INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

Bell & Howell Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
800-521-0600

UMI[®]

The role of neighbourhood density and probabilistic phonotactics in children's and adults'
spoken word recognition

by

David Ridgeway, B.A.

A thesis submitted to
the Faculty of Graduate Studies and Research
in partial fulfillment of
the requirements for the degree of

Master of Arts

Department of Psychology

Carleton University

Ottawa, Ontario

January 15, 2001

©2000, David Ridgeway

National Library
of Canada

Acquisitions and
Bibliographic Services

395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque nationale
du Canada

Acquisitions et
services bibliographiques

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

Our file Notre référence

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

0-612-57714-7

Canada

The undersigned recommend to the
Faculty of Graduate Studies and Research
acceptance of the thesis

"The role of neighbourhood density and probabilistic phonotactics in children's and
adults' spoken word recognition"

submitted by

David Ridgeway

in partial fulfillment of the requirements for
the degree of Master of Arts

Chair

Thesis Supervisor

Carleton University

January 15, 2001

Abstract

The effects of neighbourhood density (ND) and probabilistic phonotactics (PP) on spoken word recognition were examined in three experiments. Experiment 1 tested 2-, 3-, and 4-year old children using a picture pointing task. Performance was better for words with high than low probability phoneme sequences. Experiments 2 and 3 tested adults using the same stimuli in the picture pointing task and a naming task, respectively. Adult performance was better for high rather than low ND words, an effect opposite to that typically found with adults (Vitevitch & Luce, 1998, 1999). PP only affected adult performance in the naming task. The use of an orthogonal design in which ND and PP were varied independently is likely the main reason for the pattern of results obtained. The results suggest a revised role for ND and PP in models of spoken word recognition.

Acknowledgements

The author is deeply indebted to several people, without whom this research would not have been possible. Primarily, I must express my sincere gratitude to Dr. John Logan for his invaluable guidance and support. Throughout the evolution of this thesis, Dr. Logan contributed many hours of advice and editing; complemented by encouraging words which did not go unheeded. I also owe a great deal of gratitude to Diana DeStefano, who tirelessly assisted with data collection and preparation of stimulus materials, as well as providing unlimited encouragement and inspiration. Thanks also to Carol Unruh for her data collection efforts for Experiment 1. I thank Dr. Chris Herdman for the many informal, insightful conversations throughout my years at Carleton University. His friendship and support were extremely beneficial to me. Finally, I wish to acknowledge the support I received from the staff in the Psychology department, my friends, and fellow students at Carleton University.

Table of Contents

Abstract.....	iii
Acknowledgements	iv
Table of Contents	v
Table of Figures	vii
Table of Appendices	viii
Word recognition in children.....	10
Experiment 1	14
Method	14
Participants	14
Stimuli.....	15
Procedure	18
Results.....	19
Discussion	30
Experiment 2.....	32
Method	34
Participants	34
Stimuli.....	34
Procedure	34
Results.....	35
Discussion	37
Experiment 3	39
Method	40
Participants	40

Stimuli.....	40
Procedure	40
Results.....	41
Discussion	46
General Discussion.....	48
Comparison with previous research	49
Adaptive Resonance Theory	52
Summary and conclusions	55
References	57
Appendix A – Child-derived lexical statistics	62
Appendix B – Adult-derived lexical statistics	64
Appendix D - ANOVA Tables for Experiments 1, 2, & 3	67
Appendix E - Error rates per word for each experiment	68
Notes.....	72

Table of Figures

Figure 1. Percentage error by word type for Experiment 1 verification task.	22
Figure 2. Percentage error by word type for Experiment 1 picture-pointing task.	26
Figure 3. Experiment 1 – 2-year olds' picture-pointing percentage error by word type.	27
Figure 4. Experiment 1 – 3-year olds' picture-pointing percentage error by word type.	28
Figure 5. Experiment 1 – 4-year olds' picture-pointing percentage error by word type.	29
Figure 6. Percent error by word type for Experiment 3.	43
Figure 7. Mean median reaction times for each word type in Experiment 3.	45

Table of Appendices

Appendix A – Child-derived lexical statistics.....	62
Appendix B – Adult-derived lexical statistics.....	64
Appendix D - ANOVA Tables for Experiments 1, 2, & 3.....	67
Appendix E - Error rates per word for each experiment.....	68

The role of neighbourhood density and probabilistic phonotactics in children's and adults' spoken word recognition

Several variables have been shown to have an effect on the speed and accuracy with which adults can recognize and repeat spoken words and non-words. These variables include the frequency with which a word appears in print or speech (e.g. Balota & Chumbley, 1984), neighbourhood density (e.g. Luce & Pisoni, 1998; Vitevitch & Luce, 1998) and the phonotactic probabilities of the phonemes that comprise a word (e.g. Vitevitch & Luce, 1999; Vitevitch, Luce, Charles-Luce, & Kemmerer, 1997).

Neighbourhood density is the number of lexical neighbours a word has – a neighbour is a word which differs from another word by substituting one phoneme with another phoneme, adding a phoneme, or deleting a phoneme. Phonotactic probability refers to the likelihood phonemes either appearing individually or in combination with each other in a particular language. Data from studies investigating the roles of these variables allow for the building of models of how the speech perception system works (e.g., the Cohort Model, Marslen-Wilson & Tyler, 1980; the Neighbourhood Activation Model [NAM], Luce & Pisoni, 1998; the TRACE model, McClelland & Elman, 1986; and Shortlist, Norris, 1994).

Models of human speech perception need to account for the way in which humans accomplish the processing of a transient, acoustic signal into recognizable speech. These models have traditionally postulated two levels of representation for sounds entering the speech perception system– a lexical level and a sublexical level. The lexical level allows a listener to recognize a spoken stimulus as a word by matching input to a representation in that part of long-term memory devoted to information about words, or, more simply,

the lexicon. Recognition of a word is the result of competitive processes among these representations in the lexicon (Luce & Pisoni, 1998; Norris, 1994). The level of representation can be thought of as independent representations of the parts of the spoken word relating to how it sounds (McClelland & Elman, 1986; Vitevitch & Luce, 1999). Sublexical could refer to syllabic, phonemic, phonetic, or featural descriptions of the sound patterns in words. In these ways, words, both old and new, and non-words can be recognized and/or repeated by a person. Thus, the lexical level is thought to be responsible for the recognition of familiar words while the sublexical level is responsible for new words or non-words. The distinction between lexical and sublexical levels of representation in spoken word recognition is similar to aspects of dual route models of visual word recognition (e.g. Coltheart, Curtis, Atkins, & Haller, 1993; Paap & Noel, 1991).

Connectionist models propose that recognition is the result of activation of units within a network. Units in connectionist models (e.g. TRACE, McClelland & Elman, 1986; Seidenberg & McClelland, 1989; SHORTLIST, Norris, 1994) encode a word's actual letter composition (orthography), its meaning (semantics), and its phonology (sound). Activation at one level spreads across the network and leads to activation at another level. It is the sum total of this activation which facilitates word recognition. Connectionist models differ from other types of models mainly in their view of how information is processed once it enters the recognition system.

In the TRACE model, for example, there is a high level of interaction between the phonemic and lexical processing levels. Featural information in TRACE is extracted from the input and encoded first. Activation spreads from the feature level to the

phoneme level and on to the word level. Word recognition is aided by feedback from word level to the phoneme level – lexical information can influence phonemic analysis. Recognition occurs when the activation of units in the network reaches a threshold level within a given time. Failure to reach threshold results in a non-word response. Words are recognized faster than some non-words because they receive feedback from the lexical level. As a result, their activation levels are increased. Non-words resembling words receive some feedback. Their levels of activation remain lower and they take longer to identify than words. Non-words that are not at all like real words do not receive feedback – their activation levels remain too far below threshold to be discriminated as words. It is the interactive nature of levels in TRACE determines the speed of recognition of input.

Norris' (1994) Shortlist model is a two-stage connectionist account of speech recognition. Information flow is entirely bottom-up, with no feedback from higher levels. The first stage generates a set of candidate words from the input. Input is encoded and candidates are generated based on the incoming phonemic information. As more input is encoded, candidates matching it more closely are sent to compete with each other in working memory. Those candidates that do not match input are removed from the candidate set. The candidate set increases until all input has been encoded. If there is a complete mismatch between input and stored lexical representations, a non-word response is given.

Stage two of Shortlist involves a process of competition between activated lexical items. Lexical items inhibit one another in competition for recognition. Shorter words are inhibited by input from other words occurring subsequent to their completion. As

words become more and more inhibited, their levels of activation decrease and their chances of being recognized decrease concomitantly. If a candidate matches a lexical representation, recognition occurs. Shortlist does not address frequency since there is no role for such high-level information. Instead, candidates are considered more - or less - plausible based on context. The activation levels of implausible candidates are reduced through inhibition from more plausible candidates, based on the context in which they are heard.

Until recently, research on spoken word recognition had been almost exclusively focussed on the roles of word frequency and neighbourhood density, either individually or in combination. The role of word frequency has been well documented in visual word recognition (e.g. Balota & Chumbley, 1985; Herdman, LeFevre, & Greenham, 1994), with high-frequency words generally receiving faster responses in various tasks than low-frequency words. In dual-route models, frequency serves to lower the level of needed activation in the lexicon for high-frequency words. Similarly, frequency has generally been an important component of spoken word recognition models. Frequency effects in connectionist models such as TRACE (McClelland & Elman, 1986) are found in the strengths of the connections between different units. Frequently encountered words have stronger connections, experience more activation between units and are recognized faster.

Neighbourhood density has also been shown to have an effect on the speed and accuracy of word recognition in various tasks (e.g., Grainger, 1990; Luce & Pisoni, 1998). In auditory word recognition, words that activate a number of similar patterns in memory (have more neighbours) are more confusable and are responded to more slowly and less accurately than those with few neighbours. Luce and Pisoni's (1998)

Neighbourhood Activation Model (NAM) incorporates both word frequency and neighbourhood density in its account of spoken word recognition. The model uses a computational formula called a Neighbourhood Probability Rule (NPR), which takes into account stimulus intelligibility, the number of neighbours a stimulus word has (its confusability), the frequency of the stimulus and the frequency of its neighbours. In short, NAM attempts to account for the structural as well as the higher-order relationships among lexical items in memory.

In NAM, recognition of a word is achieved by discriminating among items activated in memory by auditory input. Whereas frequency is paramount in many models, activation in NAM is the result of the “acoustic-phonetic similarities among the activated items” (p. 12). Input leads to activation of acoustic-phonetic patterns in memory. These patterns, if the stimulus is a word in memory, activate word-decision units. These decision units are viewed as identical to neighbourhood activation – the more neighbours a word has, the greater number of decision units it will activate.

Once activated, decision units begin monitoring higher-level word information, such as frequency, while continuing to monitor the input that activated them. Information such as word frequency serves to bias the decision units by adjusting their activation levels. The decision units continuously compute decision values using the NPR. Word recognition occurs once the output of a given decision unit reaches criterion and passes information to working memory. Thus, decision units act as intermediaries between input and higher-level word information, continuously computing new values using the NPR until criterion is reached.

In NAM, frequency and neighbourhood density have different effects on word recognition. High-frequency words are responded to and recognized more quickly than low-frequency words because they have higher activation levels in the decision system. Words with many neighbours are responded to more slowly than words with few neighbours since the overall activity levels in the decision system will be higher for the former. The criterion level for a given decision unit will take longer to reach for both low frequency and high neighbourhood density words. Thus, in NAM, frequency and neighbourhood density both affect recognition speed through increased activation, but with opposite effects on response time.

In the past five years, phonotactics has increasingly been viewed as an important variable in research on spoken word recognition. As described by Vitevitch and Luce (1998), "probabilistic phonotactics refers to the frequencies of segments and sequences of segments in syllables and words" (p. 325). Research has shown that listeners are sensitive to phonotactic information (Vitevitch & Luce, 1998, 1999). This sensitivity has even been shown in infants (Jusczyk, Luce, & Charles-Luce, 1994) and this has implications on how such information is to be represented in models of spoken word recognition.

