DOCTOR OF PHILOSOPHY

The acquisition of phonemic constraints : implications for models of phonological encoding

Taylor, Conrad Frazer

Award date:
2003

Awarding institution:
University of Wales, Bangor

University of Wales, Bangor

Prifysgol Cymru

The Acquisition of Phonemic Constraints: Implications for Models of Phonological Encoding

By

Conrad F. Taylor

A thesis submitted to the School of Psychology, University of Wales Bangor, in partial fulfilment of the requirements for the degree of Doctor of Philosophy. The research was sponsored by the Economic and Social Research Council through a Ph.D. Studentship: Award Number R42200034409.

October 2003

Acknowledgements

Firstly, I would like to thank Dr. George Houghton for his supervision, and for developing my academic skills. I would also like to thank the other members of my Ph.D. committee, Dr. Charles Leek, Dr Neil Dugdale, and Dr. Emmanuel Pothos. Thanks are also due to Dr. Paloma Mari-Beffa for her advice on statistics, and for giving me the chance to develop my teaching skills. I would also like to thank the administration team, especially Linda Shutt, Sue Peet, Kath Chitty, and Dilys Hughes, whose help during the latter stages of my Ph.D. was much appreciated.

Of course, no one gets through a Ph.D. without the support of friends, so thank you (in alphabetical order) to, Caroline Bowman, Neil Fisher, Harris Ingle, Sam Johnson, Dan Loach, Matthew Paul, Val Randle, Ryan Taylor, Enlli Thomas, Annalise Whittaker, Jenny Wilson, and anyone who I've forgotten (Sorry!). My liver may never forgive you though!

I could not of finished this thesis without the support of my family. Especially my Dad and my Wife, Yvonne. Thank you to you both for putting up with me (a.k.a. The Grumpy Old Hector), and for telling me that I could do it when I had my doubts.

Thanks are also due to Ken Kutaragi for providing the entertainment when I wasn't working on the thesis, and to www.hardradio.com and www.bbc.co.uk/radio3 for entertainment when I was working on the thesis.

Finally, thanks to "The Boss" (No, not Bruce Sprigsteen!). Give me a hefty shove in the right direction, okay!?

Contents

Summary

The thesis presented here is a study of the acquisition of phonemic constraints, and of the rates at which the constraints are upheld, as demonstrated in speech errors. These experiments were motivated by Dell et al. (2000), who tested whether phonotactic constraints could be learnt through recent linguistic experience (i.e. reading nonsense syllables). The results of these studies showed that the phonotactics present in the stimuli were followed, and that the speech production system is sensitive to recent experience. The experiments also showed that the syllable position constraint was not upheld as strongly as expected.

The present study began by replicating these experiments. The general findings of the Dell et al. (2000) study were supported, but other aspects of the study were not. Weaknesses were found in the statistical analyses provided by Dell and colleagues, and the syllable position constraint was upheld at a lower rate than that found in the Dell et al. study. Dell et al. provide no data regarding either the rate of acquisition, or the durability, of the learned constraints.

The next set of experiments explored the time-course and durability of the learning. This was done by adding a further section to the experiment that reversed the previously learned constraints, and analysing the speech errors for signs of "confusion" (i.e. continuing to use the previous constraints). These experiments showed that there was a period following the reversal of the constraints in which participants followed the previous constraints.

The final part of this study modified the paradigm for the auditory modality. The stimuli were played to the participants through headphones, and the participants repeated the syllables that they heard. This resulted in an increase in the error corpus, while still producing the same results as the visual version of the paradigm. The results of these experiments are considered in relation to models of speech production, and the implications for these models are discussed.

Chapter 1: Introduction

It was Lashley (1951) who first drew our attention to the problem of serial order in behaviour. He suggested that ordered behaviour could not be purely the result of associations between elementary responses, but that there has to be a hierarchical plan or schema that determines the order of these responses. Serial order is particularly relevant to language (i.e. the constituent phonemes of a syllable, word, phrase, etc. must be produced in the correct order – this is discussed in more detail in the next section), and this has led to several models of phonological encoding that propose solutions to this problem (e.g. Dell, 1986, 1989; Dell, Juliano, & Govindjee, 1993; Elman, 1990; Hartley & Houghton, 1996; Jordan, 1986; Lapointe & Dell, 1989; Roeloffs, 1996, 1997; Shattuck-Hufnagel, 1979; Vousden, Brown, & Harley, 2000; Wickelgren, 1969).

Although these models all include a serial-order mechanism, they have different approaches to the modelling of phonological encoding. More specifically, they begin the representation at different points. For example, Dell's (1986) model begins with the concepts associated with an object that is to be named (e.g. the concepts "furry", "tail", "purr", etc. being associated with "cat"). But, the Hartley and Houghton (2000) model is a representation of word repetition, and so begins with the input of a string of phonemes. Therefore, there is no need for lexical access to be performed. However, all the models must produce a string of phonemes in the correct order, whether it is in order to name an object, or to reproduce a sequence of words.

In addition to serial-order, models of phonological encoding usually represent the *phonotactics* of a given language (a notable exception being the model by Vousden, Brown, & Harley, 2000). Phonotactics are statements of regularities that define how phonemes may be combined to form legal words within a given language. Phonotactics exist within every language. For example, in English the phonemes /p/ and /k/ cannot occur next to each other within a syllable, and in Japanese every syllable must end with a vowel or the consonant /n/. A variety of methods are used to represent knowledge of phonotactics, but they all place constraints on the order and positioning of phonemes, and attempt to reproduce the type and frequency of errors demonstrated in behavioural studies.

A problem would arise for models of phonological encoding if it could be demonstrated that the frequency of some types of speech error were different from that previously recorded. A series of experiments conducted by Dell, Reed, Adams, and Meyer (2000) suggested that the *syllable position constraint* (the tendency for phonemes, when moved in error, to occupy the same syllabic position as in the intended syllable) was not upheld as strongly as had previously been thought. Ellis (1980) reported that this constraint was upheld on 84 – 96% of occasions. However, Dell et al. suggest that this constraint is upheld between 68 – 77%.

Therefore, this thesis intends to demonstrate that some of these models are inadequate explanations of phonological encoding, due to the way in which they represent the phonotactics of the language being spoken, and the frequency at which constraint-breaking errors occur.

In order to do this, I will first discuss the problem of serial order and the various categories of models that incorporate such a mechanism. The methods used to represent the phonotactics of a language within these models will also be considered. Specific examples of these models will be given where appropriate. And finally, we will look at the paradigm developed by Dell et al. (2000), and examine the results from this present study that demonstrate the possible short-comings of some models of language.

Serial Order and Models of Phonological encoding

Serial-Order Mechanisms

Producing a sequence of phonemes, or carrying out a series of actions, involves activating a series of elements in the correct order. And, within a model of such behaviour, it is required that the representations of these elements are also activated in the correct order. In order to describe the necessary functions of such a mechanism, I shall borrow several definitions from Dell, Burger, and Svec (1997). Thus, in the following description, *present* will refer to the representational unit that should currently be active; *past* will refer to the unit immediately preceding the present; and *future* to the unit immediately following the present.

The essential functions of a serial-order mechanism are: 1). *Turn on the present unit.* The mechanism must have a way of identifying and activating the present unit. 2). *Turn off the past unit.* The mechanism must deactivate the past unit, and as most activation-based theories assume that activation is associated with a large positive number that does not decay immediately after deactivation, there needs to be a way of countering the past units. 3). *Priming the future*. The mechanism must prepare to activate the future unit. Dell, Burger, and Svec (1997) refer to this as the *throw-away* principle, meaning that it is easier to throw away what you have than to search for something that you do not have. Adopting this principle allows the turning off of the past unit to be relatively easy, as the network already has access to this unit, due to its activated state as the present unit. But, activating the future unit is a more difficult problem. Having all units at least partially activated according to a *plan representation*, surmounts this problem. This is a set of units that can influence the activation of the other units in the sequence. The use of such a plan means that the future unit is already partially activated, and so does not need to be searched for when it has to actually be produced.

The activation of future elements is important in language, as the appropriate form for the present may depend upon the future. For example, the use of *a* or *an* depends on the initial phoneme of the following noun, e.g. *an <u>o</u>pen space*. There are also many examples of the future form depending upon the past, as in English subject-verb number agreement (e.g. *The <u>cat is</u> sitting on the mat,* or *The <u>cats are</u> sitting on the mat*). This demonstrates that there is a need to co-represent the past, present, and future, simultaneously within a serial-order mechanism.

Models of Serial Order in Language

Models of language fall into three main categories. Those that use a form of *associative chaining* to build a representation of a sequence of elements, and recent models of this sort incorporate a parallel-distributed-processing (PDP) recurrent network such as that of Elman (1990); *frame-based* models that impose a pattern into which the elements must be correctly fitted (e.g. Dell, 1986); and control signal models such as the *competitive queuing* model of Houghton (1990), and the subsequent models of Hartley and Houghton (1996), and Vousden, Brown, and Harley (2000). I will discuss each of these types of model in turn, with particular attention to any difficulties they may present for the representation of phonotactic information.

Associative Chaining Models

These models are the simplest of the serial order models. The sequential control of the elements to be produced is achieved through a left to right chain of association. Retrieval of the correct order is achieved by following the link from one element to the next. Figure 1 gives an example of this for the production of the word *dogma*, in which following the chain from one element to the next produces the word.

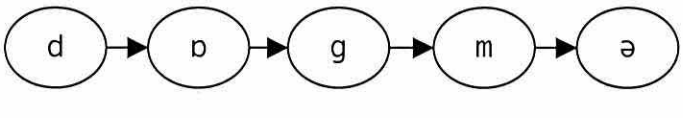

Figure 1: An associative chaining representation of the serial order for the word *dogma* (Adapted from Vousden, Brown & Harley, 2000).

The use of chaining, in the context of phonological encoding, was first suggested by Wickelgren (1969), and this has led to more elaborate models based upon connectionist modelling techniques, including Dell, Juliano, and Govindjee (1993), Elman (1990), and Jordan (1986). I will look at the model by Elman in more detail, and then discuss the problems with associative chaining models.

Elman's model (1990) achieved an internal representation of time, through the use of a simple recurrent network. Figure 2 gives a simplified architecture for such a network. Within a feed-forward network, a learning algorithm is employed that allows the hidden units to recode the input patterns in such a way that the appropriate output patterns are produced. In the model proposed by Elman, the context units "remember" the previous internal state of the network (the activation levels of hidden units) on a one-to-one basis. Thus, the following output will be based upon the present input activations plus the "memory" of the previous internal state. In this way, the changing contextual representation allows the network to map between a plan representation and a sequence of outputs, based upon knowledge of the network's previous state.

Such PDP explanations do solve some of the problems of simple associative chaining accounts. Recurrent networks can represent strings of phonemes in which the next phoneme can be predicted by one other than that which was most recently produced (e.g. *The fish is swimming* compared to *Lots of fish are swimming).* In these models the cue for the next element is not the immediately preceding element, but a combination of several of the previous elements (e.g. Dell et al., 1993; Elman, 1990; Jordan, 1996). These connectionist accounts may also model coarticulation effects (e.g. Jordan, 1996). But, such models have difficulty accounting for speech errors in which the source of the error is in a future element, such as would be the case for anticipations and transpositions. For example, chaining accounts would have difficulty explaining such errors as *aminal* being produced instead of *animal.* In a simple associative chaining model, a broken link would make it impossible for the sequence to continue, unless this led to yet further errors. Such mechanisms contain no mechanism that allows nonadjacent elements to become sufficiently active to be produced (e.g. Dell et al., 1993).

