf	dgɑf	rdɑf
dmip	dgip	mdip
dmʊp	dgʊp	mdʊp
dnʊp	dbʊp	rdʊp
dniʃ	dgiʃ	rbiʃ
gmɛp	gdɛp	lgɛp
gmɑn	gbɑn	lfɑn
gmɛf	gbɛf	rgɛf
gmit	gbit	mgit
kmɛf	ktɛf	lkɛf
kmæf	kpæf	rgæf
knik	ktik	rkik
knʊk	kpɑk	mkʊk
kmʊp	ktɑp	ltɑp
kmɛp	ktɛp	rkɛp
pnik	pkik	ltik
pnæf	ptæf	rpæf
tlʊf	tkʊf	rtʊf
tlɛp	tkɛp	mtɛp
tnɑk	tkɑk	rtɑk
tmæf	tpæf	mtæf
tnɛf	tpif	rtɛf
tnʊk	tgʊk	mgʊk
tmæp	tpæp	rpæp
tmɑk	tpɑk	mtɑk

Appendix B

The nonwords used in Experiments 5–6

Sonority plateaus	Sonority falls
bdɪf	lbɪf
bdœp	rgœp
kpɪm	lpɪm
ktɛg	rtɛg
dbɪf	rdɪf
dgœf	rdœf
dgɪb	mdɪb
dgʌp	mdʌp
dbʌp	rdʌp
dgɪʃ	rbɪʃ
gdɛb	lgɛb
gbœn	lfœn
gbɛf	rgɛf
gbɪd	mgɪd
ktɛθ	lkɛθ
kpœf	rgœf
ktɪg	rkɪg
kpœg	mkʌg
ktœb	ltœb
ktɛp	rkɛp
pkɪg	ltɪg
ptœf	rpœf
tkʌf	rtʌf
tkɛp	mtɛp
tkœg	rtœg
tpœf	mtœf
tpɪf	rtɛf
tgʌk	mgʌk
tpœb	rpœb
tpœg	mtœg

References

Anderson, S. R. (1974). *The organization of phonology*. New York: Academic Press.

Berent, I., Marcus, G. F., Shimron, J., & Gafos, A. I. (2002). The scope of linguistic generalizations: Evidence from Hebrew word formation. *Cognition, 83*, 113–139.

Blevins, J. (1995). The syllable in phonological theory. In J. Goldsmith (Ed.), *The handbook of phonological theory* (pp. 206–244). Cambridge, MA: Blackwell.

Broselow, E., & Finer, D. (1991). Parameter setting in second language phonology and syntax. *Second Language Research, 7*, 35–59.

Chomsky, N., & Halle, M. (1968). *The sound pattern of English*. New York: Harper & Row.

Clements, G. N. (1990). The role of the sonority cycle in core syllabification. In M. Beckman (Ed.), *Papers in laboratory phonology I: Between the grammar and physics of speech* (pp. 282–333). Cambridge: Cambridge University Press.

Davidson, L. (2000). Experimentally uncovering hidden strata in English phonology. In L. Gleitman & A. Joshi (Eds.), *Proceedings of the 22nd annual conference of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates.

Davidson, L., & Stone, M. (2004). Epenthesis versus gestural mistiming in consonant cluster production. In *Proceedings of the West Coast Conference on Formal Linguistics* (*WCCFL*), San Diego.

Dell, G. S., Reed, K. D., Adams, D. R., & Meyer, A. S. (2000). Speech errors, phonotactic constraints, and implicit learning: A study of the role of experience in language production. *Journal of Experimental Psychology: Learning, Memory & Cognition, 26*, 1355–1367.

Dupoux, E., & Green, K. (1997). Perceptual adjustment to highly compressed speech: Effects of talker and rate changes. *Journal of Experimental Psychology: Human Perception and Performance, 23*, 914–927.

Dupoux, E., Kakehi, K., Hirose, Y., Pallier, C., & Mehler, J. (1999). Epenthetic vowels in Japanese: A perceptual illusion?. *Journal of Experimental Psychology: Human Perception and Performance 25*, 1568–1578.

Dupoux, E., Pallier, C., Kakehi, K., & Mehler, J. (2001). New evidence for prelexical phonological processing in word recognition. *Language and Cognitive Processes, 5*, 491–505.

Gierut, J. A. (1999). Syllable onsets: Clusters and adjuncts in acquisition. *Journal of Speech, Language and Hearing Research, 42*, 708–726.