Vitevitch and Luce (1998) investigated the role of phonotactic information in the spoken word recognition process. As Vitevitch and Luce observe, a confound exists between neighbourhood density and phonotactics: Words with many neighbours tend to have higher phonotactic probabilities, whereas words with few neighbours have less probable phonotactic patterns. Vitevitch and Luce dealt with the confound by controlling both phonotactic probability and neighbourhood information within words and non-

words. In an effort to understand the relative contribution of these two variables to lexical and sublexical levels of processing, four lists of stimuli were constructed: words with high phonotactic probabilities and many neighbours (high probability-density words), words with low probabilities and few neighbours (low probability-density words), non-words composed of high probability phonotactic patterns and derived from words with high neighbourhood density (high probability-density words), and non-words composed of low probability phonotactic patterns and derived from words with low neighbourhood density (low probability-density non-words).

Vitevitch and Luce (1998) manipulated both neighbourhood density and probabilistic phonotactics using monosyllabic words and non-words. In a shadowing task paradigm, lexicality (word vs. non-word) was blocked by group to encourage listeners to use either the lexical route (for words) or sublexical route (for non-words) to process the stimuli. While no significant effect was found for accuracy, listeners responded more slowly to words from high-density neighbourhoods that had high phonotactic probabilities than those from low-density neighbourhoods that had low phonotactic probabilities. For non-words, the opposite pattern of reaction times was obtained. Those non-words composed of highly probable phonotactic sequences with high-density ratings were responded to more quickly than those composed of low probability and low-density information. Thus, depending on the task the listeners are asked to perform and the lexicality of the stimuli used (i.e., whether the presented utterance is a word or a non-word), neighbourhood density and phonotactics can either facilitate or inhibit the recognition and pronunciation of stimuli.

These results implicate separate roles in the word recognition process for both neighbourhood density and phonotactics. The greater the number of lexical representations in a neighbourhood, the greater the level of inhibition among them. As a result, reaction times are longer. At the sublexical level, phonotactics determines recognition speed since there is assumed to be no lexical activation. As a result of having no lexical competition occurring amongst activated items, non-words with high-probability phonotactic patterns are responded to more quickly than those non-words with low-probability phonotactic patterns. While these results offer evidence that both neighbourhood density and phonotactics play differing, opposing, roles in spoken word recognition, a limitation lies in the lack of orthogonal manipulation of these variables.

Vitevitch and Luce (1999) conducted a further series of experiments in an effort to illuminate the roles of neighbourhood density and phonotactics and their effects on the lexical and sublexical levels of processing. In a comprehensive set of six experiments employing various paradigms, participants were encouraged to process stimuli either lexically or sublexically. In these experiments Vitevitch and Luce provided further support for the differential effects of neighbourhood density and phonotactics – facilitatory effects were found for probabilistic phonotactics and competitive effects from neighbourhood density.

Vitevitch and Luce used Grossberg's Adaptive Resonance Theory (ART) of speech perception (Grossberg, 1986; Grossberg, Boardman, & Cohen, 1997; Grossberg & Stone, 1986) to illustrate how phonotactics can be combined with lexical information in speech perception. The decision to use ART to explain the results of their experiments was facilitated by a contradiction between their results and the Neighbourhood Activation

Model. The NAM predicts that high-density stimuli are processed more slowly than low-density stimuli. As a result, words from high-density/high probability neighbourhoods are predicted to be recognized faster than those from low density/low probability neighbourhoods. However, words in high-density neighbourhoods typically have highly probable patterns, whereas words in low-density neighbourhoods typically have low probability patterns. Thus, NAM predicts that stimuli (both words and non-words) with high probability phonotactic patterns should be processed more slowly than low probability phonotactic stimuli. The result of their experiments, in which a different pattern of results than that predicted by NAM was found, and the absence of a segmental level of representation in NAM, led Vitevitch and Luce (1999) to use ART to explain their results.

In ART, input activates 'items' in working memory. Items are presumed to be made up of feature clusters (roughly corresponding to phones or phonemes, although Vitevitch and Luce intentionally remain equivocal about this correspondence). These items activate 'list chunks' in short-term memory. List chunks correspond to possible groupings of items, such as segments, syllables and words. In their account, Vitevitch and Luce use the term lexical to refer to list chunks that correspond to words and sublexical for those list chunks that correspond to components of those words. List chunks compete by laterally inhibiting one another, with longer ones masking shorter ones. List chunks in short-term memory that receive activation from items in working memory "send excitatory signals back to the items, establishing a *resonance*, between list chunks in short-term memory and items in working memory" (p. 377, italics in original). Over time, a resonant state develops between the best-matching, most predictive list

chunk and the items in working memory. Responses are based on resonances between the most active list chunks. In ART, these resonances are what comprise the speech percept. Responses are based on resonances between the most active list chunks and items in working memory. Thus, ART does not hold that responses are based on specific level of representation but instead on the developed resonant state.

Density and phonotactics play differing roles in the framework offered by ART. High-density and high probability words activate list chunks based on the items contained in the word and these experience inhibition from the other words in their neighbourhood. The list chunk corresponding to the high-density/high probability word dominates its inhibitors because its resonance (percept) is strongest. The resonant state developed when a low-density/low probability word serves as input is similar. However, since there is little inhibition from other list chunks, these words are processed faster than those from a high-density neighbourhood.

For non-words, list chunks are activated by items in working memory, as before. However, list chunks will only match segments or sequences of segments, as there is no existing lexical representation for a non-word. According to ART, stronger resonances will arise between items and list chunks that occur more frequently. As a result, low-density/low probability non-words will achieve weaker resonances and will be processed more slowly. Thus, ART provides a theoretical framework in which to understand the lexical role of neighbourhood density and the sublexical role of phonotactics.

Word recognition in children

Theoretical accounts of spoken word recognition in adults are based on a substantial body of experimental findings. However, research to understand the nature of lexical organization and its development in children has not advanced to the same level of

maturity. As a consequence, efforts to model word recognition in children are much less sophisticated and much less detailed. Nonetheless, sufficient experimental results are currently available to begin the task of linking adult performance to its precursors in childhood. For example, research has shown that infants are sensitive to phonotactic information contained in their native language (e.g. Friederici, & Wessels, 1993; Jusczyk et al. 1994). Other work has highlighted the role of vocabulary growth and its effect on the nature of a child's lexical representations (for a review, see Walley, 1993).

Developmental research has direct implications for segmentally based models of speech perception derived from adult data because children's lexical representations may not be the same as those of adults. In short, the task facing researchers studying spoken word recognition is to integrate adult and child data in the form of a developmentally plausible model.

Developmental models of spoken word recognition posit changes in the format of lexical representation as a consequence of new words being added to the lexicon. Early lexical representations are generally considered to be more holistic, or whole-word based, in the early stages of lexical development (Studdert-Kennedy, 1987; Walley, 1993; for an opposing viewpoint, see Gerken, Murphy, & Aslin, 1995). Though holistic, these representations are not construed as lexical. Rather, they are based on whole-word information such as prosody and phoneme structure. In conjunction with the development of the lexicon, the child moves from this holistic representation to the more adult-like segmental representation. Thus, the child begins life sensitive to the phonotactic patterns of his/her native language, moves to a more holistic representation

of speech, and progresses to a segmentally based representational structure as his/her vocabulary develops.

Analyses of children's lexicons have aided in extending understanding of the roles of probabilistic phonotactics and neighbourhood density. Until recently, phonotactic and neighbourhood density statistics have only been available for adult lexicons. However, Logan (1992) obtained neighbourhood density statistics from a computational analysis of children's lexicons. Logan computed confusability ratings for a large number of words using a database of children's speech. Confusability was based on neighbourhood density and word frequency. These analyses showed that a computational approach to children's lexicons could be applied to further delineate how words are represented and stored by children.

Using words selected from Logan (1992), Pitrat, Logan, Cockell, and Gutteridge (1995) attempted to determine the role of neighbourhood density in spoken word recognition in preschool children using a picture-pointing task. In this task, children aged two, three, and four years listened to words presented over headphones and pointed to a computer screen to indicate which one of four presented pictures represented the word they heard. Results indicated that words with relatively many neighbours were identified more accurately than words with few neighbours, and that older children identified words more accurately than younger children. However, the effect of neighbourhood density was in stark contrast with adult studies, in which words from high-density neighbourhoods are identified more slowly and less accurately than those from low-density neighbourhoods. Pitrat et al.'s results suggested that a child's lexicon might be changing as new words are added to it, perhaps requiring a different method of access.

The apparent conflict between the effect of neighbourhood density in children (Pitrat et al., 1995) and adults (e.g., Luce & Pisoni, 1998) led Logan (1999) to investigate whether another factor besides density was affecting performance. Based on the work of Vitevitch and Luce (1998) he considered the potential role of phonotactics as mitigating the effect of neighbourhood density. Children identify words with many neighbours more accurately than those with few of them. Since a neighbour is any new word which can be made from an existing word by changing one or more phonemes, then "having many neighbours would increase the frequency of the phoneme sequences common to those words" (Logan, 1999, p. 5). Thus, phonotactics could be confounded with neighbourhood density in experiments such as Pitrat et al (1995), just as Vitevitch and Luce (1998, 1999) had shown in work with adults.

To address the possibility that the probability of phonotactic sequences was the source of differences between adults and children, Logan used a partial correlation approach to analyze three child lexicons from children ages 1 ~ 2, 2 ~ 3, and 3 ~ 4. With the effect of word length removed, he found that density and phone and biphone probability were correlated within each lexicon. Pearson correlations ranged from $r = .17$ to $r = .33$. Moreover, the highest correlations were found in the 3 ~ 4 year old lexicon. An analysis of the subset of these words that were used by Pitrat et al. revealed an even more striking finding: partial correlations between neighbourhood density and phone and biphone probability for low-density words were $r = .09$ and $r = .19$, respectively, whereas they were $r = .57$ and $r = .53$, respectively, for high-density words. These results provided strong evidence for the potential role of probabilistic phonotactics in explaining the different role of neighbourhood density in adults and children. Logan's

(1999) analysis of phonotactic effects could explain why Pitrat et al.'s neighbourhood density results for children were opposite to the results for adults.

The present research was designed to further elucidate the role of probabilistic phonotactics and neighbourhood density in children's and adults' speech perception. Experiment 1 was conducted as a replication of Pitrat et al. (1995) using a) stimuli in which both density and phonotactics were orthogonally varied and b) a greater number of stimuli. Stimuli were chosen from the Logan (1992) database based on high- and low-density and high- and low- probabilistic phonotactic statistics. Experiment 2 served as an adult control in which the same picture pointing task and stimuli used in Experiment 1 were used with adults as listeners. The only difference was the use of response time as the dependent measure for adults (compared to identification accuracy for children). Experiment 3 served as an additional control in which adults were presented the same stimuli as in experiments 1 and 2 but using a naming task that is more typical of adult research.

Experiment 1

Method

Participants

Fifty-four English-learning (native English) preschool children served as listeners in this experiment. Participants attended local daycares in the Ottawa area and ranged in age from 2 to 4 years. All listeners had normal hearing and speech abilities, as determined by parental report. Based on information from a language experience questionnaire filled out by parents and caregivers, the data from 8 participants were not used because they used another language in addition to English more than 25% of the time. The data from an additional 2 participants were not used because they were 5-years

old, older than the age limit for the study. Thus, the data from fourteen 2-year olds, seventeen 3-year olds, and thirteen 4-year olds were used in the analyses. All 44 of these participants used English in the home at least 75% of the time.

Stimuli

The words used in Experiment 1 and subsequent experiments are listed in Appendix A. They were derived from words selected by Logan (1992) from the CHILDES database (McWhinney & Snow, 1985). The CHILDES database is a record of spoken interactions between young children and their caregivers or experimenters. Words from five children were selected and phonemically transcribed. Various lexical statistics were then derived, including neighbourhood density and frequency. Paul Luce calculated phonotactic probabilities for the words. Each of the lexical statistics was based on a lexicon of approximately 1200 words. The words are representative of words used by children aged 18-24 months and this corresponds to the lowest age of the children tested in the present investigation.