Chaining models would also seem poorly equipped to explain the *tip-of-the-tongue* phenomenon in which a speaker may be able to report some aspects of the sounds contained in a word, but be unable to complete the utterance (Brown & McNeil, 1966; Jones & Langford, 1987). If a speaker can report that a word ends in /ɪŋ/, but is unable to report the beginning of the word, then a chaining explanation would seem unlikely.

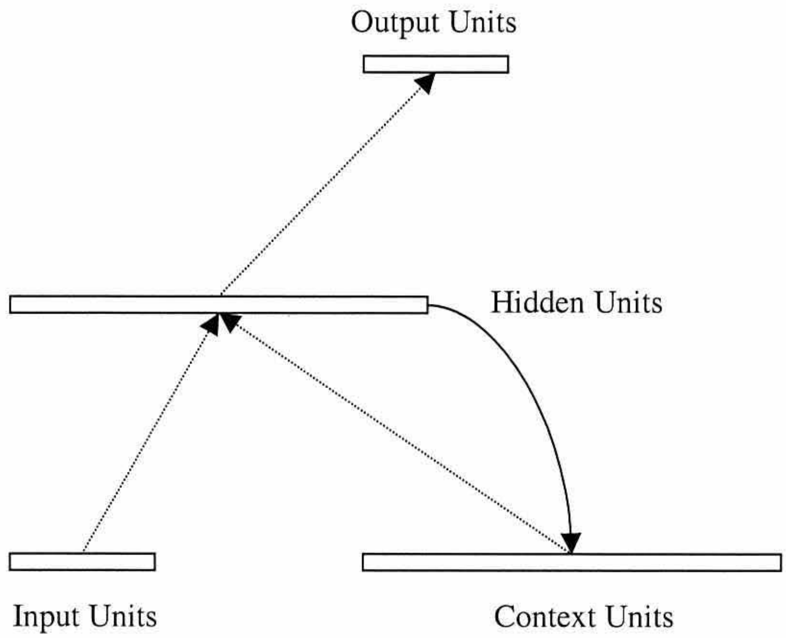

Figure 2: A simple recurrent network in which activations are copied from hidden layer to context layer on a one-for-one basis. Dotted lines represent trainable connections (Adapted from Elman, 1990).

However, it is evidence that the syllabic structure tends to be retained, even when phonological content is jumbled (Treiman & Danis, 1988), which is the major problem for associative chaining. Treiman and Danis asked whether syllables are coded in terms of onset and rhyme units, and whether the rhyme was then coded in terms of subunits. Although the ideas explored in the Treiman and Danis paper are of interest in terms of phonological encoding per se, they are not of particular relevance to this thesis, and so will not be discussed in detail. However, in the Treiman and Danis study, participants were asked to repeat strings of CVC syllables, with the responses being recorded and checked for errors. Within the data corpus only 42% of responses had the correct phonological content, but 96% of errors retained the original syllabic structure. Such evidence supports frame-based models, in which an abstract syllabic structure is imposed, to a greater or lesser degree, upon the order that the phonemes may legally be produced.

Frame-Based Models

The main difference between frame-based and associative chaining models is that frame-based models produce serial order by associating phonemes with positional markers or a frame of some kind, rather than it being produced by relationships between the phonemes themselves (e.g. Dell, 1986, 1989; Hartley & Houghton, 1996; Lapointe & Dell, 1989; Roeloffs, 1996, 1997; Shattuck-Hufnagel, 1979).

Although there are individual differences between these frame-based models, I will concentrate on two in particular that differ in two important respects. Dell's model (1986) adopts a *hierarchy* with *forward lateral inhibition*, and has a tightly constrained representation of a syllabic frame. On the other hand, the model by Hartley and Houghton (1996) uses *competitive queuing (CQ)* (Houghton, 1990), and uses two separate routes in encoding the phonetic information and the syllabic frame. I will first look at the details of the model developed by Dell, discussing any shortcomings that the model may have, before moving on to look at the model developed by Hartley and Houghton.

The Frame-Based Model of Dell (1986)

Before examining the model developed by Dell (1986), it is necessary to look at a hypothesis that is central to the model; the *two-step lexical access hypothesis.*

- *The Two-Step Lexical Access Hypothesis*

This hypothesis offers an explanation for the process involved in the naming of pictures. And, although this thesis is not concerned with this task, it is worth describing this hypothesis in order to clarify the architecture of the model proposed by Dell.

Picture naming tasks involve the translation of a visual stimulus into a conceptual representation, the retrieval of the name associated with the image, and then articulation of the name (e.g. Glaser & Glaser, 1989; La Heij, 1988, Potter & Faulconer, 1975; Theios & Amrhein, 1989). In the model proposed by Dell (1986), it is assumed that the conceptual representation is a set of features associated with the object (e.g. a cat would be represented by features such as PET and FURRY), and that this forms the input to lexical access.

Accounts of lexical access show that mapping a conceptual representation on to a phonological form is an extremely complex process. It involves the manipulation of many types of information: conceptual, pragmatic, syntactic, and phonological (Levelt, 1989). It also involves mapping between two unrelated spaces, as with the exception of morphologically related words, and the rare case of phonetic symbolism, there is no tendency for words that are similar in form to also be similar in meaning.

It is this difficulty that has led many theories to adopt a two-stage mapping process (e.g. Butterworth, 1989; Dell, 1986; Fay & Cutler, 1977; Fromkin, 1971; Garrett, 1975; Kempen & Huijbers, 1983; Levelt, 1989; Levelt et al., 1991a, Roeloffs, 1992). The first step of this two-stage process involves the mapping of the conceptual representation on to a *lemma*. A lemma is a nonphonological representation of a word that is associated with semantic and grammatical information. The second step involves mapping the lemma on to the phonological form of the word. This step is known as *phonological access*.

Both empirical evidence and functional considerations motivate the use of the two-step process. From a functional perspective, the fact that there is an arbitrary relationship between form and meaning makes the intermediate step a necessity if the mapping is carried out by spreading activation. According to Dell et al. (1997), if the phonological form and the meaning of a word were both patterns of activation across a series of nodes, then a one step mapping, utilising direct connections, would not be possible.

Non-decomposed theories of picture naming have been proposed (e.g. Collins & Loftus, 1975; Fodor, 1976; Fodor, Fodor, & Garrett, 1975; Fodor, Garrett, Walker, & Parkes, 1980; Garrett, 1982; Kintsch, 1974), and these suggest that a word is retrieved directly from an abstract representation. For example, an abstract representation CAT is used to retrieve the word *cat*, and the properties PET and FURRY are represented separately (Roelofs, 1992).

Empirical evidence for the two-step process comes from various sources: speech errors in aphasic and nonaphasic speakers (Garrett, 1975, 1980, 1984), analysis of the time-course of lexicalisation (e.g. Schriefers, Meyer, & Levelt, 1990), experimental studies of the production of multiword utterances (Ferreira, 1993; Kempen & Huijbers, 1983; Meyer, 1994; Schriefers, 1992), and tip-of-the-tongue (TOT) phenomenon (e.g., Meyer & Bock, 1992).

Speech errors can be divided into two categories: *lexical errors*, which involve whole words; and *sublexical errors*, which involve the sounds that make up words. Lexical errors can be associated with lemma access, and sublexical errors with phonological access. The two types of error may occur simultaneously. For example, "Unicorn" may be spoken as "House". This involves a lexical error, *Unicorn* to *Horse*, followed by a sublexical error, *Horse* to *House* (Martin, Dell, Saffran, & Schwartz, 1994). The important aspect of lexical errors is that whole words are most often replaced by whole words from the same syntactic category (Fay & Cutler, 1977; Garrett, 1975, 1980; MacKay, 1982). Thus, for example, nouns replace nouns, and verbs replace verbs, but these replacements are not necessarily similar in sound (Garrett, 1975). This syntactic similarity, without phonological influence, gives support for the existence of a lemma-like representation.

Several experiments have attempted to examine the time-course of the picture naming process by determining when semantic and phonological information is active. In these experiments, participants are asked to name a picture but, while they are doing this, they also see a word that may be semantically or phonetically related to the picture name. In some experiments, participants were told to ignore the word, and the interference with picture naming was assessed (Schriefers, Meyer, & Levelt, 1990). In other experiments, participants had to respond to the word by either naming it (Peterson & Savoy, 1998), or by making a lexical decision about it (Levelt et al., 1991), and the response time to the word was taken as the dependent measure. All of these studies manipulated the time of picture onset relative to the presentation of the word. The results of these studies showed that semantically related words had an effect early in the naming task, and that phonologically related words had an effect later on. If it is assumed that semantic errors are occurring during lemma access (see Schriefers et al., 1990), and the phonological errors during phonological access, then these results provide evidence for the two-step hypothesis.

The TOT phenomenon also supports the two-step hypothesis as, when a participant is in a TOT state, they know that a word exists (lemma access has been successful), but they cannot complete phonological access. This is consistent with a two-step model. It is interesting to note that in languages with grammatical gender, the participants often know the gender of the TOT word (Badecker, Miozzo, & Zanuttini, 1995; Vigliocco, Antonini, & Garrett, 1997). This lends support to the concept of grammatical information being associated with lemmas.

I will now look in detail at the model developed by Dell (1986), and discuss its viability as a model of phonological encoding.

The model developed by Dell (1986) combines the two-step hypothesis with an interactive retrieval mechanism. Lexical knowledge is embedded in a three-layer network: a semantic layer, a word or lemma layer, and a phoneme layer (See Figure 3 below).

The units in the semantic layer represent the concept of the object, with each object being represented by 10 semantic units. These 10 units connect to the object's word node through bi-directional excitatory connections. The word node then connects to its phoneme nodes through bi-directional excitatory connections. There are no inhibitory connections in the network. The strength of the connections is assumed to be the product of learning and recent experience.

As the connections run in top-down and bottom-up directions, the network is a form of interactive activation model. This makes it similar to a number of interactive models of lexical access (Berg, 1988; Dell, 1985, 1986; Harley, 1984; MacKay, 1987; Stemberger, 1985, 1990, 1991), but it differs from the two-step models of lexical access (e.g., Butterworth, 1989; Fay, & Cutler, 1977; Garrett, 1980; Levelt et al., 1991). The combination of interaction with two distinct steps allows the model to account for the data that motivate the two-step model, as well as for a variety of other error phenomena that suggest that there is interaction amongst levels.

To look at this model in more detail, I will start with the process of lemma access. If, for example, a picture of a cat is presented, then visual processes that are outside of the model identify the picture. The 10 concept units in the semantic layer are then given a jolt of activation. The activation level was arbitrarily set to 100, which when divided between the ten concept units, gave each unit a jolt of 10. This activation then spread throughout the network for a number of time steps, according to a linear activation function. Each node's activation is also perturbed by normally distributed noise during each time step. This noise is the sum of two components, *intrinsic noise* and *activation noise.* This results in the noise increasing as the activation level of the node increases.