Goldinger, S. D. (1996). Words and voices: Episodic traces in spoken word identification and recognition memory. *Journal of Experimental Psychology: Learning Memory and Cognition, 22*, 1166–1183.

Goldinger, S. D. (1998). Echoes of echoes? An episodic theory of lexical access. *Psychological Review, 105*, 251–279.

Gouskova, M. (2002). *Exceptions to sonority distance generalizations*. 38th meeting of the Chicago Linguisitc Society, Chicago.

Gow, D. W. J. (2001). Assimilation and anticipation in continuous spoken word recognition. *Jounal of Memory and Language, 45*, 133–159.

Greenberg, J. H. (1978). Some generalizations concerning initial and final consonant clusters. In E. A. Moravcsik (Ed.), *Universals of human language* (Vol. 2, pp. 243–279). Stanford, CA: Stanford University Press.

Halle, M. (1971). *The sound pattern of Russian*. The Hauge: Mouton.

Hallé, P., Segui, J., Frauenfelder, U. H., & Meunier, C. (1998). The processing of illegal consonant clusters: A case of perceptual assimilation?. *Journal of Experimental Psychology: Human Perception and Performance 24*, 592–608.

Hammond, M. (1999). *The phonology of English: A prosodic optimality-theoretic approach*. Oxford, New York: Oxford University Press.

Hayes, B., Kirchner, R. M., & Steriade, D. (Eds.). (2004). *Phonetically based phonology*. Cambridge: Cambridge University Press.

Hayes, B., & Steriade, D. (2004). A review of perceptual cues and cue robustness. In D. Steriade (Ed.), *Phonetically based phonology* (pp. 1–33). Cambridge: Cambridge University Press.

Hume, E. V., & Johnson, K. (2001). *The role of speech perception in phonology*. San Diego: Academic Press.

Jakobson, R. (1941). *Kindersprache, Aphasie und allgemeine Lautgesetze [Child language, aphasia and phonological universals]*. Frankfurt: Suhrkamp.

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

Kawasaki-Fukumori, H. (1992). An acoustical basis for universal phonotactic constraints. *Language and Speech, 35*, 73–86.

Kenstowicz, M. (1994). *Phonology in generative grammar*. Cambridge, MA: Blackwell.

Kiparsky, P. (1979). Metrical structure assignment is cyclic. *Linguistic Inquiry, 10*, 421–442.

Kucera, H., & Francis, W. (1967). *Computational analysis of present-day American English*. Providence, RI: Brown University Press.

Kuhl, P. K., Williams, K. A., Lacerda, F., Stevens, K. N., & Lindblom, B. (1992). Linguistic experience alters phonetic perception in infants by 6 months of age. *Science, 255*, 606–608.

Ladefoged, P. (1975). *A course in phonetics*. New York: Harcourt Brace Jovanovich.

Levin, J. (1985). *A Metrical Theory of Syllabicity. Ph.D. dissertation*. Unpublished doctoral dissertation, MIT.

LoCasto, P. C., & Connine, C. M. (2002). Rule-governed missing information in spoken word recognition: Schwa vowel deletion. *Perception & Psychophysics, 64*, 208–219.

Manuel, S. Y., Shattuck-Hufnagel, S., Huffman, M., Stevens, K. N., Carlson, R., & Hunnicut, S. (1992). Studies of vowel and consonant reduction. In *Proceedings of the 1992 international conference on spoken language processing*.

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

Mattys, S. L., Jusczyk, P. W., Luce, P. A., & Morgan, J. L. (1999). Phonotactic and prosodic effects on word segmentation in infants. *Cognitive Psychology, 38*, 465–494.

Maye, J., Werker, J. F., & Gerken, L. (2002). Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition, 82*, B101–B111.

McLennan, C. T., & Luce, P. A. (2005). Examining the time course of indexical specificity effects in spoken word recognition. *Journal of Experimental Psychology: Learning Memory and Cognition, 31*, 306–321.

McLennan, C. T., Luce, P. A., & Charles-Luce, J. (2003). Representation of lexical form. *Journal of Experimental Psychology: Learning Memory and Cognition, 29*, 539–553.

McMurray, B., & Aslin, R. N. (2005). Infants are sensitive to within-category variation in speech perception. *Cognition, 95*, B15–B26.

Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology, 90*, 227–234.

Moreton, E. (2002). Structural constraints in the perception of English stop-sonorant clusters. *Cognition, 84*, 55–71.

Ohala, D. K. (1999). The influence of sonority on children's cluster reductions. *Journal of communication disorders, 32*, 397–421.