Frequency. Frequency information for the words used in this experiment was derived from the database of children's speech compiled by Logan (1992). Both high- and low-frequency words were used in the experiment. Words in the various neighbourhood density and phonotactic conditions were balanced with respect to mean frequency such that stimulus conditions did not differ significantly on the frequency measure.

Phonotactic Probability. Phonotactic probabilities were calculated using a frequency-weighted measure (i.e., probabilities were obtained by counting the frequency of phonemes and/or combinations of phonemes and then multiplying these values by the

frequency of the words in which they occurred). Probabilities for both individual phonemes and biphones were calculated. In the present experiment, only the biphone probabilities were used to select stimuli. Using the list of high- and low-frequency words, a median split was performed; those words having a biphone probability of greater than the median were categorized as high probability words while words which had a biphone probability of the median value or lower were categorized as low-probability words.

Neighbourhood Density. Neighbourhood density information was also derived from the Logan (1992) database. Density was calculated using a one phoneme substitution, addition, or deletion metric used in previous work (e.g., Luce & Pisoni, 1998, Pitrat et al., 1995). A median split determined high and low neighbourhood density categories. Those words having a neighbourhood density value above the median were categorized as high neighbourhood density stimuli while words having a value on this variable at or below the median were categorized as low neighbourhood density stimuli.

The stimuli were divided into categories reflecting their phonotactic probabilities and neighbourhood density information: 14 words had high phonotactic probabilities and high neighbourhood densities (High-High), 14 words had high phonotactic probabilities and low neighbourhood densities (High-Low), 14 words had low phonotactic probabilities and high neighbourhood densities (Low-High), and another 14 words had low phonotactic probabilities and low neighbourhood densities (Low-Low). Hence, four sets of experimental stimuli were created. An additional 40 words (ten targets and 30 foils) were selected at from a set of words that corresponded to pictures in a picture

database (see below) and used as practice words. None of these practice words also appeared in the experimental trials.

One word selected from the four stimulus conditions (High-High, High-Low, Low-High, or Low-Low) served as a target for a trial. Another three words, drawn from the remaining stimuli, were presented as foils. Thus, all 14 words from each of the four conditions were used as targets and the remaining words served as foils. Each target was randomly assigned three foils and placed as a list item in an array. Each list item was randomly selected by a computer from the set of 56 possible items (14 items containing targets from each word type with three foils for each list item). No word was used more than once as a target. In addition, no target was used as a foil within seven trials of its appearance as a target.

A female native speaker of American English recorded all stimuli in a sound-attenuated booth. The stimuli were low-pass filtered at 4.8 kHz and digitized at a sampling rate of 11.25 kHz. All sound stimuli were edited into individual files and stored on a computer disk.

Fifty-six different black and white line drawings corresponding to the words were used in this experiment, the majority (38) chosen from the Snodgrass and Vanderwort (1980) collection of picturable nouns. A smaller number of pictures (18) were chosen from other collections that had similar kinds of line drawings. All 40 pictures (ten targets and 30 foils) used in the practice trials were taken from the Snodgrass and Vanderwort collection and corresponded to the practice words. None of these words appeared in the experimental trials.

Procedure

All children were tested individually. Each child was seated at a table next to the experimenter in a room in the daycare. On the table in front of the child was a Power Macintosh 5300c computer. The resolution for the experiment was set to 640 by 480 pixels with a refresh rate of 72 kHz. Delivery of stimuli and collection of responses were controlled by Psyscope (v. 1.2.5) software (Cohen, MacWhinney, Flatt, & Provost, 1993). The experimenter initiated the presentation of the stimuli on each trial by pressing a key on the keyboard. The child and the experimenter each wore a pair of headphones connected to the computer. Audio files of words used in the experiment were presented over these headphones at a comfortable listening level.

The experimenter initiated each trial. A trial began with a fixation point consisting of a plus sign ('+') located in the centre of the screen. When the experimenter was sure the child was ready to proceed, a mouse key was pressed once to start the trials. Four randomly selected pictures then appeared on the screen, each placed randomly in one quadrant of the computer screen. One of the pictures was the target word for the trial with the three other pictures serving as foils. A sound file corresponding to the target picture played over the headphones at the same time as the pictures appeared on the screen. The child was then asked to point to the picture matching the word he/she just heard. In the event the child did not hear the word or did not respond, the experimenter pressed a key on the keyboard to repeat the word and the child could try again.

Once the child made his/her response, the experimenter pressed a numerical key - '1', '2', '3', or '4' - to indicate which picture the child pointed to. In addition, the experimenter entered the number of times the audio stimulus had to be repeated before

the child responded. This information was then stored in a data file on the computer.

After the information about a trial had been entered, the experimenter pressed the mouse button to move on to the next trial. The screen again cleared to white and the fixation point again appeared to indicate the start of a new trial.

Practice trials. There were 10 practice trials in the experiment. All practice trials were conducted in the same manner as the experimental trials but used different target and foil words, as outlined above. All target and foil pictures used in the practice trials were drawn from the Snodgrass and Vanderwort (1980) collection.

Experimental trials. There were 56 trials in the experiment proper. On each trial one of the 56 target words plus three foils were presented. Each word appeared as a target only once. In sum, all children saw all targets in a different, random order. Each target had the same randomly selected foils for each presentation.

Verification trials. The 56 pictures corresponding to target words were presented after the experimental trials to verify that the children knew all the target words used in the experiment. Each picture was presented, one at a time, in the centre of the screen. The child was asked to name the picture. No sound files played in this part of the experiment. Once the child named the picture, the experimenter clicked the mouse button and the next picture was presented. The experimenter recorded all correct and incorrect responses by hand. In this manner, all children saw and attempted to name all 56 targets used in the experiment.

Results

Verification data from Experiment 1 was first analyzed to determine if the children were familiar with the words/pictures used. The criteria for evaluating the

verification results were as follows. If the word used by the child to identify the picture was identical to the name for the picture used in the picture-pointing task then the trial was categorized as a match (e.g., rooster - rooster). If the two words were semantically equivalent (e.g., cube - square, rooster - chicken) the trial was considered a mismatch, as were any unrelated words produced (e.g., giraffe - animal). An identical match was required in order to make the verification results meaningful for purposes of excluding trials in the picture pointing task analysis. A semantically equivalent or related word produced by a child does not indicate whether the child knows the specific word presented. In short, it was essential to the logic of Experiment 1 that only exact matches be considered as evidence that the child knew a word.

Error rates for each child for each word type were calculated and converted to percentage of trials for which the child gave a response. This was necessary because, due to an inability of some children to complete the task, not all participants were able to respond to all 56 pictures in the verification trials. Thus, only those trials in which a child gave a response were used to calculate the percentage of error for each word type. A 3 (age) by 2 (ND) by 2 (PP) mixed analysis of variance (ANOVA) was conducted on the error rates for the verification trials. The analysis revealed a significant effect only for the ND by PP interaction, $F(1, 41) = 11.81$, $MSE = 916.47$, $p = .001$, $\eta^2 = .224$, observed power = .918. Participants made the fewest errors on words with high density and high phonotactics ($M = 13.5\%$, $SE = 1.45\%$). The next lowest error rate was for LDLP words, ($M = 16.12\%$, $SE = 2.05\%$), followed by HDLP words ($M = 18.35\%$, $SE = 1.7\%$). The highest error rate for the verification trials was found in the LDHP word category ($M = 20.53\%$, $SE = 2.12\%$). Figure 1 shows error rates for each word type.

The pattern of results obtained in the verification task does not lend itself to a simple interpretation. The lexical and sublexical variables compared in this analysis were unlikely the primary reason for the pattern of results observed. Probably the biggest factor underlying error in the verification task was the fact a child would provide a name for a picture but not the name that was associated with the picture in this experiment. For example, the picture corresponding to the word 'cube' was often named as 'block' or 'square'. These alternate names were not wrong per se but they did not match the actual word used in the picture-pointing task. Because the primary goal of the verification task was to ensure that the child actually knew the word and its picture referent used in the picture-pointing task, a conservative criterion was used to score the verification task data. In turn, this made 'errors' more prevalent because synonyms for the pictures were counted as correct responses. Overall, the pattern of results observed in the verification task likely reflects more an idiosyncratic choice of label for a picture than the lexical characteristics of the words used in the experiment.

Insert Figure 1 about here

Figure 1. Percentage error by word type for Experiment 1 verification task.

HD = High (Neighbourhood) Density, LD = Low-density

HP = High (Phonotactic) Probability, LP = Low Probability

The dependent variable for the picture-pointing task in Experiment 1 was the error rate associated with the identification of pictures corresponding to the spoken words presented. No item analyses were conducted for two reasons. First, the set of stimuli used in Experiment 1 exhausted the possible stimuli available for children in this age group meeting the word type criteria. The usual reason for an item analysis is to assure that the effects observed could be generalized to a broad set of stimuli that would share the same characteristics as the particular words used in an experiment. The words used in this experiment are the only words that fit the criteria for density and phonotactics in this model of a preschool child's lexicon. Therefore, an item analysis was deemed unnecessary. Second, data from the verification task serve as a form of item analysis, albeit for the purpose of eliminating from the ANOVAs individual words with excessive error rates. (see Raaijmakers, Schrijnemakers & Gremmen, 1999, for an argument against item analysis as a method for assessing differences among stimuli.)

Errors for each stimulus condition in the picture-pointing task were calculated for 44 children and analyzed in subgroups according to age categories: 2-year-olds ($n=14$), 3-year-olds ($n=17$), and 4-year-olds ($n=13$). The verification trials were used to measure whether the child could name the pictures presented, thus indicating whether he/she held that particular item in his/her lexicon. Any items that could not be named were excluded from the analysis of picture pointing data. Because this led to some conditions having fewer than 14 stimuli, errors were converted to percentages for each child according to the number of the total pictures for a word type (HDHP, HDLP, LDHP, and LDLP) that he/she failed to identify in the picture-pointing task.

Once these initial steps were completed, an 2 X 2 X 3 mixed analysis of variance (ANOVA) with neighbourhood density (high, low) and phonotactic probabilities (high, low) as within groups variables, and age (2-, 3-, 4-years) as a between groups variable was carried out. In this and subsequent analysis an alpha of .05 was used. A significant main effect of age was obtained, $F(2,41) = 7.95$, $MSE = 103.22$, $p < .001$, $\eta^2 = .28$, observed power = .941. Two-year-olds, as expected, had the highest error rate ($M = 8.4\%$, $SE = 1.36\%$), with three-year-olds at approximately half that error rate ($M = 4.4\%$, $SE = 1.23\%$). Four-year-olds made the fewest errors ($M = 0.61\%$, $SE = 1.41\%$). Post hoc tests conducted on age revealed a significant difference in the error rates between 2- and 4-year olds, Mean Difference = 7.797%, $SE = 1.96\%$, $p = .001$. The Bonferroni correction was used to control alpha for these tests.

No significant main effect for neighbourhood density (ND) was observed, $F(1,41) = 2.60$, $MSE = 36.09$, $p = .115$, $\eta^2 = .06$, observed power = .35. In addition, the interaction between ND and age was not significant, $F(2,41) = 2.53$, $MSE = 36.09$, $p = .092$, $\eta^2 = .11$, observed power = .478. Despite the lack of a statistically significant interaction between ND and age, a suggestive pattern is present in the data that may be obscured due to a floor effect for the 4-year-olds. The interaction between ND and age was most evident in the 3-year-olds, whose mean error rate ($M = 6.458\%$, $SE = 1.616\%$) in the high-density condition was almost triple than that in the low-density condition ($M = 2.313\%$, $SE = 1.219\%$). Error rates were equally high for 2-year-olds in both high- and low-density neighbourhoods ($M = 8.865\%$, $SE = 1.781\%$ and $M = 7.939\%$, $SE = 1.344\%$, respectively). Again, 4-year-olds had few errors in either condition ($M = .275\%$, $SE = 1.848\%$ and $M = .937\%$, $SE = 1.394\%$, respectively).