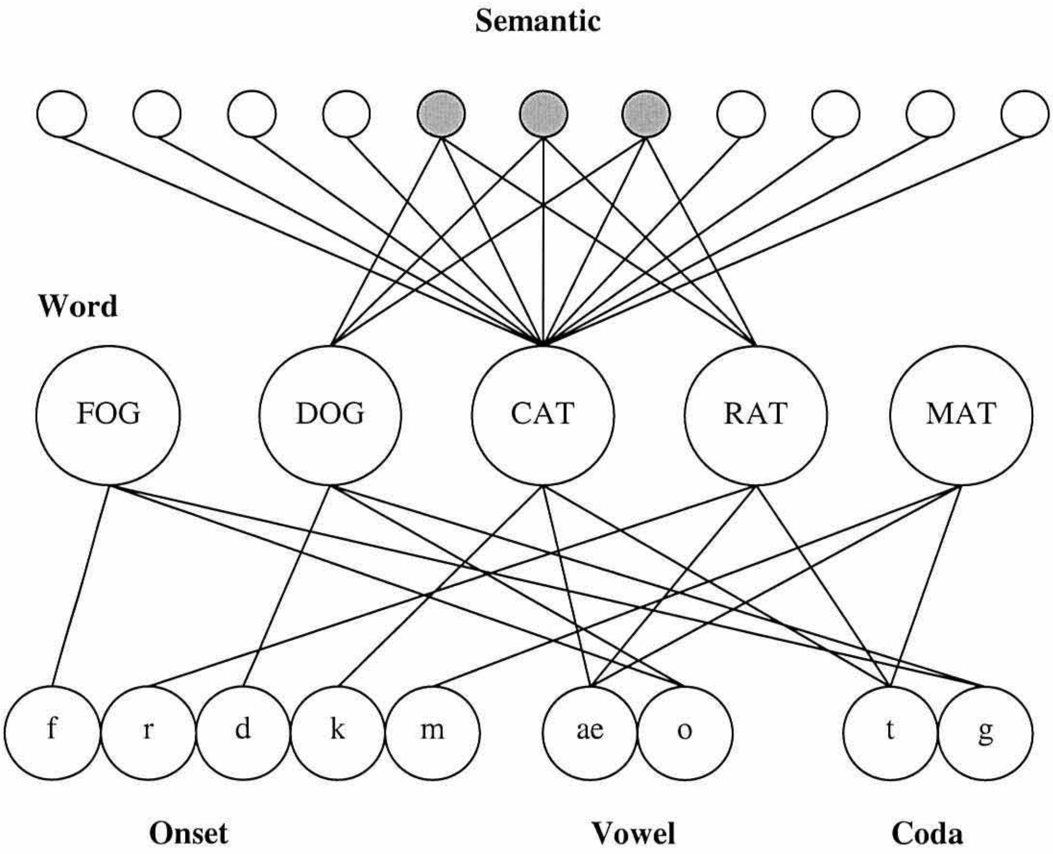

Figure 3: Illustration of the lexical network for the interactive two-step model of naming developed by Dell (1986). Connections are excitatory and bi-directional. The common semantic features of *cat, dog,* and *rat* are shaded in grey (Adapted from Dell, Schwartz, Martin, Saffran, & Gagnon, 1997).

The same equation applies to all nodes in all layers of the network, and it is applied at every time step in both lemma and phonological access. This means that the semantic nodes for any given word will be subject to decay (which is a function within the activation function), input from neighbours, and noise. This also allows the phonological nodes to receive some activation during lemma access. As the connections within the network are bi-directional, this activation creates positive feedback, and so the phonological units provide feedback for the word units, which then provide input for the semantic units. The existence of feedback within the network means that during lemma access the phonological and semantic neighbours of a given word will also become activated. This means that the most active word nodes at a given time step, are the target and the phonological and semantic neighbours.

The lemma access procedure concludes when the most active word node, of the correct syntactic category, is selected. During the production of a sentence, a selection is made based upon the linking of a word to a slot in a syntactic frame. Such frame and slot approaches to grammatical encoding during sentence production have empirical support (Bock & Loebell, 1990; Garrett, 1975; Levelt, 1989). Dell (1986) assumes that the task of object picture naming, as modelled in the network, uses a degenerate frame that merely requires the selection of a noun. Thus, the most active noun is selected at the end of lemma access.

Following this, *Phonological Access* takes place. This begins when the most active word node is given a large jolt of activation; 100 units worth, as in the original activation of the semantic level. This large jolt is important, as it introduces nonlinearity, and allows the word nodes to become a useful hidden layer. This allows the meaning-to-form mapping to be achieved. This enhancement of the "winner" is sometimes performed by lateral inhibition among the competitors (e.g., Feldman & Ballard, 1982; Grossberg, 1982; Harley, 1990), or by an absolute threshold that boosts the output of the node when the threshold is crossed (e.g., MacKay, 1987). In the Dell model, this process is tied to syntactics. The source of the jolt is the syntactic slot to which the word is linked (e.g., Berg, 1988; Dell & O'Seagdha, 1991; Eikmeyer & Schade, 1991; MacKay, 1982, 1987; Stemberger, 1985).

After the jolt, the activation continues to spread. And, as with lemma access, the activation spreads both upward and downward through the network so that nodes, other than that originally selected, can be activated. However, the goal during this process is to retrieve the appropriate phonemes. At the end of this process, the most highly activated phoneme nodes are selected and linked with the slots in the phonological frame. Most current theories of lexical access hypothesise that phonological access consists of retrieving phoneme-size units that are then slotted in to a frame. Evidence for phonological frames comes from speech errors (e.g. Shattuck-Hufnagel, 1979; Stemberger, 1990), and experimental studies showing that frame structures can be primed (Meijer, 1994; Romani, 1992; Sevald, Dell, & Cole, 1995).

The model, as implemented, is simplified in that it only has a frame for single-syllable CVC words. Each phoneme is labelled according to whether it is an onset consonant, a vowel, or a coda consonant. This limits the model's ability to predict more complex error patterns. Selection consists of picking the most highly activated phoneme from each category. This "categorical" selection is suggested by patterns of sound substitutions in phonological speech errors (e.g. MacKay, 1970, 1972; Shattuck-Hufnagel, 1979). Models have been developed which do handle more detailed error patterns (Berg & Schade, 1992; Dell, 1986, 1988; Eikmeyer & Schade, 1991; Hartley & Houghton, 1996; MacKay 1987; Stemberger, 1990, 1991).

The model allows for the basic kinds of errors that occur during lexical access. These are divided into five categories: semantically related word errors; formally related word errors; mixed semantic and formal errors; unrelated word errors; and nonwords or neologisms.

A semantic error occurs because the words share semantic nodes. For example, the nodes for *cat* and *dog* share semantic nodes, and dog could be activated in error if it becomes the most active node. This can happen due to noise in the system.

A form-related error, such as *mat* for *cat*, can occur either as an error of lemma access or of phonological access. During lemma access, nodes for words that share phonemes with the target become activated by feedback from the target phonemes to the word layer. Therefore, both *mat* and *sat* would gain some activation. The wrong word could be selected through noise affecting the activation level. Form-related errors that occur during lemma access would, therefore, be expected to be nouns in a picture-naming task, and that these errors would follow the syntactic class constraint. Form-related errors can arise during phonological access. The correct word may have been selected but, due to interference from other activated words or from noise, one or more phonemes may be replaced. The difference here is that a form-related error that does not follow the syntactic class constraint may occur.

Mixed semantic-formal errors, such as *rat* for *cat*, are particularly important because they may reveal the joint effects of semantic and phonological similarity, and hence, speak to the model's assumption that semantic and phonological information are active at the same time. In this model, the *rat* word node obtains activation directly from shared semantics, and from feedback from shared phonemes. The combination of top-down and bottom-up information gives *rat* a much better chance of occurring as an error than a purely semantic or purely formal neighbour. In the model, mixed errors are expected to be more likely than semantic errors that happen to be phonologically related, or formal errors that happen to be semantically related. That is, semantic and formal influences should not be independent. Analyses of normal speech error collections (Dell & Reich, 1981; del Viso, Igoa & Garcia-Albert, 1991; Harley, 1984), and experimental studies of normal speakers' picture naming (Bredart & Valentine, 1992; Martin, Weisberg, & Saffran, 1989), have consistently shown that there is a true mixed-error effect (i.e. semantic and formal influences are not independent). This model attributes this nonindependence to interaction among network layers during lemma access. Alternative accounts, which do not hypothesise phonological activation during lemma access, attribute mixed-error effects to the action of late editorial processes in production (e.g. Levelt et al., 1991).

Unrelated word errors are word substitutions that are neither semantically, nor formally related to the target. Unrelated errors occur at lemma access, and are attributable to noise and any small amounts of activation obtained from distant relations to the target. In principle, unrelated word errors, like any word error, can also occur during phonological access, because of the remote possibility that a correct word selected at lemma access is phonologically coded as an unrelated word (i.e. the correct word is selected but, due to competition within the network, phonemes are selected that produce an unrelated word). Alternatively, an unrelated error may reflect trouble at both lemma and phonological access. In the case of *log* for *cat*, a semantic error at lemma access, *dog*, could be phonologically encoded as *log*, resulting in an unrelated word error.

The final error category allowed in the model, is for neologisms or nonwords. A nonword, such as *lat* or *cag*, indicates a problem at phonological access. Noise combines with interference from other activated words, resulting in the replacement of one or more target phonemes. A nonword that resembles the target, or target-related neologism, would likely reflect correct lemma selection that was followed by incorrect phonological access. A nonword that does not resemble the target, an abstruse neologism, could arise from either a severe disruption of phonological access, or difficulties at both lemma and phonological access.

At this point, it is also worth considering the argument given by Caramazza (1997), who takes issue with the two-step lexical hypothesis, and questions whether the lemma level or representation is necessary. The evidence used by Caramazza (1997) is taken from the pattern of dissociations of lexical errors in speaking and writing, the main points of which I will summarise here.

Caramazza notes that there are brain-damaged patients who make semantic errors in only one modality of output, and that the deficit in these patients is located in the language production system, at a stage beyond the lexical-semantic level. The example of HW (Caramazza & Hillis, 1990) is given, amongst others, in support of this claim. HW, when asked to read aloud and then define a word, would often make semantic errors in oral reading, but would then give the correct definition of the word. The fact that semantic errors were restricted to one modality of output implies that, since the patient had selected the correct lexemes in one modality, their associated lemmas must have been correctly accessed or they could not have produced the correct lexemes. This rules out a post-lexical deficit as the explanation for the differing performance in oral and written naming and, therefore, the locus of the damage in such patients must be at a point between the correctly selected lemma-level representation and the modality-specific lexeme representations (i.e. a phonological lexeme or a orthographic lexeme).

Having confined the locus of the damage to a point between the lemma-level representation and the modality specific lexeme representation, Caramazza (1997) addresses the question of how the inaccessibility of a modality-specific lexeme can result in a *semantic* error, if the correct lemma had already been selected. It would be expected that the failure of a lemma node to activate a modality-specific lexeme node would result in the absence of a response. However, Caramazza offers a possible explanation. He suggests that the failure to select a modality-specific lexeme could result in another concept node being selected, because of the spreading activation within the system, and this would therefore lead to the selection of a new lemma etc.

However, Caramazza (1997) suggests that this explanation is not feasible when data from other patients is considered. For example, he quotes the performance of WMA (Miceli, Benvegnu, Capasso, & Caramazza, 1997) and PW (Rapp, Benzing, & Caramazza, 1997). Both of these patients made different oral and written semantic errors when naming the same object. In this case, Caramazza suggests the following explanation. If a semantic error is made in one modality, then the process outlined above can account for the mistake, but, if a semantic error is made in both modalities, then it would have to be argued that in this modality, too, the target lexeme could not reach threshold, and would therefore lead to another cycle of lexical selection for production in that modality. The question asked by Caramazza is, why is a different lexical concept selected the second time? To illustrate the problem he uses data from the patient PW. PW was shown a picture of tweezers, and was then asked to name them, then write down the name, and finally, name them again. PW's responses were "pliers", "needle", and "pliers" respectively. This illustrates that the selection of the concept cannot be random, as the selection of alternative concepts is consistent from one trial to another, it is only across modalities that different responses were produced for the same picture. Caramazza argues that in the above case the phonological-lexeme representing "tweezers" could not reach threshold, and so the next concept node was selected, "pliers". The phonological-lexeme for "pliers" could reach threshold, and so the process was completed. It would then have to be argued that when trying to write the word "tweezers", PW selected the correct concept node, but then encountered difficulty, because the orthographic-lexeme for "tweezers" could not reach threshold. It would therefore be assumed that the same alternative concept node as the previous trial was chosen in its place (pliers). But, as PW produced a different semantic error in writing, it must be assumed that the orthographic-lexeme for "pliers" could not reach threshold and, therefore, a completely new cycle of lexical selection had to be performed. This new cycle led to the selection of the concept node "needle", which could be completed.