Ohala, J. J. (1990). Alternatives to the sonority hierarchy for explaining segmental sequential constraints. *Papers from the regional meetings, Chicago Linguistic Society* (Vol. 2, pp. 319–338).

Parker, S. (2002). *Quantifying the Sonority Hierarchy*. Unpublished doctoral dissertation, University of Massachusetts, Amherst, MA.

Patterson, D., LoCasto, P. C., & Connine, C. M. (2003). Corpora analyses of frequency of schwa deletion in conversational American English. *Phonetica, 60*, 45–69.

Pertz, D. L., & Bever, T. G. (1975). Sensitivity to phonological universals in children and adolescents. *Language, 51*, 149–162.

Pitt, M. A. (1998). Phonological processes and the perception of phonotactically illegal consonant clusters. *Perception & Psychophysics, 60*, 941–951.

Prince, A., & Smolensky, P. (1993/2004). *Optimality theory: Constraint interaction in generative grammar*. Malden, MA: Blackwell Pub.

Romani, C., & Calabrese, A. (1998). Syllabic constraints on the phonological errors of an aphasic patient. *Brain and Language, 64*, 83–121.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science, 274*, 1926–1928.

Selkirk, E. O. (1982). The syllable. In N. Smith (Ed.), *The structure of phonological representations* (pp. 337–383). Dordrecht, Holland: Foris Publications.

Smolensky, P. (1996). On the comprehension/production dilemma in child language. *Linguistic Inquiry, 27*, 720–731.

Smolensky, P. (2006). Optimality in phonology II: Markedness, feature domains, and local constraint conjunction. In P. G. L. Smolensky (Ed.), *The harmonic mind: From neural computation to optimality-*

theoretic grammar Vol. 2. Linguistic and Philosophical Implications (pp. 27–160). Cambridge, MA: MIT Press.

Steriade, D. (1982). *Greek prosodies and the nature of syllabification; Ph.D. dissertation*. Unpublished doctoral dissertation, MIT, Cambridge (available from MITWPL, Department of Linguistics and Philosophy, MIT).

Steriade, D. (2003). The syllable. In W. Frawley (Ed.), *International encyclopedia of linguistics*. Oxford: Oxford University Press.

Treiman, R. (1984). On the status of final consonant clusters in English syllables. *Journal of Verbal Learning and Verbal Behavior, 23*, 343–356.

Treiman, R., Bowey, J., & Bourassa, D. (2002). Segmentation of spoken words into syllables by English-speaking children as compared to adults. *Journal of Experimental Child Psychology, 83*, 213–238.

Treiman, R., & Danis, C. (1988). Syllabification of intervocalic consonants. *Journal of Memory and Language, 27*, 87–104.

Treiman, R., Straub, K., & Lavery, P. (1994). Syllabification of bisyllabic nonwords: Evidence from short-term memory errors. *Language and Speech, 37*, 45–60.

Trubetzkoy, N. S. (1938/1958). *Grundzuge der Phonologie*. Gottingen: Vandenhoeck & Ruprecht.

Vitevitch, M. S., Armbruster, J., & Chu, S. (2004). Sublexical and lexical representations in speech production: Effects of phonotactic probability and onset density. *Journal of Experimental Psychology: Learning Memory and Cognition, 30*, 514–529.

Vitevitch, M. S., & Luce, P. A. (2004). A web-based interface to calculate phonotactic probability for words and nonwords in English. *Behavior Research Methods Instruments, & Computers, 36*, 481–487.

Warker, J. A., & Dell, G. S. (2006). Speech errors reflect newly learned phonotactic constraints. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32*, 387–398.

Walter, M. A. (2006). *Repetition avoidance in language*. Unpublished doctoral dissertation, Massachusetts Institute of Technology, Cambridge, MA.

Werker, J. F., & Tees, R. C. (1984). Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant Behavior and Development, 7*, 49–63.

Wilson, C. (2003). Experimental investigation of phonological naturalness. In M. Tsujimura (Ed.), *Proceedings of the 22nd west coast conference on formal linguistics* (pp. 533–546). Somerville, MA: Cascadilla.

Wilson, C. (in press). *Learning phonology with substantive bias: An experimental and computational study of velar palatalization. Cognitive science*.

Wright, R. (2004). A reivew of perceptual cues and robustness. In D. Steriade, R. Kirchner, & B. Hayes (Eds.), *Phonetically based phonology* (pp. 34–57). Cambridge: Cambridge University Press.

Zuraw, K. (2005). *The role of phonetic knowledge in phonological patterning: Corpus and survey evidence from Tagalog*. (submitted for publication).