A significant main effect of phonotactic probability was obtained, $F(1,41) = 5.59$, $p = .023$, $MSE = 41.44$, $\eta^2 = .12$, observed power = .637. Errors were higher in the Low Probability (LP) condition ($M = 5.62\%$, $SE = 0.67$) than in the High Probability (HP) condition ($M = 3.31\%$, $SE = 1.10$). The interaction between phonotactic probability and age was not significant, $F(2,41) = 1.583$, $MSE = 41.44$, $p = .218$, $\eta^2 = .072$, observed power = .316, nor was the interaction between ND and phonotactic probability significant, $F(1,41) = 3.77$, $MSE = 58.26$, $p = .059$, $\eta^2 = .084$, observed power = .474. Participants across all age groups made significantly more errors when low probability words were presented than high probability words, regardless of their ND status.

Figure 1 shows the interaction between ND and PP. Despite a non-significant interaction, a pattern emerges that may be obscured statistically due to low power (i.e., a small number of participants in each condition with relatively high variability associated with each condition). The effect stems primarily from the difference between the HDLP condition and the other conditions, suggesting that high probability phoneme sequences may mitigate confusability in high-density neighbourhoods. As illustrated in Figure 2, there were relatively high error rates in the HDLP condition. Three-year olds had consistently fewer errors than 2-year-olds in all other conditions. Four year olds had so few errors (none in the HDHP condition) that it could be considered a floor effect. Figures 3, 4, and 5 illustrate the error rates for the picture-pointing task for 2-, 3-, and 4-year olds, respectively.

Insert Figures 2, 3, 4 & 5 about here

Figure 2. Percentage error by word type for Experiment 1 picture-pointing task.

HD = High (Neighbourhood) Density, LD = Low-density

HP = High (Phonotactic) Probability, LP = Low Probability

Figure 3. Experiment 1 – 2-year olds' picture-pointing percentage error by word type.

HD = High (Neighbourhood) Density, LD = Low-density

HP = High (Phonotactic) Probability, LP = Low Probability

Figure 4. Experiment 1 – 3-year olds' picture-pointing percentage error by word type.

HD = High (Neighbourhood) Density, LD = Low-density

HP = High (Phonotactic) Probability, LP = Low Probability

Figure 5. Experiment 1 – 4-year olds' picture-pointing percentage error by word type.

HD = High (Neighbourhood) Density, LD = Low-density

HP = High (Phonotactic) Probability, LP = Low Probability

Finally, the expected three-way interaction among age, ND, and PP was not significant, $F(2,41) < 1$, $MSE = 58.26$, $p = .466$, $\eta^2 = .037$, observed power = .173.

In an effort to compare the results of Experiment 1 to those of earlier research with adults by Vitevitch and Luce (1998, 1999), a paired samples t-test was conducted on the HDHP and LDLP words. Vitevitch and Luce used only words chosen to be either high-density/high-probability or low-density/low-probability. Although participants made more errors in the LDLP condition, ($M = 3.737\%$, $SE = .973\%$) than in the HDHP condition ($M = 2.948\%$, $SE = 1.223\%$), a paired samples t-test revealed no significant difference in error rates for the participants in this study, $t(1, 43) = -.785$, $p = .437$.

Discussion

The anticipated pattern of results was to have a three-way interaction among age, density, and phonotactics such that for younger children (2- and 3-year-olds), words with less frequent phonotactic sequences would have a higher error rate than words with frequent phonotactic sequences. The results for Experiment 1 show main effects for both age and probabilistic phonotactics, but no main effect of neighbourhood density. Density effects were expected to be minimal for younger children as the size of their lexicons may limit lexical competition effects (Logan, 1992; 1999). It was expected that older children would be affected by neighbourhood density more than phonotactics, as their lexicons would be larger. Specifically, older children were predicted to have a higher error rate for words that had many neighbours compared to words that had few neighbours, as is the pattern of results observed in adults.

The underlying reason for these predictions was that as children increase in age, the size of their lexicons grows, enabling lexical effects (i.e., neighbourhood density) to

override sublexical effects. The small size of the younger children's lexicons may minimize lexical effects and enable sublexical effects (i.e., phonotactic probabilities) to predominate. The nonsignificant finding that words with low probability phoneme sequences produced more error only in high-density neighbourhoods (see Figure 1) suggests that high probability phoneme sequences may mitigate the effects of confusability in high-density neighbourhoods where its effects should be most evident. Because this finding was nonsignificant, such accounts must be viewed as speculative. Nonetheless, they are suggestive of an interaction that parallels adult findings obtained with words and non-words.

Vitevitch and Luce (1998, 1999), in their studies with adults, showed a pattern of results different from the present research. Vitevitch and Luce (1999) found that, for adults processing both words and non-words, neighbourhood density and probabilistic phonotactics have independent effects. Density has an effect at the lexical level, with high-density words showing a slower reaction time and greater error rate. Probabilistic phonotactics affects the sublexical level, with non-words composed of highly probable sequences being named faster than those composed of less probable sequences. Thus, the role of one variable or another may shift depending on the characteristics of the stimuli.

The present results indicate no effect of neighbourhood density for children for the words in this study, but a strong effect of phonotactic probabilities. This may be interpreted in two ways. Either the effect of density only occurs when the lexicon reaches a certain size or phonotactics continues to play a role in the processing of words during childhood. It may be impossible to choose between these explanations because of the difficulty associated with holding lexical size constant during development. Whether

phonotactics plays a role mitigating confusability among words in high-density neighbourhoods for adults, as it seems to for children, also remains unclear. While based on a marginal interaction, the suggestion is intriguing. Models such as ART do not rule out the possibility that phonotactic effects could occur in adults, only that they are likely to be overwhelmed by lexical effects.

Methodological differences may also account for the differences observed between Experiment 1 and previous work with adults. Variations in task and stimulus selection procedures could be responsible for the presence of phonotactic effects in words that were found in Experiment 1. Experiments 2 and 3 attempted to address some of these issues using adults as listeners. Experiment 2 used the same procedure and stimuli as Experiment 1, while Experiment 3 used the same stimuli but a different and more typical adult procedure.

Experiment 2

Experiment 2 was designed to provide data from adults to compare with the data obtained from children in Experiment 1. The rationale for an adult comparison using the same experimental procedure was threefold. First, earlier work by Pitrat et al. (1995) suggested that adults and children were differentially affected by neighbourhood density, such that children recognized high-density words more easily whereas adults recognized low-density words more easily. However, no direct comparison between children and adults was actually done by Pitrat et al., only a comparison of the children's performance in their experiment with results from adults obtained by other researchers (e.g., Luce, 1986). Experiment 2 provided an opportunity to compare the same stimuli and virtually the same procedure for both adults and children.

Second, a picture identification task had not been previously used to assess phonotactic effects in adults. Thus, it was useful to know whether such a task would produce effects similar to those obtained in other investigations of phonotactics and neighbourhood density (e.g., Vitevitch & Luce, 1999). The dynamics of the picture-pointing task differ in many dimensions from typically used tasks, such as naming tasks. The most obvious difference between the picture pointing task and other tasks is the use of the picture referents. Participants are required to integrate the information they hear with its referent in a picture display, potentially creating a more complex task than a naming task, for example. Also, the mapping of the picture referent to a manual response may also affect performance differently compared to simply repeating a spoken word.

Finally, the stimuli chosen for use in Experiment 1 were derived from a database of words used by children. It may be possible that the range of densities and phonotactic probabilities are more limited than if words were derived from an adult database (e.g., Kucera & Francis, 1967). Moreover, it is also possible that the specific values of the density and phonotactic statistics derived from child lexicons may differ from the corresponding values derived from an adult lexicon. For example, it is conceivable that a word with a relatively large number of neighbours during childhood could have relatively few neighbours in an adult lexicon. It is even more plausible that a word with few neighbours in childhood could have many neighbours in an adult lexicon. Therefore, another role for Experiment 2 was to validate the stimuli chosen for use in Experiment 1.

The major difference between Experiment 1 and Experiment 2 was the dependent measure. Whereas Experiment 1 used percent correct identification (and number of repetitions) as the dependent variables, Experiment 2 used RT. The presence of ceiling

effects associated with the task requires this change: adults naming pictures of real objects do not make any significant number of errors. The assumption in Experiment 2 was that adults' reaction times would vary as a function of the word's lexical characteristics in a manner comparable to a naming or lexical decision task. Although not identical to the dependent measure obtained in Experiment 1, it was assumed that Experiment 2 would enable comparison between performance of adults and children such that fast RTs in adults would be associated with correct responses in children.

In short, the goal of Experiment 2 was to collect RT data from adults to compare to the error data obtained from the children tested in Experiment 1. Experiment 2 served as a conceptual replication of Experiment 1 in order to assess the effects of neighbourhood density and phonotactic probabilities in adults.

Method

Participants

Twenty-six native speakers of English served as participants in this study. Participants ranged in age from 16 to 51 years. All participants were volunteers drawn from the Carleton University community. All listeners had normal hearing and speech abilities, as determined by self-report.

Stimuli

Words. Experiment 2 used the same words as Experiment 1 (see Appendix A).

Pictures. Experiment 2 used the same picture stimuli as in Experiment 1.

Procedure

Experiment 2 used the same procedure as in Experiment 1, with the exceptions that adult participants were tested in a sound-attenuated booth in a laboratory setting and

indicated their responses with the dominant hand using a keyboard. Participants pressed a key on the keyboard labeled with '1', '2', '3', and '4' to make their responses. Keys labeled with these numbers were chosen so as to be as equidistant as possible from the idle position (index finger of dominant hand resting in the centre of the four keys) and to match the pattern of quadrants on the screen during a trial. The upper left quadrant of the screen was '1', the upper right '2', the lower left '3' and the lower right '4'. The 'o', 'p', 'l' and ';' keys were used for right-handed participants and the 'w', 'e', 's', and 'd' keys were used for left-handed participants. The experimental session took approximately 20 minutes to complete.

Results

Reaction times were screened prior to analysis. Median RTs exceeding 2.5 standard deviations above or below the mean were excluded from further analysis. Responses were also excluded from the analysis if participants either entered a single incorrect response, indicating they pressed the wrong key and were not aware of the error, or they incorrectly responded and followed this response by a second keystroke reflecting the correct response. As it is unlikely any of the participants would not know the words used in the experiment, these errors were treated as equivalent and removed from further analysis. Overall, the error rate for Experiment 2 was 2.6%. A 2 X 2 within groups ANOVA conducted on the error data between word types (HDHP, HDLP, LDHP, LDLP) revealed no significant effects for Neighbourhood Density, $F(1, 25) < 1$, $MSE = .0018$, $p = .746$, $\eta^2 = .004$, observed power = .061, or Probabilistic Phonotactics $F(1, 25) < 1$, $MSE = .003$, $p = .615$, $\eta^2 = .01$ observed power = .078. The interaction (ND X PP)

was also not significant, $F(1, 25) < 1$, $MSE = .0012$, $p = .424$, $\eta^2 = .122$ observed power = .122.

As reaction times tend to exhibit a positive skew, median RTs were used in the analysis as the measure of central tendency for performance in each condition for each participant. RTs meeting the requirements outlined above were submitted to a 2 X 2 within-groups ANOVA, with neighbourhood density (high, low) and phonotactic probabilities (high, low) as within groups variables and RT as the dependent measure. A main effect of density was found, $F(1,25) = 7.03$, $MSE = 5259.97$, $p = .014$, $\eta^2 = .219$, observed power = .722. Responses to low-density words ($M = 1333$ ms, $SE = 37.55$) were significantly slower than those to high-density words ($M = 1296$ ms, $SE = 38.96$). This effect of density is opposite to that typically observed in adults.