Caramazza admits that these two cases do not constitute a "knock-down" argument against the two-stage hypothesis. But, he argues that the way in which the two-stage hypothesis has been forced to account for the results, provides an argument against the hypothesis. Caramazza points out that had the lemma node been removed completely from the explanations for the data, then they would still constitute a formally equivalent argument. Therefore, argues Caramazza, the lemma level is superfluous.

- *Further Problems with the Dell (1986) Model*

There are two main problems with the model proposed by Dell (1986). Firstly, because each naming attempt is assumed to be separate, perseverative errors are not allowed for. Secondly, and more importantly for this thesis, because the model makes assumptions about the association of phonemes to slot positions, it does not allow for speech errors in which phonemes do not maintain their syllabic position. For example, the model assumes that if the phoneme /f/ was intended to be produced as an onset, but was actually produced in another syllable, then it could only be moved to another onset position. The model would not allow for the onset representation of /f/ to be moved to a coda position, due to the "labelling" of each phoneme as an onset, vowel, or coda. Dell claims that such strong restrictions on the movement of phonemes are correct for normal speech errors (i.e. phonemes produced in error follow the syllable position effect). However, if the syllable position effect could be shown to be weaker than previously reported (e.g. Ellis, 1980), then it would pose insurmountable problems for the model.

Control Signal Models

In order to offer a possible solution to the difficulty caused by the rigid association of phonemes to slots, I will now consider a control-signal model developed by Hartley and Houghton (1996) that uses a dual route to encode the syllabic frame and phonetic information, and so potentially allow for such errors. I will also consider a model by Vousden, Brown, and Harley (2000) that uses a different implementation of the control-signal. Unlike the model by Dell (1986), these models do not attempt to include an explanation of the retrieval of words from the mental lexicon. Therefore, they do not include any semantic representation, and so make lemma access unnecessary. Both these models assume that the phoneme sequence has already been constructed, and so begin the simulation at this point.

The Dual-Route Model of Hartley and Houghton (1996)

The model proposed by Hartley and Houghton (1996) learns and recalls unfamiliar words, which follow the linguistic rules of English, by using an approximation of short-term memory, i.e. it can repeat a sequence of nonwords that has been presented once. The repetition of strings of nonwords is a task that normal speakers can accomplish from an early age (Gathercole & Baddeley, 1989; Gathercole & Adams, 1993). The ability to repeat strings of nonwords seems to be an easy task for normal speakers. However, the mechanism behind this task is a complex system that controls the representation and recall of verbal stimuli that is serially ordered. This task is difficult, as spoken words are distributed across time, and so at no point is all the information concurrently available. Thus, the mechanism must be capable of tracking the input in real-time, and then, on the basis of a single trial, repeating the input in the correct order. The model is the result of several empirical observations and theories, and I shall discuss each of these in turn.

- Empirical Background

The repetition of novel nonwords is a frequently used method for testing short-term memory. However, it has been largely confined to developmental and clinical studies. Within developmental studies it has been shown that, by the age of 8 years, a child can correctly reproduce 90% of two-syllable nonwords, and that the more syllables there are, the lower the percentage of correct recall. Within the errors that are made, the majority are single phoneme substitutions (e.g. Gathercole, Willis, Baddeley, & Emslie, 1994).

There is evidence from clinical studies to suggest that errors are constrained by linguistic factors. For example, Caramazza, Miceli, and Villa (1986) reported on a patient with global aphasia, and Bisiacchi, Cipolotti, and Denes (1989) reported a patient with impairment to nonword reading and a severe short-term memory deficit. The phonemic errors of both these patients were more likely to be the substitution of a phoneme by another phoneme that shared the same manner of articulation.

A similar task, which is sufficiently difficult to test normal speakers, is the serial recall of lists of nonwords. Hulme, Maughn, and Brown (1991) compared memory span for words and nonwords, and found that span varied linearly with speech rate for both words and nonwords. The span for words was 5.0 and that for nonwords was 3.5. They suggested that both word and nonword recall are supported by a short-term memory store that uses an articulatory control process, and so leads to the consistent effect of speech rate upon span. It was suggested that the difference in span was, therefore, due to the availability of long-term phonological representations for the words.

The errors found when conducting such tests would seem to be qualitatively different from those found when using words. Ellis (1980) showed that the majority of errors when using unfamiliar stimuli involve phonemes from different target syllables recombining to form new syllables. These recombination errors were shown to be highly constrained, with phonemes maintaining their syllabic position, despite the error, on 84 – 96% of occasions.

Treiman and Danis (1988) made a detailed analysis of the error types in nonword recall. This showed a tendency for syllabic structure to be maintained; even when phonological content was confused (96% of responses maintained the syllabic structure, but only 42% contained the correct order of phonemes). This shows that the most common type of errors are phonemic substitutions, with insertions and deletions being relatively rare. Most of these errors were recombinations of phonemes from different items in the same list, but there were some that did not originate in the target list. A similar pattern was found for various syllable structures (CVC and VCC), and so it was clear that consonants and vowels rarely substitute each other.

Such recombination errors are comparatively rare in the recall of familiar words. But, the misordering of entire items can be found in both word and nonword recall. Hartley and Houghton (1996) suggest that the recombination errors can be explained as a failure of a subsystem dedicated to the retention of phonological forms. When the stimuli are familiar words this system is largely redundant, due to the ability to draw upon long-term phonological knowledge (Hulme et al., 1991). When the words are novel, there is no long-term representation to call upon, and so the speaker is reliant upon short-term memory, with the result that the error rate increases. Patterson, Graham, and Hodges

(1994) give data on three patients (PP, FM, and JL) who suffered from semantic dementia, a condition that results in a gradual degradation in semantic knowledge. The patients were shown lists of *known* words (those for which they still had semantic knowledge) and *unknown* words (those whose meaning they could no longer retrieve), and asked to pronounce them. The resulting pattern of speech errors was very similar to that reported by Treiman and Danis (1988), in that there were large numbers of phonological errors on the words that the patients could no longer recall. Patterson et al. propose an explanation for these data, suggesting that the word's long-term phonological representation is normally supported by the semantic representation. And, as this has been disrupted in these patients, recall must depend upon short-term memory, as would seem to be the case for the recall of nonwords with normal speakers.

Caramazza et al. (1986) proposed an alternative to the above suggestion, in that long- and short-term memory stores have separate phonological output systems. They based this proposal upon acquired (Caramazza et al., IGR, 1986; Bisiacchi et al., RR, 1989) and developmental cases (Campbell & Butterworth, RE, 1985; Hulme & Snowling, JM, 1992), in which the patients demonstrated particular problems with nonwords (compared to familiar words), in a range of tasks assumed to involve short-term memory. In these patients, performance on nonwords repetition (and other tasks requiring the processing of nonwords) shows a greater degree of disruption than for similar tasks involving words. However, measures of span for familiar words, still shows some degree of impairment.

Hartley and Houghton (1996) suggest that the existence of specialised systems for the output of familiar forms is not intuitively appealing, questioning why there should be different processes for the production of novel and familiar forms. They point out that this idea is also at odds with data showing that phonological errors in nonword recall are qualitatively similar to those seen in spontaneous phonological encoding (Ellis, 1980). Also, such a theory could not explain why difficulties in processing nonwords are often associated with a short-term memory deficit. Hartley and Houghton suggest that it would be more parsimonious to explain this data in terms of a general short-term memory deficit, and the processing demands of the two tasks (i.e. nonword recall requires retention of phonological information, whereas word recall draws upon long-term phonological storage).

-Theoretical Background

When developing the model, Hartley and Houghton (1996) wanted it to fit in with existing models of related processes, so that it would provide insight into the interaction of these subsystems, as well as how they operate in isolation. From this point of view they embraced models from two areas: short-term memory, and spontaneous phonological encoding. I will now look at each of these areas in turn.

1. Models of short-term memory. Baddeley and Hitch (1974) proposed the articulatory loop model of verbal memory, and this provides an account of many well-known phenomena (e.g. the reduction of span under articulatory suppression). But, until more recently, it lacked a computational specification, including the ability to represent serial order. Burgess and Hitch (1992, 1999) addressed this problem, proposing a connectionist model of the articulatory loop that used the *Competitive Queuing* mechanism (Houghton, 1990) to control serial order. In competitive queuing models, the items become associated during the learning phase with the states of an internally generated dynamic context, which provides a "distributed" representation of the serial position of the items to be recalled. Recovery of the context, during the recall process, leads to the parallel activation of all items that occurred in close proximity during learning. The sooner an item is to be output, the more strongly it is activated, and all activated responses compete for control of output at a "competitive filter". This filter picks out the most active response and then suppresses it, allowing generation of the next response (c.f., Estes, 1972; Rumelhart & Norman, 1982). In this way, competitive queuing models effectively separate the preparation of responses, which may happen in parallel, from the selection of the next response, which must necessarily happen serially. Thus, competitive queuing models are consistent with the observation by Lashley (1951) that responses must be active at some level before they are generated (Houghton & Hartley, 1996).

Hartley and Houghton (1996) took the model by Burgess and Hitch (1992) as the starting point for the representation of serial order in verbal short-term memory. This model has a number of properties that make it suitable. Firstly it has the property of being able to learn and recall an ordered list in a single trial. And secondly, it exhibits a number of traits, including a tendency to create ordering errors such as paired transpositions. But, the Burgess and Hitch model is limited in two ways: a). it can only use familiar words, and b). it is only concerned with the ordering of words, and not of the phonemes within the word.

2. *Models of spontaneous phonological encoding.* Within their model Hartley and Houghton (1996) wanted to account for the pattern of constraints that

dictate the type of phonological errors that can occur within nonwords. Within models of phonological encoding, this is usually achieved through separate representations of phonological structure and content. Among the models that display these characteristics, Hartley and Houghton single out that of Dell (1986,1988) that I discussed earlier. However, I will briefly describe this model again, so as to indicate the reasons that it was of interest to Hartley and Houghton. In Dell's model, priming effects are responsible for phonological order errors in spontaneous speech. Activation is given to utterances prior to their production (as in competitive queuing models). The activation spreads through the network to the phonological output layer, which may occasionally lead to phonological errors in the output. Of particular interest, is the fact that the model shows anticipatory and perseveratory substitutions of single phonemes, which are the most frequent types of spontaneous speech error.

In the Dell (1986) model, the order in which the phonemes are selected is constrained by a syllable schema. Phonemes are represented by different nodes that further constrain their syllable position. (e.g. /k/ would have separate nodes for onset and coda positions). During the production of a syllable, the schema selects the most active onset, peak, and coda nodes. In a later model (Dell, 1988) a different mechanism was employed that used a number of wordshape (CV-structure) representations. However, in both cases the models represent content and structure separately, with the two representations interacting at the production stage. Dell (1988) suggested that this interaction might be achieved through the activation of the phoneme category nodes in series, with each sending an increasing amount of activation to all of the phoneme nodes, until one of them, starting with the one with the highest activation, reaches some threshold.

This model offers the possibility that there is a competitive process at the output level that is based upon relative activation levels. This is the same as that suggested by the competitive queuing models, but, in order to account for the observed error patterns, competition is restricted to phonemes that can occupy a given position, as dictated by a syllabic template.