No main effect of probabilistic phonotactics was found, $F(1,25) < 1$, $MSE = 5120.29$, $p = .461$, $\eta^2 = .022$, observed power = .111.

The interaction of neighbourhood density and probabilistic phonotactics was also not significant, $F(1,25) < 1$, $MSE = 3009.73$, $p = .659$, $\eta^2 = .008$, observed power = .071.

To determine if the position in which the target appeared was a factor in the RT data, means were calculated for each listener for each quadrant. A 2 X 2 repeated measures ANOVA in which quadrant the target appeared (Horizontal position, Vertical position) was conducted. A significant main effect for Horizontal position (whether the target appeared in quadrants 1 or 2, or in quadrants 3 or 4, i.e. Upper or Lower) of the target was found, $F(1, 25) = 7.702$, $MSE = 14034.84$, $p = .01$, $\eta^2 = .236$, observed power = .76. Targets that appeared in the upper half of the screen ($M = 1282.548$ ms, SE

= 36.744 ms) were responded to significantly faster than targets appearing in the lower half of the screen ($M = 1347.029$ ms, $SE = 41.951$ ms).

The main effect of Vertical position (whether the target appeared in quadrants 1 or 3, or in quadrants 2 or 4, i.e. Left or Right) was not significant, $F(1, 25) < 1$, $MSE = 5987.03$, $p = .465$, $\eta^2 = .236$, observed power = .11. The interaction between Horizontal position and Vertical position was also not significant, $F(1, 25) = 4.148$, $MSE = 8079.81$, $p = .052$, $\eta^2 = .142$, observed power = .499.

The presence of a position effect in the picture-pointing task is an artifact of the task requirement and a potential confound. However, an analysis of target type as a function of position indicated that targets were evenly distributed across all four positions, minimizing any effect of target position on reaction times for a particular word type.

A paired groups t-test was also conducted on the two word types at the extremes of each lexical variable, namely HDHP words and LDLP words. This was done to have a comparison with the results of earlier research, which used just these two categories of stimuli (see Vitevitch and Luce, 1998; 1999). The analysis revealed a significant difference in RTs for these two groups, $t(1, 25) = -2.424$, $p < .023$. Participants recognized HDHP words ($M = 1292.69$ ms, $SE = 41.38$ ms) faster than LDLP words ($M = 1340.9$ ms, $SE = 39.69$ ms).

Discussion

The results of Experiment 2 revealed a main effect for neighbourhood density but no significant effect for probabilistic phonotactics or the interaction between these two variables. However, the effect of neighbourhood density obtained in Experiment 2 was

opposite to that typically obtained for adults – participants in Experiment 2 responded more slowly to words from low-density neighbourhoods than to words from high-density neighbourhoods.

Several reasons may account for this result. In previous research (Vitevitch & Luce, 1998; 1999), stimuli were not selected so as to independently manipulate neighborhood density and phonotactic probability information. Stimuli in these studies have typically been chosen to be either high or low on both neighbourhood density and probabilistic phonotactics. These stimuli would correspond to the HDHP and LDLP stimuli in the present experiments.

Stimuli in the present study were selected so as to differentiate ND and phonotactics while balancing other important variables such as word frequency and word length. Thus, when stimuli are orthogonally manipulated, the effect of neighbourhood density may be such that low-density words are prone to slower reaction times.

In light of considering the possibility of stimulus effects being responsible for the reversed effect of density in Experiment 2 the role of range effects needs to be addressed. Because these stimuli were derived from a children's lexicon, and the statistics used to select them were also based on this lexicon. The effect of variation in lexical parameters for these words may be attenuated for an adult listener. Nonetheless, despite a more limited range of variation, an effect of density was present in Experiment 2. This suggests that, at least as far as density is concerned, the use of stimuli based on a child's lexicon was not a problem per se.

A second possibility to explain the greater number of errors for low-density words in this experiment is the nature of the task. A picture-pointing task has not been used in

previous research with adults to evaluate the roles of neighbourhood density and probabilistic phonotactics. Participants had to rely on information from both the visual and auditory domains to be successful at this task, and this may have contributed to error rates and/or slower reaction times. The analysis of target position implicates the role of visual information in response times. Those targets appearing in a position in the upper half of the screen (Quadrants 1 or 2) were responded to more quickly than those appearing in the lower half. In the current experiment, participants saw four pictures appear simultaneously on the screen and heard a target word presented at the same time. Ideally, a picture-pointing task using words derived from an adult lexicon to examine density and phonotactics is necessary to resolve some of these issues.

Overall, the unusual effect of density observed in Experiment 2 is difficult to interpret without additional research. Experiment 3 provides an additional adult comparison for the results obtained in experiments 1 and 2. If the same pattern of results obtained in Experiment 2 is also obtained in the naming task used in Experiment 3, then it is possible to attribute the effects observed in Experiment 2 to stimulus characteristics rather than task differences.

Experiment 3

The goal of Experiment 3 was to use the same stimuli used in experiments 1 and 2 in a different task with adults as listeners. The rationale for Experiment 3 was to provide a link from experiments 1 and 2 to experiments done by other researchers investigating phonotactics and neighbourhood density (e.g., Luce & Vitevitch, 1999). A word-naming (shadowing) task was employed in Experiment 3. The naming task is typically used in word recognition research and is sensitive to most lexical variables. Lexical effects were expected to predominate because of the size of the lexicon in adults compared to

children. Phonotactic effects were expected to be minimal because previous research (Vitevitch & Luce, 1998; 1999) has shown that these effects are present in adults only for non-words.

Method

Participants

Twenty-nine native speakers of English served as participants in this study. Participants ranged in age from 16 to 32 years. All participants were volunteers drawn from the introductory psychology subject pool at Carleton University and received partial credit for a course requirement. All listeners had normal hearing and speech abilities, as determined by self-report.

Stimuli

The stimuli used in Experiment 3 were the same as those in experiments 1 and 2. However, only the auditory stimuli were used: the same 56 words used as experimental stimuli and the same 10 words used as practice stimuli in experiments 1 & 2. Thus, there was a total of 66 items; 10 practice words and 56 experimental words.

Procedure

All participants were tested individually. Each participant was seated in a sound-attenuated booth in a laboratory setting. On the shelf in front of the participant was a Power Macintosh 5300c notebook computer on which instructional text messages directed to the participant appeared. Participants wore a pair of Genexxa HPM010 headphones with a microphone attached to a voice-operated relay box. A Power Macintosh 5300c computer controlled the presentation of the auditory stimuli, prompts, and instructions, and recorded reaction times. The participant controlled presentation of

the stimuli by using the computer mouse button. The participant pressed a button on the mouse to initiate each trial.

A trial proceeded as follows: A message appeared telling the participant to press the mouse button to start the trials. After this, the screen cleared and a fixation point consisting of a plus sign ('+') appeared in the centre of the screen. When the participant was ready to proceed, he/she pressed the mouse key once to start the trial. A randomly selected sound file corresponding to a target picture in experiments 1 and 2 played over the headphones. The participant's task was to repeat aloud the word he/she heard as quickly and as accurately as possible. Once the participant had made a response, the screen again cleared and the fixation point reappeared to indicate the start of a new trial. Participants were recorded in order to monitor production errors and inadvertent tripping of the voice key by non-speech sounds. This information was used to eliminate error trials from analyses of the RT data and to determine error rates.

Results

RTs were screened prior to analysis. Median RTs exceeding 2.5 standard deviations above or below the mean were excluded from further analysis. RTs from correct responses were submitted to a 2 X 2 repeated measures ANOVA with neighbourhood density (high, low) and phonotactic probabilities (high, low) as within groups variables and RT as the dependent measure. Error data was also analyzed using an ANOVA with the same design. Errors caused by equipment failure were noted at the time of the experiment and excluded from the final analysis. In addition, the experimenter verified all participants' responses by listening to the responses from the experiment on audiotape. Responses were screened to ensure that the response that

triggered the voice key was the word presented over the headphones. This was necessary as the voice-operated relay is incapable of distinguishing word responses from noise or one word from another when registering a reaction time.

Overall, the error rate for errors not involving equipment failure across conditions in Experiment 3 was 2.03%.ⁱ This error rate is comparable to most naming tasks, where error rates are typically no greater than 5%. An ANOVA conducted on the error data revealed no significant main effect of density, $F(1, 28) < 1$, $MSE = .0017$, $p < .663$, $\eta^2 = .007$, observed power = .071 or probabilistic phonotactics, $F(1, 28) = 3.59$, $MSE = .0009$, $p < .068$, $\eta^2 = .114$, observed power = .449. Participants did make a slightly greater percentage of errors on words with low probability phoneme patterns ($M = 2.53\%$, $SE = .55\%$) than words with high probability patterns ($M = 1.47\%$, $SE = .0049\%$). Figure 6 shows the error data for word types.

The interaction between neighbourhood density and probabilistic phonotactics was also not significant for the error analysis, $F(1, 28) < 1$, $MSE = .0019$, $p < .473$, $\eta^2 = .019$, observed power = .108.

Insert Figure 6 about here

As in Experiment 2, median RTs were used as the measure of central tendency. A 2 X 2 repeated measures ANOVA on the RT data revealed a significant main effect of neighbourhood density, $F(1, 28) = 4.501$, $MSE = 2069.49$, $p = .043$, $\eta^2 = .138$,

Figure 6. Percent error by word type for Experiment 3.

HD = High (Neighbourhood) Density, LD = Low-density

HP = High (Phonotactic) Probability, LP = Low Probability

observed power = .535. Participants named words from high-density neighbourhoods ($M = 1019.103$ ms, $SE = 37.631$ ms) faster than those from low-density neighbourhoods ($M = 1037.026$ ms, $SE = 35.517$ ms). Note that this effect is opposite to that typically obtained with adults but is consistent with the density effect obtained in Experiment 2 using these stimuli.

In addition, a strong main effect of phonotactic probability was found, $F(1, 28) = 48.325$, $MSE = 1724.86$, $p < .001$, $\eta^2 = .633$, observed power = 1.0. Participants named words with highly probable phonotactic patterns ($M = 1001.259$ ms, $SE = 35.597$ ms) significantly faster than those with less probable patterns ($M = 1054.871$ ms, $SE = 37.476$ ms).

The 2-way interaction between neighbourhood density (high/low) and phonotactic probability (high/low) was also significant, $F(1, 28) = 10.693$, $MSE = 1846.77$, $p = .003$, $\eta^2 = .276$, observed power = .884. Figure 7 shows the mean median reaction times for ND and PP across word type.

Insert Figure 7 about here

Words from low-density neighbourhoods with high probability phonotactic patterns (LDHP) were named the fastest overall ($M = 997.172$ ms, $SE = 35.81$ ms) whereas words from low-density neighbourhoods with low probability phonotactic patterns (LDLP) were slowest to be named ($M = 1076.879$ ms, $SE = 36.261$ ms).

Figure 7. Mean median reaction times for each word type in Experiment 3.

HD = High (Neighbourhood) Density, LD = Low-density

HP = High (Phonotactic) Probability, LP = Low Probability

Words from high-density neighbourhoods with highly probable phonotactic patterns (HDHP, $\underline{M}$ = 1005.345 ms, $\underline{SE}$ = 36.187 ms) were named faster than high-density words with less probable patterns (HDLP, $\underline{M}$ = 1032.862 ms, $\underline{SE}$ = 39.643 ms).

A paired groups t-test was also conducted on the two word types at the extremes on each lexical variable, namely HDHP words and LDLP words. This was done to compare with the results of earlier research, which used just these two categories (see Vitevitch & Luce, 1998; 1999). The analysis revealed a significant difference in RTs for these two groups, $t(1, 28) = -6.569$, $p < .001$. Participants were faster to name HDHP words ($\underline{M}$ = 1005.34 ms, $\underline{SE}$ = 36.19 ms) than LDLP words ($\underline{M}$ = 1076.89 ms, $\underline{SE}$ = 36.26 ms).

Discussion

The goal of Experiment 3 was to investigate the roles of neighbourhood density and probabilistic phonotactics for the words used in experiments 1 & 2 in a task that was more typical of adult speech recognition research. The speeded naming (shadowing) task requires participants to name spoken words as quickly and as accurately as they can. Lexical (ND) effects were expected to predominate because of the size of the lexicon in adults compared to children. Phonotactic effects were expected to be minimal because previous research has shown that these effects are present in adults only for non-words (Vitevitch & Luce, 1999). Moreover, because the naming task used in Experiment 3 is more typical of tasks used with adults than the picture-pointing task used in Experiment 2, the anomalous density effect that was observed in Experiment 2 could reasonably be expected to be absent in Experiment 3.

In contrast to the expected pattern of results, the effect of neighbourhood density in Experiment 3 showed the same pattern of results as it did in Experiment 2. This effect is the opposite of that typically observed in research on the effect of ND in adults. In the present experiment and in Experiment 2, high-density words were named faster and more accurately overall than low-density words. The replication of this ‘reversed’ density effect in both the picture-pointing task and the naming task suggests that characteristics of the stimuli, rather than the task, underlie this effect. In addition, Experiment 3 also yielded an effect of PP in which words with high probability phoneme sequences were named faster than words with low probability sequences, an effect that previously had only been observed in non-words (Vitevitch & Luce, 1999). Finally, the effect of phonotactics was contingent on density. The difference in RT between words with high and low probability phoneme sequences were larger for words from low-density neighbourhoods than for words from high-density neighbourhoods.

Experiment 3 resulted in a variety of unexpected effects. Because at least one of these effects, the ‘reversed’ density effect, was also found in Experiment 2, stimulus characteristics are implicated as the reason for the effects more so than task differences between the picture-pointing and naming tasks. The words used in experiments 1, 2, and 3 were drawn from a database of words used by children 2 years of age. The statistics used to select and balance the stimuli were also derived from the same children’s database (see Logan, 1992). The range of values underlying the high and low dimensions for ND and PP were selected to be as far apart as possible on these dimensions. In addition, words were balanced for frequency and length. In contrast, stimuli used in previous work were selected from an adult lexicon using adult-derived lexical statistics.

As a result of using child-derived statistics to determine stimulus categories for adults, the range for ND and PP values would be truncated compared to words selected from an adult lexicon. It is also possible that words in a child lexicon that have many neighbours could have only a relatively small number of neighbours in an adult lexicon or vice versa. Similarly, words with high probability phoneme sequences in a child lexicon may not have high probability phoneme sequences in an adult lexicon. In addition, the words used in experiments 1, 2, and 3 utilized ND and PP variables in an orthogonal design whereas previous work (e.g., Vitevitch & Luce, 1999) used only the HDHP and LDLP categories. All of these differences are possible reasons for the different results obtained in the present investigation compared to previous work.

In order to partially address the potential differences between stimuli selected from child and adult lexicons in terms of ND and PP values, the correlation between the ND values derived from the child lexicon (i.e., the values used to select the stimuli used in the present experiments) and the ND values derived from an adult lexicon was calculated. Similarly, the correlation between child-derived PP and adult PP values was also calculated. In both cases, a positive correlation was obtained for ND and PP, $r = .452$, $p < .001$ and $r = .505$, $p < .001$, respectively. These positive correlations suggest that the stimuli selected from the child lexicons generally sorted themselves into high/low categories that corresponded to adult lexical statistics. The mean and standard deviation for frequency, density, biphone probability, and length for both high-and low-density words are provided in Appendix B.

General Discussion

The present set of experiments sought to further clarify the roles of neighbourhood density and probabilistic phonotactics in spoken word recognition in

adults and children. In previous research, neighbourhood density and probabilistic phonotactics have been shown to exhibit independent effects on spoken word recognition and production. Density has been shown to dominate the lexical level of recognition, whereas phonotactics dominates the sublexical level. Words from high-density neighbourhoods typically cause more confusability than words from low-density neighbourhoods. High phonotactics, on the other hand, aids the processing of non-word stimuli such that non-words with high probability sequences are processed more quickly than non-words with low probability sequences.

Comparison with previous research

The results of the present research seem to run contrary to the findings of other researchers (Vitevitch & Luce, 1998, 1999), though no direct comparison can be made between the adult and child data, due to different dependent measures. Vitevitch and Luce (1998, 1999) used only adult listeners and a more limited type of stimuli in which neighbourhood density and probabilistic phonotactic information were not manipulated independently. Instead, Vitevitch and Luce used words high on both variables and low on both variables. The reason cited for this is the difficulty in finding stimuli that can be matched on other lexical variables while controlling for density and phonotactics. While acknowledging that this is an important issue, stimuli chosen under such constraints may inadvertently affect the outcome of experiments. The present research sought to independently manipulate these variables, thus creating four categories of stimuli, HDHP words, HDLP words, LDHP words, and LDLP words.

In addition, Vitevitch and Luce used non-words constructed to closely resemble words in their two categories, and blocked by lexicality in naming and categorization

tasks. Non-words were not used in the present research. Finally, whereas Vitevitch and Luce used stimuli based on an adult lexicon, the present investigation used words based on child-derived statistics selected from a child language database.

The present Experiment 1 used a picture-pointing task with 2-, 3- and 4-year old participants identified pictures corresponding to high probability words almost twice as accurately as low probability words. The absence of the anticipated three-way interaction among age, density and phonotactics may have been due to the minimal error rate for the 4-year old participants. An alternative explanation is that the effect of density may only manifest itself in older children. In either case, further work is required to evaluate these possibilities.

Experiment 2 employed the same task and the same stimuli, but used adult listeners as participants and RT as the dependent measure instead of error rates. Adults, unlike children, did not make significantly more errors on any type of word, and error rates were very low overall (2.6%). More importantly, however, pictures matching words from high-density neighbourhoods were identified faster than those from low-density neighbourhoods. This is opposite the effect typically found for ND in adults. However, as noted above, this 'reverse' effect of density may be the result of characteristics of the stimuli used, as the words were selected based on child-derived lexical statistics. A replication of Experiment 2 using words selected to be at the extreme range of values for both ND and PP in an adult database of spoken words is necessary to assess the role of phonotactics in naming words. In addition, the orthogonal design used in the present work should be employed using adult-derived stimuli (see below).

Experiment 3 employed the same stimuli as experiments 1 and 2 in a naming task more typical of research with adults. As in Experiment 2, error rates across word types were relatively low (2.05%) and there were no significant main effects or interactions for the error data. However, the same 'reverse' density effect as in Experiment 2 was found for the RT data. Though not the same type of task as Experiment 2, participants responded to words from high-density neighbourhoods significantly faster than those from low-density neighbourhoods.

The significant effect of phonotactics indicates that it might play a mitigating role when naming words in either high- or low-density neighbourhoods. Words with high phonotactics were named faster than those with low phonotactics. Although the effect of phonotactics was only significant in the naming task, it suggests a role for phonotactics in the recognition of words beyond its more limited role in non-word recognition, as proposed by Vitevitch and Luce (1998, 1999).

The interaction of ND and PP in Experiment 3 is of particular interest. Differences between words with high and low probability phoneme sequences were larger when they came from low-density neighbourhoods than when they came from high-density neighbourhoods. High probability phoneme sequences facilitated processing more in low-density neighbourhoods than in high-density neighbourhoods. Conversely, low probability phoneme sequences slowed processing more in low-density than in high-density neighbourhoods. Again, further studies are needed to clarify the role of phonotactics and density in word naming, using words selected based on adult-derived, orthogonally varying ND and PP statistics.

Adaptive Resonance Theory

Vitevitch and Luce (1999) used Grossberg's (1997) Adaptive Resonance Theory (ART) to account for the effects of density and phonotactics in words and non-words. In their adaptation of ART, neighbourhood density and phonotactics both play a role in spoken word recognition. In ART, items, consisting of feature clusters in working memory, and list chunks, consisting of possible groupings of those items, establish a resonance through a series of excitatory signals. Activation of list chunks is assumed to be a result of their frequency. Over time, a strong resonance develops between the list chunks in short-term memory and the items in working memory. List chunks can consist of words, or parts of words, and longer list chunks (such as words) dominate shorter list chunks. Thus, list chunks can be lexical or sublexical in nature. Items are, by their nature, sublexical.

Vitevitch and Luce (1999) argue that, for words, neighbourhood density dominates any facilitatory effects phonotactics may provide. This is so because the items in short-term memory form the strongest resonance with the list items in working memory. High-density words have multiple lateral inhibitory signals converging on them and, thus, take longer to establish a dominant resonant state. Low-density words experience less lexical competition for a resonant state as there are fewer inhibitory links converging on them. Larger list chunks (words) will dominate smaller ones. As a result, any benefit of a word having high probability phoneme sequences is overshadowed by this dominant resonance for words, an effect associated with density.

Vitevitch and Luce (1999) argue that phonotactics is the major determinant of response times for non-words. Non-words with common phonotactic patterns are named

faster than those with less common patterns. The resonant state that develops for non-words is between the items and the list chunks, as before, but there is no competition from lexical entries, although some weak lexical activation may occur. Those sublexical chunks occurring more frequently establish stronger resonances with items in working memory and are not inhibited by longer list chunks corresponding to words. Thus, the strongest resonance develops between list sublexical list chunks that occur most often. In this way, phonotactics displays the opposite pattern of density for non-words, with high phonotactics resulting in faster and more accurate responses as observed by Vitevitch and Luce (1998, 1999).

It is less clear how ART could be used to explain the results of the present experiments, in which only words were presented and both neighbourhood density and phonotactics were controlled. In the present set of experiments, neighbourhood density effects were the opposite of what was found by Vitevitch and Luce (1998, 1999). Words from high-density neighbourhoods were less prone to error and were responded to more quickly than low-density words. In addition, the phonotactic information was found to play a role in the naming of words. Highly probable phonotactic patterns seem to facilitate the naming of words in both high-density and low-density neighbourhoods, although the effect is largest for low-density neighbourhoods. Thus, phonotactics appears to play a significant role in the naming of words as well as non-words.

ART does not rule out the role of phonotactics in recognizing words, nor does it dismiss the possibility of lexical effects in naming non-words. Rather, neighbourhood density is assumed to *dominate* when processing word stimuli and phonotactics to dominate when processing non-words. Common phonotactic patterns may aid the

processing and/or production of both words and non-words. Conversely, uncommon phonotactics may hinder such processing.

Accounting for the pattern of adult results is complicated because of the need to assimilate the results of experiments 2 and 3 with previous findings obtained by Vitevitch and Luce (1998, 1999). However, the results obtained from the children tested in Experiment 1 are more compatible with ART. Participants in Experiment 1 exhibited no significant effect of density, but did show an effect of phonotactics. For children, stronger resonant states appear to develop for sublexical items, much as is assumed for non-words, and these effects dominate. The limited size of a child's lexicon permits the sublexical level to dominate identification of the spoken word. Nonetheless, extending the age range tested in Experiment 1 to 5- and 6-year-olds would help to clarify the size of the lexicon required for lexical effects to play a greater role in word recognition. This would also entail using a different task or making the picture-pointing task more difficult because of the ceiling effects observed for 4-year-olds in Experiment 1.

Experiments 2 and 3 both showed an effect of density opposite of what has been found in other research. Adults in these experiments named high-density words more quickly than low-density words. As these effects were consistent across both adult experiments, even though they were different tasks, they may be a function of the stimuli themselves rather than the task. The presence of phonotactic effects for words was also anomalous. A possible reason for the difference in results is the orthogonal versus non-orthogonal manipulation of density and phonotactics. Vitevitch and Luce (1998, 1999) selected only HDHP and LDLP words for use in their experiment whereas I selected not only HDHP and LDLP words, but also HDLP and LDHP words. Perhaps the use of a

limited set of stimuli by Vitevitch and Luce created a context in which the word recognition system responded differently than when all four density and phonotactic conditions were used. The reason for such an explanation comes from experiments comparing words blocked or mixed according to their frequency. When words are presented blocked by frequency, frequency effects are minimized whereas they are present when words are presented mixed by frequency. Perhaps an analogous phenomenon occurs when stimuli are selected to include limited subsets of words (i.e., Vitevitch & Luce) compared to the full set of words when selecting stimuli based on density and phonotactic statistics. It seems prudent to explore this possibility by directly comparing these stimulus parameters within the same study.