Hartley and Houghton (1996) suggest that it is natural to suppose that the underlying articulatory control processes should be the same for errors in the recall of nonwords, and for those in spontaneous speech, given that the errors are both constrained in similar ways. They suggest that the higher rate of errors in nonword recall, when compared to errors in spontaneous speech, can be explained by the representation being more "fragile". This is because the nonwords have been learned in one trial, whereas the words in spontaneous speech have a long-term representation that is more "robust". Hartley and Houghton, therefore, take Dell's approach as the most obvious candidate for integration with the Burgess and Hitch (19992) model. I will now discuss how this integration was achieved, and the model that resulted from this.

- Description of the Hartley and Houghton (1996) Model

This model has to incorporate both phonological structure and verbal short-term memory, and therefore has to simultaneously represent stimuli at two levels. The first of these is the syllable level, at which an input stream of phonemes is parsed into syllabic chunks. This represents the list of words that are to be learned. The syllable level is responsible for remembering each syllable's position in the list. The phoneme level performs a similar function for the order and position of the phonemes within each syllable. This is important, as not only must the syllable be recalled in the right position, it must also have the correct form.

The single-trial learning and recall within this model are assumed to operate along the lines of the Burgess and Hitch (1992) network model of the articulatory loop. This is augmented by a lower-level mechanism that represents syllabic structure and content. This lower-level mechanism records the input stream in real time, and parses it into syllables; thereby creating a representation that is capable of immediate repetition.

Architecture

The basic architecture is shown in Figure 4. The model contains a set of "uncommitted" nodes for the encoding of the syllables as they appear in the input. The phonological form of the syllables is encoded in two separate pathways; 1). The *content pathway* links the syllable units directly to the units representing its constituent phonemes, and 2). The *structural pathway* connects the syllable units to the phoneme units via a *syllable template*. In this template, excitatory connections, of a fixed strength, link nodes that represent "slots" in the template to the nodes representing phonemes that can fill those slots.

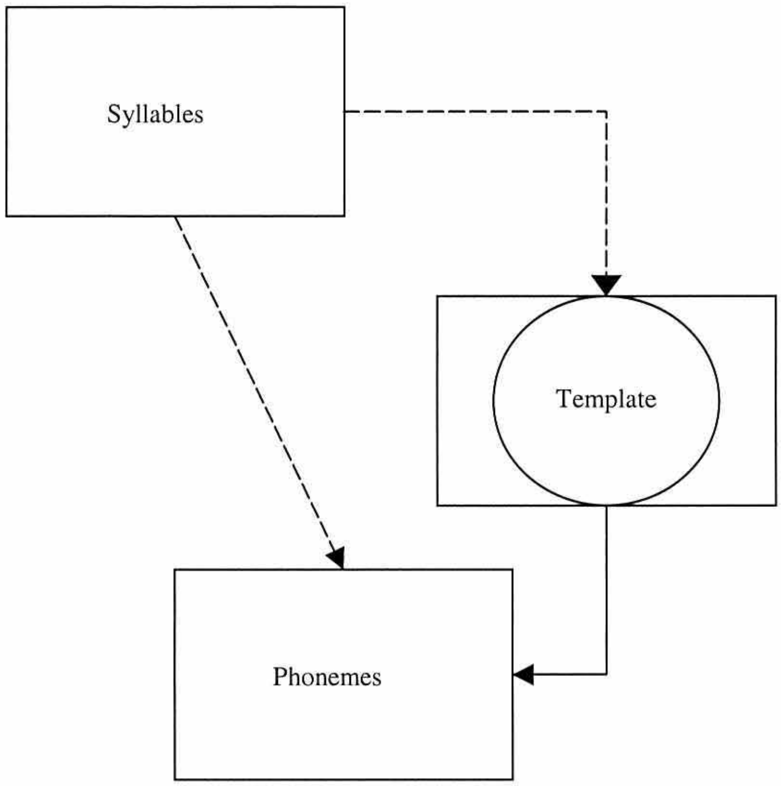

Figure 4: An outline of the model by Hartley and Houghton (1996).

All other connections in the network have variable weights, which are set during the learning stage. Hartley and Houghton (1996) follow Burgess and Hitch (1992) by making these connections, which are responsible for short-term memory, temporary. The strength of these connections decays over time to prevent saturation of the system, and therefore, the same weights can be used on other occasions. Figure 5 is a more detailed, though still simplified, diagram of the model.

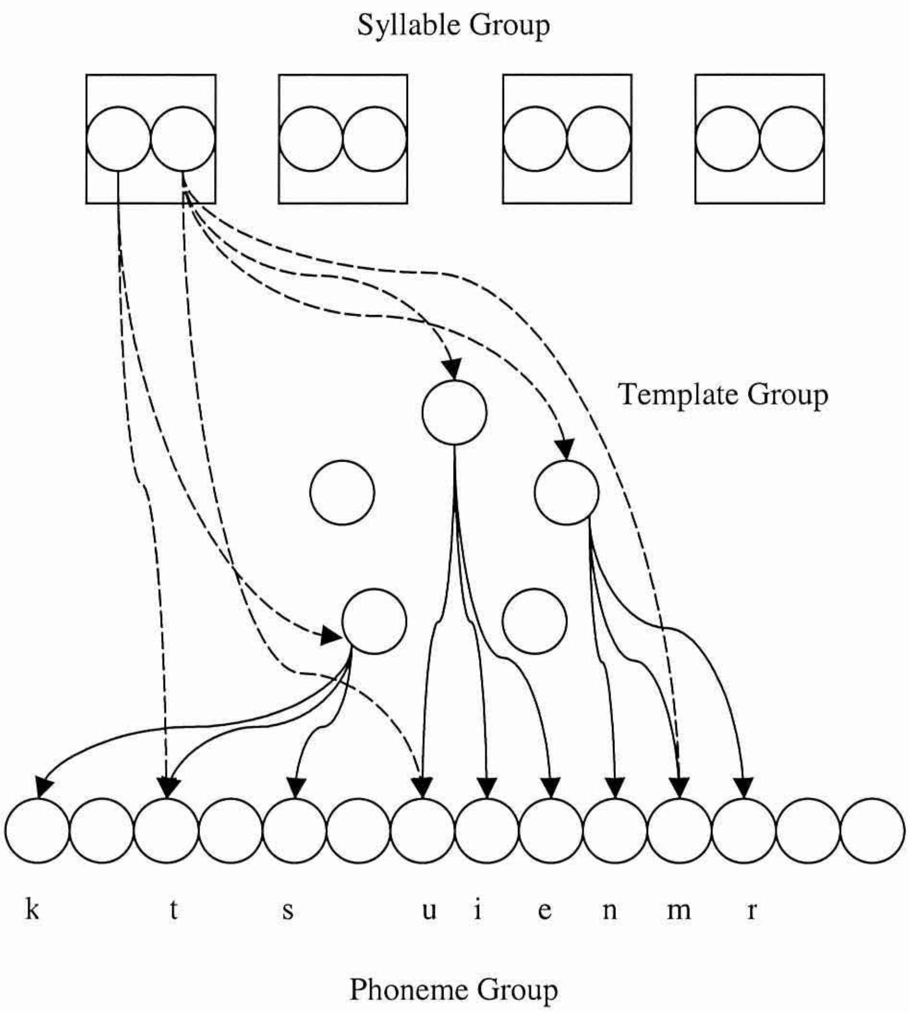

Figure 5: A more detailed diagram of the Hartley and Houghton model. The dashed lines represent temporary weights; the solid lines permanent connections (Adapted from Hartley and Houghton, 1996).

- *The Phoneme Group*

This group represents the phonemes perceived during learning, and the phonemes articulated during recall. There are a total of 47 nodes which represent 20 vowels, 24 consonants, and the consonant clusters /sp/, /st/, and /sk/, which are treated as single consonant phonemes (see *Syllable Template* p. 35).

- The Syllable Group

The syllable group consists of pairs of nodes, each representing one syllable. These two nodes represent the onset and the rhyme, respectively. At any time, only one node has all of the activation associated with the particular syllable. The division of syllables into onset and rhyme is supported by a variety of data (e.g. Fudge, 1969; Treiman, 1986; and Goswami & Bryant, 1990). In this model the division is also necessary as most of the consonants can appear in either the pre- or post-vocalic position. If this were not accounted for, then words such as "pat" and "tap" would be represented identically, as both contain the same phonemes and have the same CVC structure. In addition to this, it also prevents the system from generating a level of transpositions between pre- and post-vocalic positions, which is higher than has been reported in the previous data (Ellis, 1980).

- The Syllable Template

The purpose of the syllable template is to approximate the structure of a generalised "syllabic gesture". This is based upon the linguistic quality *sonority.* Sonority is related to the energy that a sound generates, and imposes restrictions upon the way in which phonemes can be ordered within the syllable. In a well-structured syllable, sonority increases with each phoneme until it reaches its peak with the vowel, and then decreases. Vowels have the highest sonority, nasals and liquids have intermediate sonority, and obstruents have the lowest level of sonority. Thus, the phonemic sequence /slInt/ conforms to the principle of sonority. Whereas, /lsInt/ and /slItn/ do not conform. The principle of sonority may be subject to language-specific constraints. For example, in English, the consonant cluster /tl/ cannot occur in the onset of a syllable (Fudge, 1969). But, in the model proposed by Hartley and Houghton (1996) it is the general constraints imposed by the principle of sonority that were implemented.

The template implemented within the model is based upon the work by Fudge (1969). Fudge proposed that English syllables conform to a structure, in which each part of the syllable (onset, peak, coda, and termination) has a limited number of phonemes that can legally occur in that position. The structure proposed by Fudge had six slots; with two slots each for onset and coda. And, each of these six slots had a list of possible phonemes that could fill the position. In the model implemented by Hartley and Houghton (1996), only the first five slots are represented, as the termination is usually reserved for inflectional endings and was omitted.

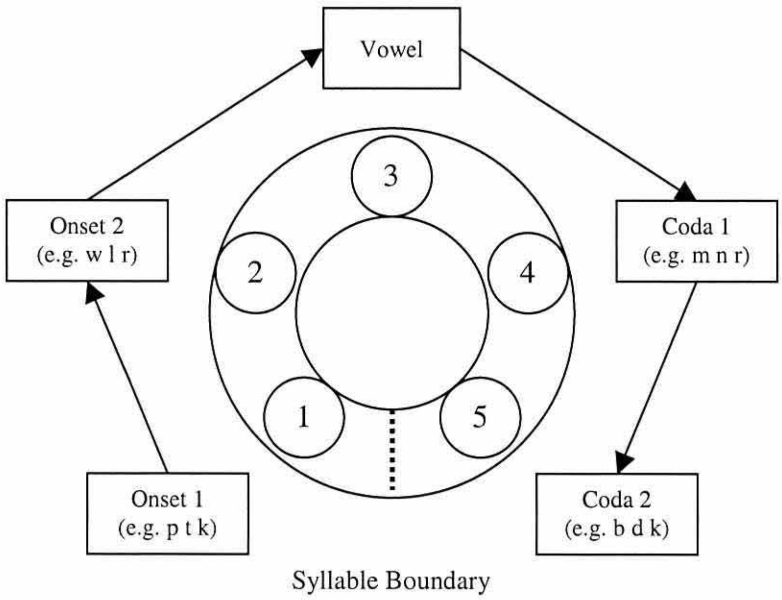

Figure 6: A simplified structure of the cyclical syllabic template in the used in the model (Adapted from Hartley and Houghton, 1996).

As can be seen in Figure 5, a node that is linked to a subset of phoneme nodes by permanent top-down connections represents each of these slots. Hartley and Houghton assume that these represent the speaker's long-term phonological knowledge. Figure 6 gives a simplified view of the relationship between the slots and the phonemes within the syllable template.