Summary and conclusions

In summary, the goal of the present investigation was to further elucidate the roles played by neighbourhood density and probabilistic phonotactics in word recognition. The present research points to a possible role of phonotactics in the naming of words. Probabilistic phonotactics was found to play a role for children performing a picture-pointing task, as well as for adults performing a naming task. A 'reverse' effect of neighbourhood density, where adults responded more slowly to words in high-density neighbourhoods, was found. Reasons for the atypical pattern of the present results are likely due to stimulus differences.

Stimuli used in the present research differ from those used in previous research in that they were all selected from a children's database and were selected to vary orthogonally on both neighbourhood density and phonotactics according to lexical statistics derived from the same database. Further studies using 4- and 5- year old

participants are needed using these words to clarify the role of ND and PP in children's word recognition. Similarly, replication of experiments 2 and 3 using words selected from an adult lexicon to vary along the same dimensions would also help to resolve some of these issues.

The unexpected pattern of results in the present study may initially appear to suggest an anomaly when viewed in light of previous work. In contrast, I would argue that the use of words in which both phonotactics and density were varied provided a unique opportunity to independently evaluate the effect of these variables that previous research did not. The consistent presence of the 'reverse' neighbourhood density effect across two very different experimental tasks suggests that the effect is real. Taken together, the 'reverse' density effect and the finding of phonotactic effects in words call into question the interpretation of previous findings regarding neighbourhood density and phonotactics. It is clear that further work is needed to sort out some of the issues surrounding these effects. Nonetheless, their existence indicates a need to re-evaluate the mechanisms thought to underlie spoken word recognition in adults and children.

References

- Balota, D. A., & Chumbley, J. I. (1984). Are lexical-decisions a good measure of lexical access? Journal of Experimental Psychology: Human Perception and Performance, 10, 340-357.
- Chomsky, N., (1965). Aspects of the theory of syntax. Cambridge, MA: MIT Press.
- Chomsky, N., and Halle, M. (1968). The sound pattern of English. New York: Harper and Row.
- Cohen J.D., MacWhinney B., Flatt M., and Provost J. (1993). PsyScope: A new graphic interactive environment for designing psychology experiments. Behavioral Research Methods, Instruments, and Computers, 25, 257-271.
- Coltheart, M., Curtis, B., Atkins, P., & M. (1993). Models of reading aloud: Dual-route and parallel-distributed processing approaches. Psychological Review, 100, 589-608.
- Friederici, A. D. & Wessels, J. M. I. (1993). Phonotactic knowledge of word boundaries and its use in infant speech perception. Perception and Psychophysics, 54, 287-295.
- Gerken, L. A., Murphy, W. D., & Aslin, R. N. (1995). Three- and Four-year old's perceptual confusion for spoken words. Perception and Psychophysics, 57, 475-486.
- Grainger, J. (1990). Word frequency and neighbourhood frequency effects in lexical decision and naming. Journal of Memory and Language, 29, 228-244.
- Grossberg, S. (1986). The adaptive self-organization of serial order in behavior: Speech, language, and motor control. In E.C. Schwab and H.C. Nusbaum (Eds.) Pattern recognition by humans and machines: Vol. 1. Speech Perception. (pp. 187-294). New York: Academic Press.

- Grossberg, S., Boardman, I. & Cohen, J. (1997). Neural dynamics of variable-rate speech categorization. Journal of Experimental Psychology: Human Perception and Performance, 23, 483-503.
- Grossberg, S., & Stone, G. O. (1986). Neural dynamics of word recognition and recall: attentional priming, learning, and resonance. Psychological Review, 93, 46-74.
- Herdman, C., & LeFevre, J. (1992). Individual differences in the efficiency of word recognition. Journal of Educational Psychology, 84, 95-102.
- Herdman, C., LeFevre, J., & Greenham, S. (1994). Implicating the lexicon: Base-word frequency effects in pseudohomophone naming. Journal of Experimental Psychology: Human Perception and Performance, 20, 575-590..
- Jusczyk, Peter W. (1997). The discovery of spoken language. Cambridge, MA: The MIT Press.
- Jusczyk, P. W., Luce, P. A., & Charles-Luce, J. (1994). Infants' sensitivity to phonotactic patterns in the native language. Journal of Memory and Language, 33, 402-420.
- Kucera, H., & Francis, W. N. (1967). Computational analysis of present-day American English. Providence, RI: Brown Univ. Press.
- Lieberman, A., and Mattingly, I. (1985). The motor theory revisited. Cognition, 21, 1-36.
- Logan, J. (1992). A computational analysis of young children's lexicons. Research On Spoken Language Processing: Technical Report 8, Indiana University.
- Logan, J. (1999). Spoken Word Recognition in Preschoolers: The Effect of Phonotactics And Neighbourhood Density. Paper presented at the annual meeting of the Canadian Society for Brain, Behavior and Cognitive Science, Edmonton, Alberta, June, 1999.

- Luce, P. A. (1986). A Computational analysis of uniqueness points in auditory word recognition. Perception and Psychophysics, 39, 155-158.
- Luce, P. and Pisoni, D. (1998). Recognizing spoken words: The Neighbourhood Activation Model. Ear & Hearing, 19, 1-36.
- MacWhinney, B., & Snow, C. (1985). The child language data exchange system. Journal of Child Language, 12, 271-296.
- McClelland, J. L. & Elman, J. L. (1986). The TRACE model of speech perception. Cognitive Psychology, 18, 1-86.
- McNeill, D. (1966). Developmental psycholinguistics. In F. Smith and G.A. Miller (Eds.) The genesis of language. Cambridge, MA: MIT Press.
- Marslen-Wilson, W. (1987). Functional parallelism in spoken word-recognition. Cognition, 25, 71-102.
- Marslen-Wilson, W. D. & Tyler, L. K. (1980). The temporal structure of spoken language understanding. Cognition, 19, 1-36.
- Morton, J. (1969) Interaction of information in word recognition. Psychological Review, 76, 165-178.
- Newman, R. S., Sawusch, J. R., & Luce, P. A. (1997). Lexical neighbourhood effects in phonetic priming. Journal of Experimental Psychology: Human Perception and Performance, 23(3), 873-889.
- Norris, D.G. (1994b) Shortlist: A connectionist model of continuous speech recognition. Cognition , 52,189-234.
- Paap, K. R. & Noel, R. W. (1991). Dual-route models of print to sound: Still a good horse race. Psychological Research, 53, 13-24.

- Pinker, S. (1994) The language instinct. New York, NY: William Morrow and Company, Inc.
- Pitrat, A., Logan, J., Cockell, J., & Gutteridge, M. E. (1995). The role of phonological neighbourhoods in the identification of spoken words by preschool children. Poster Presented at the annual meeting of the Canadian Society for Brain, Behavior and Cognitive Science, Halifax, Nova Scotia, June, 1995.
- Raaijmakers, Jeroen G. W., Schrijnemakers, Joseph M. C., & Gremmen, Frans (1999). How to deal with "The Language-as-Fixed-Effect Fallacy": Common Misconceptions and Alternative Solutions. Journal of Memory and Language, 41, 416-426.
- Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. Psychological Review, 96, 523-568.
- Snodgrass, J. G., & Vanderwort, M. (1980). A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity and visual complexity. Journal of Experimental Psychology: Human Learning and Memory, 6, 174-215.
- Studdert-Kennedy, M. (1987). The phoneme as a perceptuomotor structure. In A. Allport, D. MacKay, W. Prinz, and E. Scheerer (Eds.), Language perception and production, New York: Academic Press.
- Vihman, M. (1996). Phonological development: The origins of language in the child. Cambridge: Blackwell Publishers, Inc.
- Vitevitch, M. S. and Luce, P. A. (1998). When words compete: Levels of processing in the perception of spoken words. Psychological Science, 9(4), 325-329.
- Vitevitch, M. S. and Luce, P. A. (1999). Probabilistic phonotactics and neighbourhood activation in spoken word recognition. Journal of Memory and Language, 40, 374-408.

Vitevitch, M. S., Luce, P. A., Charles-Luce, J. and Kemmerer, D. (1997). Phonotactics and syllable stress: Implications for processing of spoken nonsense words. Language and Speech, 40(1), 47-62.

Walley, A. C. (1993). The role of vocabulary development in children's spoken word recognition and segmentation ability. Developmental Review, 13, 286-350.

Appendix A – Child-derived lexical statistics

Child-derived lexical statistics for words used in Experiments 1, 2, & 3.

Legend. HDHP = high-density, high phonotactic probability; HDLP = high-density, low phonotactic probability; LDHP = low-density, high phonotactic probability; LDLP = low-density, low phonotactic probability.

word	trans	freq	density	biphone prob	length
HDHP					
baby	bebi	35	4	0.0134	4
belt	bElt	2	3	0.0214	4
chairs	CErz	6	5	0.0336	4
chicken	CIkN	19	2	0.0224	4
clock	klak	4	2	0.0169	4
cups	k^ps	14	2	0.0190	4
fox	faks	7	5	0.0371	4
pig	pIgi	5	2	0.0208	4
puppy	p^pi	2	2	0.0163	4
socks	saks	8	5	0.0370	4
stair	stEr	2	3	0.0190	4
stick	stIk	2	2	0.0162	4
swing	swIG	6	3	0.0198	4
track	tr@k	18	3	0.0210	4
MEAN		9.29	3.07	0.0224	4
STD. DEV.		9.03	1.16	0.0074	0
HDLP					
arm	arm	12	2	0.0043	3
bird	bRd	19	2	0.0044	3
broom	brum	14	2	0.0113	4
cube	kyub	2	2	0.0056	4
ears	Irz	7	4	0.0033	3
flower	flWR	8	2	0.0117	4
food	fud	13	2	0.0043	3
foot	fUt	18	3	0.0053	3
house	hawse	2	2	0.0118	4
mouse	maws	5	2	0.0100	4
spill	spIl	4	2	0.0098	4
spool	spul	1	4	0.0116	4
spoon	spun	18	2	0.0096	4
towel	tWL	6	2	0.0034	3
MEAN		9.21	2.36	0.0076	3.57
STD. DEV.		6.14	0.72	0.0033	0.51
High Density words					
MEAN		9.25	2.71	0.0150	3.79
STD. DEV.		7.86	1.05	0.0096	0.42

Appendix A (continued)

word	trans	freq	density	biphone prob	length
LDHP					
barn	barn	26	0	0.0270	4
bottle	batL	5	0	0.0217	4
button	b^tN	6	1	0.0263	4
crayon	kreaN	1	1	0.0232	5
ducks	d^ks	9	1	0.0230	4
monkey	m^Gki	17	0	0.0231	5
muffin	m^f^N	1	0	0.0203	5
pants	p@ns	6	1	0.0358	4
pillow	pIlo	19	1	0.0219	4
pony	poni	7	1	0.0238	4
rooster	rustR	4	0	0.0224	5
sweater	swEtR	3	0	0.0202	5
teddy	tEdi	12	1	0.0221	4
trees	triz	4	1	0.0218	4
MEAN		8.57	0.57	0.0238	4.36
STD. DEV.		7.13	0.49	0.0038	0.50
LDLP					
cookie	kUki	23	0	0.0157	4
desk	dEsk	1	0	0.0071	4
drum	dr^m	8	1	0.0171	4
giraffe	JR@f	4	0	0.0028	4
glove	gl^v	3	1	0.0046	4
ladder	l@dR	8	0	0.0122	4
lion	lY^N	16	0	0.0097	4
seesaw	sisA	9	0	0.0083	4
shovel	S^vL	3	0	0.0044	4
skate	sket	7	1	0.0158	4
tiger	tYgR	29	1	0.0139	4
turtle	tRtL	6	0	0.015	4
TV	tivi	3	0	0.0064	4
wheel	hwil	3	0	0.0103	4
MEAN		8.79	0.29	0.0102	4
STD. DEV.		7.98	0.45	0.0046	0
Low Density words					
MEAN		8.68	0.43	0.0170	4.18
STD. DEV.		7.71	0.50	0.0081	0.39

trans = phonetic transcription; **freq** = frequency; **biphone prob** = biphone probability;
length = number of phonemes in word

Appendix B – Adult-derived lexical statistics

Adult-derived lexical statistics for words used in Experiments 1, 2, & 3.