The relationship between the slots and phonemes was simplified for the model, so that each consonant was only associated with one pre- and one post-vocalic slot. In line with the proposal of Fudge (1969) the sonority increases with each successive slot, up to the vowel, and then decreases again. As there are very few legal syllables that violate the sonority constraint, most can be represented in a single cycle of the template. For example, using the limited number of phonemes shown in Figure 6, the legal nonsense syllable /plink/ can be represented by using slots 1-5 inclusive, and the syllable /pIn/ by using slots 1, 3, and 4. But, the illegal syllable /lpInr/ cannot be described by a single cycle. Thus, in general, syllables that violate the sonority principle cannot be described in a single cycle. However, there are also some legal English syllables that cannot be described in this way (e.g. *hits, sniffs*). These either involve the use of Fudge's termination position, or involve successive phonemes from the same slot. Also, because language-specific constraints are not implemented in the model, some syllables that break the phonotactic constraints of English can also be represented. From these exceptions it is clear that the syllable template is a simplified representation of the phonotactics of English. But, Hartley and Houghton (1996) argue that it is sufficient to capture the general constraints of the sonority principle, and thereby represent most legal English syllables. The template can also be used to parse a continuous stream of phonemes in to syllables, as the syllable boundary (see Figure 6) represents the start of a new syllable.

- Learning and Recall

The simulations performed with this model consist of two stages: learning, and recall. During the learning stage, the model is given a sequence of phonemes that, when parsed, is a sequence of nonsense syllables. The recall stage involves the model reproducing this sequence as accurately as possible. The sequence that was presented to the model, and the sequence that it subsequently produces, are then compared, and the errors counted and classified.

Learning

During learning the syllable template parses the stream of phonemes into syllables. As each phoneme is presented to the network, one node in each group is activated; the node representing the phoneme in the phoneme layer, the node representing the next required slot in the syllable template, and one node in the syllable group. The activation of these groups is now described in more detail.

Activation of Phoneme Nodes. At any one time-step, a single node is activated in the phoneme group, which represents the current input phoneme. All other nodes in the group have zero activation.

Activation of Syllable Nodes. Uncommitted nodes within the syllable group are activated in turn, and once an onset/rhyme pair has been used it is not used again. If the template node that is active represents slot 1 or 2, then the onset node carries all the activation. If slots 3, 4, or 5 are activated, then the rhyme node carries the activation. Each time the syllable template completes one full cycle the active syllable unit is changed. Using the weight change rule, which is described in the next section, allows each onset/rhyme pair to encode a single input syllable.

Activation of the Template Nodes. The template nodes are activated "bottom-up" by the phoneme nodes. Each phoneme node could activate any legal slot. However, as consonants can generally appear in either the pre- or post-vocalic position, it is assumed that the previous template slot that was activated will influence the slot activated by the present consonant. Therefore, the template slot activated will be the next legal slot that the phoneme can occupy. For example, consonants following the vowel will activate post-vocalic slots, if possible. During this process the weights on the connections are altered according to a Hebbian learning rule, so that currently active nodes have their temporary connections strengthened. The strengths of these temporary connections also decay towards zero.

The result of learning is that, for each syllable in the input, excitatory connections simultaneously become established between (i) syllable units and phoneme units (encoding phonemic content) and (ii) syllable units and template units (encoding syllabic structure).

Recall

The aim of the recall process is to recreate the serial pattern of the activation of the nodes during learning. The serial order of the syllable nodes represents the order of the syllables in the original stream, and the sequence of activation of the phoneme and template nodes represents the phonological form of the recalled syllables. Item order errors would occur at the syllable level, and errors of syllabic form and content would occur in the phoneme or template level. No learning takes place during the recall process and, therefore, the temporary weights decay exponentially. The various groups of nodes are activated as follows during recall.

Activation of Syllable Nodes. The activation of the nodes during syllable recall is assumed to be controlled by the competitive queuing process described by Burgess and Hitch (1992). This is important, as it causes a number of syllables to be activated in parallel, with the nodes' activity increasing in direct relationship to their proximity to the current target syllable during the learning stage. After a syllable is produced, its node is suppressed. As with learning, activation is associated exclusively with either the onset or rhyme in a syllable pair. The same member of the onset/rhyme pair will be active in all simultaneously activated onset/rhyme pairs. Hartley and Houghton (1996) claim that this retrieval process is consistent with the findings of Meyer (1991), who provided evidence that it is not syllables that are encoded sequentially, but that it is the onset and rhyme. Further to this, Meyer also demonstrated that it is the onset that is activated first, otherwise the onset would become the first part of the word and, as speakers sometimes anticipate segments that were meant to appear later in the utterance, the segments of several words must be active at the same time.

Activation of the Syllable Template. Input to the template comes from the onset/rhyme nodes. Connections will have been formed between the syllable nodes and the syllabic slots that were activated by the syllables during learning. In order to reproduce the order at recall, the syllabic template must be accessed so that the slots associated with the target syllable are activated in series. This is achieved by applying input cyclically to the template group. This input is a dynamic pattern that moves serially through the template, once for each syllable to be recalled. The level of this input is set so as to offset the decay that has occurred, thereby making the net input to the template the same regardless of the duration of the original stimulus material. The input should be sufficient to activate a template node only if that node is receiving input from the onset/rhyme pair representing the current syllable. When this condition is met, the input from the syllable units, via the learned weights, combined with the cyclic input, allows the syllabic structure to be recalled.

Activation of the Phoneme Nodes. It is the activation of the phoneme nodes in series, which constitutes the network's output. Any one phoneme node receives input from the structural and content pathways. Neither pathway will, typically, have developed weights during learning which are sufficient to activate any node beyond the threshold at which it becomes a competitor for output. However, a phoneme node will become active if the summation of the input from the structural and content pathways is sufficient. Due to the competitive cueing algorithm, which is used for recall, nodes in the syllable group will become active in parallel. In addition to this, firing of any of the template nodes will also provide input to all of the phoneme nodes to which it has a connection. This creates parallel input to the phoneme nodes. Most input would be received by phonemes that can fill a syllabic slot, especially those which appeared close to the current target syllable in the list.

Noise becomes an increasingly important factor in the selection of phonemes as the gap between learning and recall increases. This is due to the decline in the input from the syllable nodes, caused by decay in the weights, which leads to errors. Therefore, the performance of the model on recall is dependent upon the amount of decay that has occurred in the strengths of the temporary connections. This is related in turn, to the length of the list of phonemes that are to be recalled, and the rate at which they are articulated during learning and recall.

Sources of Error in the Model

Within the model, three main factors are responsible for constraining errors: long-term phonological knowledge (represented by the syllable template), the structure of the target syllable (the slots used in the template), and the structure of other syllables near to the target syllable in the stream of phonemes.

Long-term Phonological Knowledge. Within the syllable template each slot is associated with a different number of phonemes. The greater the number of phonemes associated with a slot, the greater the number of competitors, and so the greater the opportunity for error.

The Structure of the Target Syllable. A substitution or deletion may occur if a particular slot in the template is used. If the slot is not used, then an insertion may occur. It is also possible, although it would be rare, that a deletion and insertion may occur in tandem, and so give the impression that a phoneme from one slot has been replaced by one from another slot (e.g. /pIn/ → /lIn/).

The Structure of other Syllables in the Target Sequence. Syllables close to the target may affect the chance for error in two ways. Firstly, if the structure of adjacent syllables is the same as the target syllable, then they will provide reinforcement for the structure of the target syllable and hence reduce the chance for error. But, if the adjacent syllables have a different structure, then they will provide priming for the different structure that will result in interference and provide a possibility for error. These three sources of error can, therefore, account for substitutions, deletions, and insertions.

I will now look at the control-signal model of Vousden, Brown, and Harley (2000), and contrast this with the model of Hartley and Houghton (1996).

The OSCAR model of Vousden, Brown, and Harley (2000)

The OSCAR (Oscillator-based Associative Recall) model developed by Vousden, Brown, and Harley (2000) is similar to the model of Hartley and Houghton (1996) in that it is also a control-signal model. However, unlike the Hartley and Houghton model, it does not contain any representation of the syllabic structure. In OSCAR, a sequence of syllables, and their constituent phonemes, are linked to the states of a dynamic control system that specifies the order of the syllables and phonemes. The "replaying" of the states of the control-signal will, therefore, result in the recall of the syllables and phonemes in the order that they were produced. The motivation for the control-signal used in OSCAR comes from evidence for oscillatory systems in the timing of human motor actions (Treisman, Cook, Naish, & MacCrone, 1994; Treisman, Faulkner, & Naish, 1992), studies of human short-term memory (Brown, Preece, & Hulme, 2000; Burgess & Hitch, 1992, 1996, 1999; Henson & Burgess, 1998; Henson, Norris, Page, & Baddeley, 1996), and is based upon a model of short-term memory for serial order (Brown, Preece, & Hulme, 2000).

The basic architecture of the OSCAR model is illustrated in Figure 7. I will begin the discussion of the OSCAR model by looking at the *phonological-context vector*. All the elements of this vector change over time as a function of the *oscillators*. These oscillators are divided into two subsets, repeating and non-repeating. The repeating oscillators all have the same frequency, whilst the non-repeating group consists of oscillators of different frequencies. The elements that make up the phonological-context vector are the product of the output from certain oscillators. Half of the elements represent the non-repeating and half the repeating oscillators.

Vousden et al. (2000) explain this mechanism in terms of a clock face analogy that was offered by Brown et al. (2000). They suggest that the oscillators are the hands on a clock, whilst the phonological-context vectors are the time displayed on the clock face. The repeating oscillators can be thought of as the minute hand that repeats every hour. On its own it can only discriminate between times within the hour, and this is analogous to encoding the information concerning the phonemes within a syllable, but provides no information to identify which syllable the phonemes were in. The non-repeating oscillators can be thought of as the hour hand, which allows for the discrimination between "times" that would otherwise appear identical. The non-repeating oscillators, therefore, allow for the encoding of information relating to the order of the syllables. In this way the combination of the repeating and non-repeating oscillators allows for the accurate positioning of each phoneme along the temporal continuum.

The final part of the model is the *phoneme feature vector*. Each phoneme in a sequence becomes associated with a state of the phonological context vector (illustrated in Figure 7) by the connections between the phonological feature and phoneme feature vectors. Each phoneme in a sequence will become associated with a different state of the phonological context vector, due to the dynamic nature of oscillators. The resulting associations are stored in a Hebbian weight matrix.

The recall of the sequence is achieved by resetting the oscillators to their initial state and, due to the fact that each oscillator possesses its own frequency, there is no need to recall all the successive states of the phonological-context vector. Allowing the oscillators to run through their cycles will naturally result in the phonological-context vectors producing all the states that had become associated with the phonemes and will, therefore, produce the phonemes in the correct order. However, to avoid the continued dominance of the currently activated phoneme, a "switch-off" mechanism is employed in the form of post-output suppression. This suppresses the currently active phoneme and so allows the next phoneme to be correctly produced. A similar technique is used in the Hartley and Houghton (1996) model.

Vousden et al. (2000) also use the clock face analogy to explain the learning and recall of a sequence. To take the simple example of encoding and recalling the phoneme sequence /tʌb/ (the word *tub*), then the production of the phoneme /t/ could become associated with the initial configuration of the clock face, e.g. 12:00. The second phoneme /ʌ/ might become associated with 12:20, and the final phoneme /b/ with 12:40. During the recall stage the clock face is reset to its starting position, 12:00, and allowed to run forward. Thus, at 12:00 the phoneme /t/ would be recalled, at 12:20 /ʌ/ would be recalled and, finally, at 12:40 /b/ would be recalled. In this example the hour hand (non-repeating oscillators) does not complete a cycle of the clock face, and so the phoneme sequence represents one syllable.