Legend. HDHP = high-density, high phonotactic probability; HDLP = high-density, low phonotactic probability; LDHP = low-density, high phonotactic probability; LDLP = low-density, low phonotactic probability.

word	trans	freq	density	biphone prob	length
HDHP					
baby	bebi	62	5	1.0050	4
belt	bElt	29	17	1.0158	4
chairs	CErz	66	20	1.0088	4
chicken	CIkN	37	2	1.0099	4
clock	klak	20	19	1.0106	4
cups	k^ps	45	0	1.0070	4
fox	faks	13	9	1.0102	4
pig	pIgi	8	19	1.0088	4
puppy	p^pi	2	7	1.0049	4
socks	saks	4	11	1.0101	4
stair	stEr	2	12	1.0226	4
stick	stIk	39	18	1.0230	4
swing	swIG	24	12	1.0093	4
track	tr@k	38	10	1.0263	4
MEAN		27.79	11.50	1.0123	4.00
STD. DEV.		21.17	6.51	0.0069	0.00
HDLP					
arm	arm	94	11	1.006	3
bird	bRd	31	23	1.0031	3
broom	brum	2	10	1.0135	4
cube	kyub	1	2	1.0069	4
ears	Irz	29	6	1.0028	3
flower	fIWR	23	5	1.0071	4
food	fud	147	11	1.0020	3
foot	fUt	70	10	1.0014	3
house	hWsi	591	7	1.004	4
mouse	mWsi	10	14	1.0038	4
spill	spIl	1	13	1.0133	4
spool	spul	1	9	1.0099	4
spoon	spun	6	11	1.01	4
towel	tWL	6	9	1.0007	3
MEAN		72.29	10.07	1.0060	3.57
STD. DEV.		155.41	4.91	0.0043	0.51
High Density words					
MEAN		50.04	10.79	1.0092	3.79
STD. DEV.		111.17	5.70	0.0065	0.42

Appendix B (continued)

word	trans	freq	density	biphone prob	length
LDHP					
barn	barn	29	11	1.0227	4
bottle	batL	76	10	1.0093	4
button	b^tN	10	6	1.0082	4
crayon	kreaN	1	2	1.0175	5
ducks	d^ks	9	4	1.0092	4
monkey	m^Gki	9	1	1.0133	5
muffin	m^f^n	1	2	1.0065	5
pants	p@ns	9	11	1.0332	4
pillow	pIlo	8	5	1.0164	4
pony	poni	10	6	1.0102	4
rooster	rustR	3	3	1.0273	5
sweater	swEtR	14	3	1.0075	5
teddy	tEdi	.	7	1.0137	4
trees	triz	59	7	1.0212	4
MEAN		18.31	5.57	1.0154	4.36
STD. DEV.		23.19	3.34	0.0081	0.50
LDLP					
cookie	kUki	1	5	1.0029	4
desk	dEsk	65	3	1.0198	4
drum	dr^m	11	10	1.0101	4
giraffe	JR@f	1	1	.	4
glove	gl^v	9	3	.	4
ladder	l@dR	19	10	1.0070	4
lion	lY^n	17	6	1.0049	4
seesaw	sisA	1	1	1.0054	4
shovel	S^vL	5	4	.	4
skate	sket	1	11	1.0100	4
tiger	tYgR	1	3	1.0042	4
turtle	tRtL	7	0	1.0026	4
TV	tivi	8	8	1.0069	4
wheel	hwil	56	7	1.0066	4
MEAN		14.43	5.14	1.0073	4.00
STD. DEV.		20.46	3.61	0.0048	0.00
Low Density words					
MEAN		16.30	5.36	1.0119	4.18
STD. DEV.		21.48	3.42	0.0079	0.39

trans = phonetic transcription; **freq** = frequency; **biphone prob** = biphone probability;
length = number of phonemes in word

Appendix C – Parental report questionnaire for Experiment 1.

Questionnaire

Child's name _____

1. What is your child's date of birth? _____ (day/month/year)

2. What language(s) do you speak in your home?

English: No ____ ; Yes ____ ; Proportion of time ____ %

French: No ____ ; Yes ____ ; Proportion of time ____ %

Other (specify): _____ Proportion of time ____ %

_____ Proportion of time ____ %

3. In the past year, has your child had any ear infections? No ____ ; Yes ____

If yes, how many? _____

4. Has your child ever been diagnosed with, or demonstrated, a hearing problem?

No ____ ; Yes ____

5. Has your child ever been diagnosed with, or demonstrated, a speech disorder?

No ____ ; Yes ____

Appendix D - ANOVA Tables for Experiments 1, 2, & 3

Experiment 1

Effect	F	MSE	p	eta ²	Observed Power
AGE	F(2,41) = 7.95	103.22	p < .001	.28	.941
ND	F(1,41) = 2.6	36.09	p = .115	.06	.35
PP	F(1,41) = 5.59	41.44	p = .023	.12	.637
ND X PP	F(1,41) = 3.77	58.26	p = .059	.084	.474
AGE X ND	F(2,41) = 2.53	36.09	p = .092	.11	.478
AGE X PP	F(2,41) = 1.583	41.44	p = .218	.072	.316
AGE X ND X PP	F(2,41) < 1	58.26	p = .466	.037	.173

Experiment 2

Effect	F	MSE	p	eta ²	Observed Power
ND	F(1,25) = 7.03	5259.97	p = .014	.219	.722
PP	F(1,25) < 1	5120.29	p = .461	.022	.111
ND X PP	F(1,25) < 1	3009.73	p = .659	.008	.071

Experiment 3

Effect	F	MSE	p	eta ²	Observed Power
ND	F(1,28) = 4.501	2069.49	p = .043	.138	.535
PP	F(1,28) = 48.325	1724.86	p < .001	.633	1.0
ND X PP	F(1,28) = 10.693	1846.77	p = .003	.276	.884

Appendix E - Error rates per word for each experiment
 Experiment 1 - Verification trials

word	% error	word	% error
Arm	11.4	Monkey	4.5
Baby	4.5	Mouse	6.8
Barn	22.7	Muffin	13.6
Belt	15.9	Pants	9.1
Bird	11.4	Pig	11.4
Bottle	20.5	Pillow	18.2
Broom	31.8	Pony	36.4
Button	15.9	Puppy	4.5
Chairs	9.1	Rooster	45.5
Chicken	11.4	Seesaw	22.7
Clock	13.6	Shovel	27.3
Cookie	11.4	Skate	47.7
Crayon	15.9	Socks	2.3
Cube	61.3	Spill	50.0
Cups	15.9	Spool	68.1
Desk	15.9	Spoon	11.4
Drum	20.5	Stair	43.2
Ducks	4.5	Stick	11.4
Ears	0	Sweater	11.4
Flower	6.8	Swing	15.9
Food	34.1	Teddy	15.9
Foot	9.1	Tiger	11.4
Fox	41.0	Towel	9.1
Giraffe	15.9	Track	20.5
Glove	11.4	Trees	2.3
House	11.4	Turtle	15.9
Ladder	18.2	TV	9.1
Lion	4.5	Wheel	6.8

Appendix E - Error rates per word for each experiment (continued)

Experiment 1 - Picture-pointing trials

word	% error	word	% error
Arm	6.8	Monkey	0
Baby	2.3	Mouse	6.8
Barn	11.4	Muffin	2.3
Belt	6.8	Pants	9.1
Bird	4.5	Pig	4.5
Bottle	11.4	Pillow	13.6
Broom	13.6	Pony	6.8
Button	9.1	Puppy	0
Chairs	2.3	Rooster	13.6
Chicken	2.3	Seesaw	15.9
Clock	11.4	Shovel	2.3
Cookie	6.8	Skate	29.5
Crayon	2.3	Socks	0
Cube	34.1	Spill	9.1
Cups	11.4	Spool	40.9
Desk	9.1	Spoon	6.8
Drum	6.8	Stair	9.1
Ducks	4.5	Stick	6.8
Ears	0	Sweater	0
Flower	2.3	Swing	9.1
Food	6.8	Teddy	4.5
Foot	6.8	Tiger	4.5
Fox	18.2	Towel	0
Giraffe	0	Track	9.1
Glove	6.8	Trees	0
House	4.5	Turtle	4.5
Ladder	13.6	TV	4.5
Lion	2.3	Wheel	0

Appendix E - Error rates per word for each experiment (continued)

Experiment 2

word	% error	word	% error
Arm	3.8	Monkey	0.0
Baby	0.0	Mouse	3.8
Barn	0.0	Muffin	0.0
Belt	0.0	Pants	0.0
Bird	3.8	Pig	0.0
Bottle	0.0	Pillow	0.0
Broom	0.0	Pony	0.0
Button	0.0	Puppy	0.0
Chairs	0.0	Rooster	0.0
Chicken	3.8	Seesaw	0.0
Clock	0.0	Shovel	0.0
Cookie	0.0	Skate	0.0
Crayon	3.8	Socks	0.0
Cube	0.0	Spill	0.0
Cups	3.8	Spool	3.8
Desk	0.0	Spoon	0.0
Drum	0.0	Stair	0.0
Ducks	0.0	Stick	0.0
Ears	0.0	Sweater	0.0
Flower	3.8	Swing	3.8
Food	0.0	Teddy	3.8
Foot	0.0	Tiger	0.0
Fox	0.0	Towel	0.0
Giraffe	3.8	Track	0.0
Glove	0.0	Trees	3.8
House	0.0	Turtle	0.0
Ladder	0.0	TV	0.0
Lion	0.0	Wheel	0.0

Appendix E - Error rates per word for each experiment (continued)

Experiment 3 (production errors only, not including equipment failure)

word	% error	word	% error
Arm	3.45	Monkey	0.00
Baby	0.00	Mouse	0.00
Barn	0.00	Muffin	0.00
Belt	3.45	Pants	0.00
Bird	0.00	Pig	0.00
Bottle	0.00	Pillow	3.45
Broom	0.00	Pony	0.00
Button	10.34	Puppy	3.45
Chairs	0.00	Rooster	10.34
Chicken	0.00	Seesaw	0.00
Clock	3.45	Shovel	0.00
Cookie	0.00	Skate	0.00
Crayon	10.34	Socks	3.45
Cube	3.45	Spill	0.00
Cups	0.00	Spool	3.45
Desk	13.79	Spoon	0.00
Drum	3.45	Stair	0.00
Ducks	0.00	Stick	0.00
Ears	10.34	Sweater	0.00
Flower	3.45	Swing	0.00
Food	6.90	Teddy	0.00
Foot	0.00	Tiger	0.00
Fox	3.45	Towel	0.00
Giraffe	0.00	Track	0.00
Glove	6.90	Trees	0.00
House	0.00	Turtle	0.00
Ladder	0.00	TV	3.45
Lion	3.45	Wheel	0.00

Notes

' Three words - Barn (LDLP), Cookie (LDHP), and Rooster (LDHP) - had a higher rate of error than other words in the experiment. Barn had an error rate of 34%, Cookie 48%, and Rooster an error rate of 52%. Taken together, these stimuli accounted for 39% of all incorrect responses in Experiment 3. Examination of the recorded responses revealed multiple equipment failure errors when these words were presented. All the errors made on these words were due to the voice-operated relay triggering a response as the word was presented over the headphones. As a result, these responses, and responses to other stimuli that resulted in the same type of error, were removed from the error analysis. Thus, the error rate reported reflects only errors where the participant did not say the word presented over the headphones.