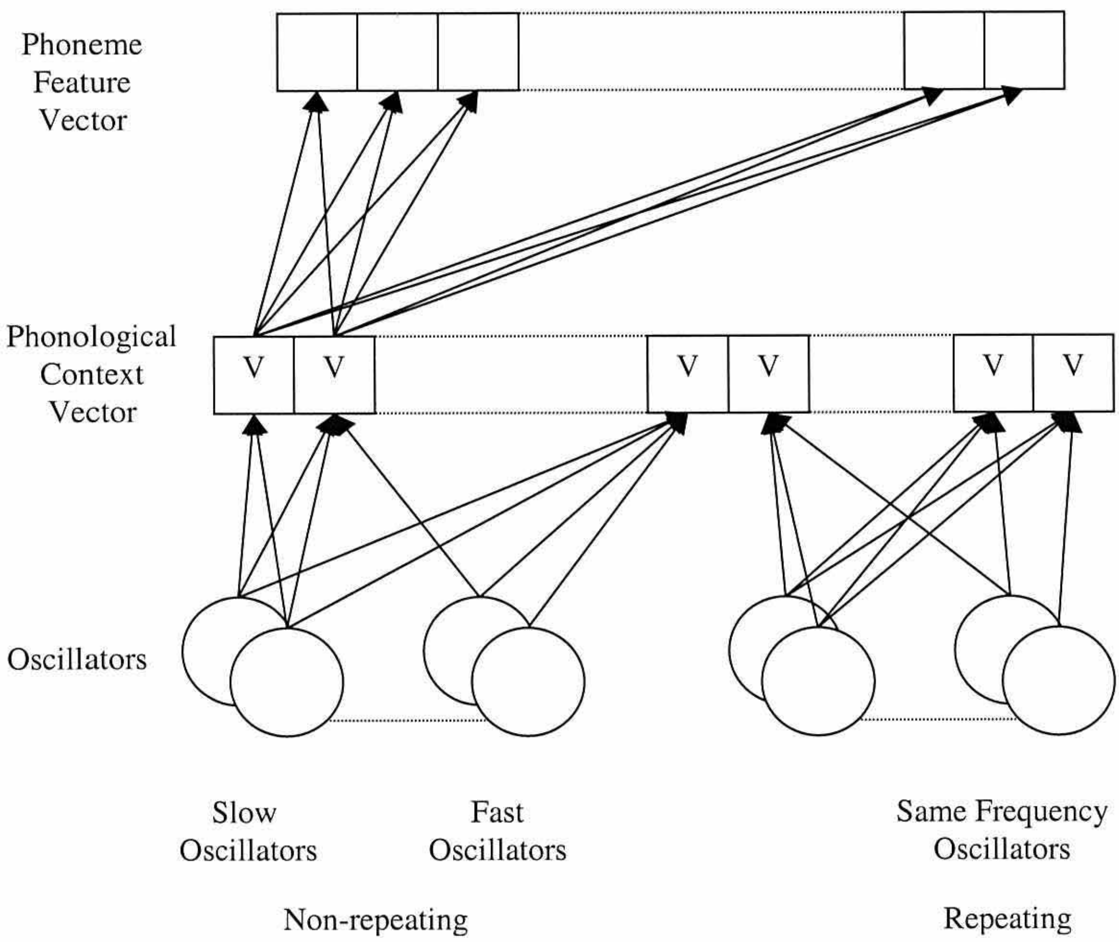

Figure 7: Simplified architecture of the OSCAR model. Circles at the bottom represent oscillators, the phonological context vectors are represented by the middle row containing the letter V, and the phoneme feature vector is represented by the top row. (Adapted from Vousden, Brown, & Harley, 2000).

In order to allow for error opportunities in the model, noise was introduced at the point when a state of the phonological context vector is used to retrieve a target phoneme. A parameter was used to decide what proportion of attempts to produce a phoneme were subject to noise. If the attempt were subject to noise, as determined by the parameter, then a non-target phoneme that was associated with a similar state of the phonological-context vector would also be activated. The more similar the two states of the phonological-context vector, the more chance that the two associated phonemes would be co-activated and, therefore, lead to an error. To use the clock face analogy, I will consider a sequence of phonemes that were associated with 12:00, 12:20, 12:40, 1:00, and 1:20. The clock face 12:20 is used to recall the second phoneme of the sequence, but the clock face 1:20 is very similar and so (if the attempt were noisy) there would be a chance that the phoneme associated with 1:20 would be produced in error.

The results of simulations run by Vousden et al. (2000) demonstrated that the model could account for order errors (anticipations, perseverations, and exchanges) and item errors (non-contextual substitutions).

A Problem for these Models

Two of these models (Dell, 1986; Hartley & Houghton, 1996) make an assumption about the phonotactic constraints of a given language (English in both cases), which dictates to a greater or lesser extent how the model is implemented. This assumption is based upon the evidence that phonemes follow the *syllable position effect* (e.g. Boomer & Laver, 1968; Fromkin, 1971; MacKay, 1970; Nooteboom, 1969). The syllable position effect states that phonemes maintain their position within the syllable, even when they are produced in a different syllable than that which was intended. If evidence could be found that this is not the case in a substantial proportion of such speech errors, then models of phonological encoding would need to take account of this. The model by Vousden et al. (2000) does not contain a representation of the phonotactics and, therefore, this would not seem to cause a problem. It is with this in mind, that I now turn to a series of experiments conducted by Dell, Reed, Adams, and Meyer (2000), which possibly provides such evidence.

Dell, Reed, Adams, and Meyer (2000) – A Study of the Role of Experience in Language Production.

Dell, Reed, Adams, and Meyer (2000) studied the role of experience in language production, and in particular, the way in which the phonotactic constraints of a speaker's native language manifested themselves in the speech errors made by the speaker.

The phonotactic constraints of a language serve a number of purposes. Firstly, they allow us to judge whether or not a word is phonologically legal in a given language. This helps us when we come across an unfamiliar word, allowing us to decide whether it is a word that we have not yet encountered, or if it cannot be a legal word in a given language. For example, in English a syllable cannot begin with the phoneme /ŋ/. Therefore, a word beginning with "ng" would be discounted as a possible word. However, in Welsh a syllable may begin with /ŋ/, and so would not be discounted. Knowledge of possible sequences also aids in the identification of ambiguous speech sounds (e.g.

Massaro & Cohen, 1983; Pitt, 1998), and helps to identify word boundaries (e.g., McQueen, 1998; Norris, McQueen, Cutler & Butterfield, 1997). It has also been shown that common sound sequences can be spoken more quickly than sequences with low frequency phonotactics (Vitevitch, Luce, Charles-Luce & Kremmerer, 1997). The ability to use this phonotactic knowledge begins very early. Infants as young as nine-months old have been shown to be sensitive to the constraints of their native language (e.g. Aslin, Saffran & Newport, 1998; Jusczyk, Frederici, Wessels, Svenkund, & Jusczyk 1993; Mattys, Jusczyk, Luce, & Morgan, 1999). And, importantly for this study, it has been demonstrated that 8-month old infants can segment a continuous stream of speech syllables into word like units after only 3 minutes of listening experience (Aslin, Saffran, & Newport, 1998; Saffran, Aslin, & Newport, 1996). This shows that a powerful mechanism for extracting statistical information from speech is present at a very early age and, as will be shown later, this mechanism is central to the present study and that of Dell et al. (2000).

Various studies have suggested that when a speaker makes an error, violations of phonotactic constraints are rare. Phonological speech errors can produce nonwords (e.g. "hymn to hing" from "hymn to sing" – from Dell, Schwartz, Martin, Saffran, and Gagnon 1997), but errors that violate the phonotactic constraints of the language constitute a very small percentage of recorded corpora (e.g. Boomer & Laver, 1968; Fromkin, 1971; Wells, 1951). However, such errors do occur. For example, Stemberger (1983) reported several examples, but these constituted less than 1% of the total error corpus. Therefore, it is generally agreed that there is a strong effect present, which produces a tendency for errors to follow the phonotactic constraints of the speaker's native language. This tendency has been termed the *phonotactic regularity effect*, and it is suggested that these constraints are used during phonological encoding (Fromkin, 1971; Motley & Baars, 1975).

The study by Dell et al. (2000) aimed to investigate the mechanisms that underlie the phonotactic regularity effect. The experiments conducted were motivated by two hypotheses: the *breadth of constraint hypothesis* and the *implicit learning hypothesis.* I will briefly discuss these hypotheses, before moving on to the details of the actual experiments.

The breadth of constraint hypothesis claims that language contains patterns at many different levels, and that the mechanisms responsible for processing language are sensitive to these patterns. For example, there are language-wide patterns such as /ŋ/ always being a coda whenever it occurs in English. There are also patterns that apply to a smaller set of words. The differing stress patterns of English nouns and verbs are an example. English nouns tend to have trochaic stress as in *IMport*, whereas verbs tend to have iambic stress as in *imPORT.* One could also find patterns that apply to only one word, for example /k/ is always the first phoneme in "keg". As Dell et al. (2000) admit, this hypothesis is not controversial. However, they intended to reinterpret speech-error phenomenon in light of this hypothesis, and identify relationships between these phenomena that have not been highlighted before.

Dell et al. (2000) suggest that we can think of a particular speech-error phenomenon in terms of very local constraints. In support of this they quote the example of the phrase "reading list" being incorrectly spoken as "leading list" (Fromkin, 1971). In this example the phoneme that has been moved, /l/, has maintained its position as an onset, thus displaying the *syllable position effect*

(Boomer & Laver, 1968; Fromkin, 1971; MacKay, 1970; Nooteboom, 1969). Exceptions to this effect do occur (MacKay, 1970; Stemberger, 1982), but they would seem to constitute a very small percentage of the recorded errors. Dell (1986) proposed a model that accounted for the syllable-position effect by associating words with position-specific phoneme nodes. For example, the word "list" would activate an onset-/l/ node, which would be separate from the coda-/l/ node. Therefore, in the example from Fromkin (1971), the error ("reading list" → "leading list") would occur if the onset-/l/ node was activated prematurely, i.e. instead of activating the onset-/r/ node for the initial phoneme of "reading", the network anticipates the initial phoneme of "list" and activates the onset-/l/ erroneously. However, the model would have difficulty if an error required the movement of a phoneme from an onset position to a coda position, or vice-versa (e.g. "red dog" → "ged rog"). This is due to the separate representations of onset and coda phonemes, and the difficulty that the network would have in turning a coda phoneme into an onset phoneme.

Dell et al. (2000) suggest that the syllable-position effect and the phonotactic regularity effect are related phenomena, and merely represent constraints at different levels of generality. For example, the fact that /ŋ/ is always a coda in English is an example of a language-wide constraint, and, due to the breadth of its applicability, errors tend to uphold the constraint very strongly. The fact that /l/ is always in the onset position of "list" is an example of a local constraint. Errors in the vicinity of "list" will tend to preserve the onset position of /l/, as in "leading list". In this way the syllable-position effect can be thought of as a local constraint. Dell et al. claim that the syllable-position and the phonotactic regularity effects can be thought of as opposite ends of a breadth of constraint continuum, with the syllable-position effect at the narrower end and the phonotactic regularity effect at the wider. In order to support their claim, Dell et al. aimed to find a speech error that occupied the middle ground of this continuum.

The fact that we are sensitive to these patterns within language sheds light upon the systems controlling speech. The implicit learning hypothesis proposes some basic properties that this mechanism should possess. Dell et al. (2000) offer the following simple description of the function of this mechanism, "The simplest answer is that the processing system learns. It experiences sound sequences and stores them in memory. Then it uses those memories in subsequent processing of sound sequences" (p. 1356). Dell and colleagues refer to three separate properties associated with this processing mechanism. They state that the mechanism should be sensitive to recent experience, be implicit, and be capable of generalisation. Dell et al. tested these three hypotheses, along with the breadth of constraint hypothesis, in a series of experiments.

Participants in the study conducted by Dell et al. (2000) were required to pronounce sequences of four consonant-vowel-consonant (CVC) syllables, such as "neng kef mes heg". These sequences were printed on paper and viewed, one at a time, through a slot cut in a piece of card. The sequences always included one /h/ and one /ŋ/. These phonemes were always required to be onsets and codas, respectively, in English. Therefore, they expected any speech errors that included these phonemes to follow the phonotactic regularity effect, i.e. /h/ would always be an onset and /ŋ/ would always be a coda. This would represent the wide end of the breadth of constraint hypothesis. Within the study, the sequences also contained two phonemes (e.g. /f/ and /s/ in the first experiment) that had constraints placed upon them which only applied within the experiment. For example, in the first experiment, half of the participants were given sequences in which /f/ was always an onset and /s/ was always a coda. The converse was true for the other participants. It was suggested that errors that involved these phonemes would follow the experiment-wide constraints, thereby providing a speech-error that occupied the middle ground of the breadth of constraint continuum. Errors involving the other 4 phonemes that were used in the sequences (/m/, /n/, /k/, /g/) could occur in either the onset or coda position, as is true for English, and were expected to follow the syllable–position constraint, and would, therefore, demonstrate the narrow end of the continuum. The three hypotheses connected with the implicit learning hypothesis were tested as follows. Firstly, the suggestion that the mechanism should be sensitive to recent experience was tested with the experiment-wide constraints. If the participants' errors followed the constraints imposed in the experiment, then the mechanism could be said to be sensitive to recent experience. If this were not the case, then the participants' errors would include a high proportion of errors that broke the experiment-wide constraints. Secondly, questioning the participants at the end of the study tested the implicit nature of the mechanism. Half of the participants were not informed of the experiment-wide constraints, and these participants were asked to report anything they had noticed about the pattern of the sounds within the sequences they had just read. If they did not report the constraints placed upon the /f/ and the /s/, Dell et al. claimed that this would be evidence for the mechanism being implicit. The implicit nature of the learning would also be supported, if there were no significant difference in the rates at which the experiment-wide constraints were upheld by the informed and uninformed participants. Finally, the generalisability of the mechanism was tested, in a later experiment, by varying the vowel in the sequence (it was always /ɛ/ in the original experiments), and having separate constraints on the consonants depending on the vowel used in the syllable. This is not of interest in the present study, but the results of the Dell et al. study did not provide conclusive evidence for the generalisability of the mechanism.

Dell et al. (2000) claim that support was found for the middle ground of the breadth of constraint hypothesis, due to the rate at which the experiment-wide constraints were upheld (94.7 – 97.7%) falling between that for the language-wide constraints (100%) and the local positional constraints (68.2% – 77.5%). They also claim to have found support for the sensitivity of the mechanism to recent experience, due to the fact that the experiment-wide constraints were upheld at a higher rate than that for the local positional constraints. And, they also claim support for the implicit nature of the mechanism, due to there being no significant differences between participants that were, and were not, made aware of the experiment-wide constraints, $p > .05$. However, the local positional constraints were upheld at a much lower rate than expected which, as will be discussed later, is of particular interest to this thesis.

Is the Paradigm of Dell et al. (2000) a Test of Reading Aloud?

The main problem with the paradigm developed by Dell et al. (2000) is that it would appear to be closer to a test of reading aloud than a test of phonological encoding, although it is impossible to test spontaneous phonological encoding in its true form without using naturalistic methods.

In addition to the obvious fact that the test involves the participants reading nonsense words from a sheet, looking at models of reading aloud, and the way that they propose that grapheme-to-phoneme conversion is handled, also support the concerns about this limitation.

There is debate concerning models of reading aloud. In particular, there is disagreement as to whether a single-route model is sufficient, or whether a dual-route model is required to represent the data for human performance on this task (e.g. Plaut, & McClelland, 1993; Plaut McClelland, Seidenberg, & Patterson, 1996, who propose single-route models, and Zorzi, Houghton, and Butterworth, 1998, who argue for a dual-route model). This argument is not of concern in this thesis, but I will briefly examine dual-route models of reading aloud, in order to demonstrate the problems inherent in the paradigm developed by Dell et al. (2000).

Dual-route models propose that there are two possible ways that written words (in alphabetic scripts) may be pronounced. Firstly, pronunciation may be achieved through a *look-up* procedure, in which the form of the word as a whole is considered, and the visual form used to retrieve the pronunciation from an internal lexicon. This method could not be used for the pronunciation of nonwords. Secondly, they may be pronounced through spelling-to-sound (grapheme-to-phoneme) conversion, which may also be used for reading nonwords. This second method, also referred to as the *sublexical* route, examines each grapheme, and produces the dominant pronunciation based upon the relationship between graphemes and letters. In this way the look-up procedure would provide the unusual, or exception, pronunciations for known words (such as the /æ/ in *have*, but the grapheme-to-phoneme conversion route would provide the /eI/ for *mave* based upon the regular pronunciations of *gave*, *save*, *pave* etcetera). Several models propose such a representation (e.g., Baron & Strawson, 1976; Besner & Smith, 1992; Coltheart, 1978, 1985; Morton & Patterson, 1980; Paap & Noel, 1991; Patterson & Morton, 1985; Zorzi, Houghton, & Butterworth, 1998).

The problems arise for the paradigm of Dell et al. (2000) when it is also proposed that the phonological output is the result of a "horse-race" between the two routes, especially considering that the stimuli in the Dell et al. paradigm are mainly nonwords. If the two routes are acting in parallel, then the exception words will lead to disagreements, and so lead to errors or a difference in pronunciation of the same nonwords. Most of the stimuli used by Dell et al. would seem to be relatively straight-forward pronunciations, and so would not seem to cause any such difficulty. But, this extra step in the task, along with the increased possibility of an error not associated with phonological encoding per se, is undesirable, and so should be removed if possible. With this in mind, I will now turn to two studies of phonological encoding that successfully remove the necessity to convert graphemes to phonemes.

In the first of these studies, Treiman and Danis (1988) asked whether syllables are coded in terms of onset and rhyme units, and whether the rhyme was then coded in terms of subunits. Although the findings of this study are of interest in terms of phonological encoding per se, the ideas explored in this paper were not incorporated into the present study. This was due to the differences in the nature of the stimuli in the two studies. In the Treiman and Danis paper it was necessary for the vowel in the CVC syllable to vary, yet in the Dell et al. (2000) study it was important to keep the vowel constant. Thus, it is only the manner in which the stimuli were presented that is of interest to the present study. Participants in the Treiman and Danis study were played lists of phonemes through headphones. After they had heard the list of phonemes, they were asked to repeat the phonemes in the order that they had heard them. Their responses were recorded, and the recordings were then checked for errors. These were counted and classified in accordance with the requirements of the experiment. This simple technique could easily be adapted to the paradigm developed by Dell et al., and so remove the reading element from the study.

The second study I will look at is a direct extension of the Dell et al. (2000) study, and so is of particular interest. However, it uses reaction times as its main measure of the participants learning of phonotactic information, so does deviate slightly from the present studies.

Onishi, Chambers, and Fisher (2002) asked whether phonotactic rules, that were not present in English, could be acquired by listening to streams of syllables in which artificial constraints had been imposed upon the consonants. To do this they modified the paradigm developed by Dell et al. (2000).

The stimuli used were similar to those used by Dell et al., i.e. they were CVC syllables in which the phonotactics of English applied, but in which novel constraints were also applied to certain phonemes. Unlike the study by Dell et al. the stimuli were spoken and recorded. Therefore, any knowledge acquired about the constraints had to be from auditory information and not from graphemic knowledge.

The procedure involved the participants listening to the stimuli through the headphones, with each stimulus repeated between four and six times. During this time, the participants were asked to rate the clarity of pronunciation on each trial (this had nothing to do with the collected data). After this the participants completed a distracter task (simple mental arithmetic), before moving on to the final stage of the study. This involved the participants listening to syllables and repeating them as quickly and accurately as possible. These syllables were either: items from the previously heard lists; items that were not in the previous list, but which followed the same phonotactic rules; or items that broke the phonotactic constraints. The time between stimulus offset and response onset was recorded by using a voice-activated response key. This response time was used to produce a mean response time for each of the three types of stimulus, with the hypothesis that response times would be significantly quicker for items that were in the original stimuli list, or which followed the constraints of the original stimuli, than for the items that broke the phonotactic constraints. This was found to be the case in a series of experiments, which provided support for the use of reaction times as a measure of the learning of phonotactic constraints.

Three experiments were conducted using this technique, which addressed three different questions. The first experiment asked (in line with the first studies of Dell et al., 2000) whether new phonotactic rules could be acquired through brief listening experience (e.g. /p/ is always an onset and /b/ is always a coda). The second experiment asked whether second order constraints could be acquired in the same way (e.g. /p/ is always an onset when the vowel is /I/, but is a coda when the vowel is /i/ - Dell et al. also explored this, but the results were inconclusive). The final experiment asked whether second order phonotactic constraints could be acquired based upon changes in the speaker, rather than the vowel within the syllable (e.g. /p/ is always an onset when the syllables are pronounced by speaker A, but /p/ is a coda when the syllables are pronounced by speaker B). The results of the first two experiments showed that both first and second order constraints can be acquired through listening to the syllables being spoken. However, the second order constraints could not be acquired through changes in the speaker, even though the task is no more difficult than that posed in the second experiment (i.e. acquiring the constraints through changes in the vowel).

These studies (Treiman & Danis, 1988; Onishi et al., 2002) show that presenting stimuli through the auditory modality can be successfully applied to such studies, and could prove a valuable modification to the paradigm developed by Dell et al. (2000).

The Present Study

The present series of studies used the paradigm developed by Dell et al. (2000) to explore the acquisition of the phonotactics of a language, and the effects of phonemic constraints upon speech errors. In addition to the exploration of phonotactics per se, the paradigm was also employed to gather data that may lead to the development of a more accurate model of phonological encoding. In particular, evidence that the syllable position effect may not be as strong as previously thought. It was also the intention to further develop the paradigm of Dell et al., in order to make it a more accurate tool for the collection of speech errors, by modifying it for the auditory modality. The studies conducted fell into four main areas:

1. *Replication of the Dell et al. Experiment and the transfer of the paradigm to computer control.* This was done in order to verify that the paradigm did work, and to check that similar results could be achieved (especially in relation to the syllable position effect). The experiment was then transferred to computer, so that the order of presentation could be more easily randomised, the speed of presentation could be kept constant, and to verify that the results achieved by Dell et al. were not an artefact of their method of presentation.
2. *Bilingual study.* A bilingual (Welsh-English) version of the paradigm was developed in order to test whether (as suggested by Dell et al.) the strength of the language-wide constraints was due to the exposure to the phonotactics of a speaker's native language. If this were the case, then the constraints should not be as strong in the speaker's second language.
3. *Time-course studies.* A series of studies were conducted to test the time-course and robustness of the learning that took place when the participants were exposed to the novel phonemic constraints. This was done by reversing the constraints on two phonemes once participants had learnt the previous constraints. The confusion caused by this reversal (e.g. continuing to use the previous constraints after they had been reversed) and the length of time that the confusion continued, would give an insight into the process of learning phonemic constraints.

4. *Development of an auditory version of the paradigm.* Studies were conducted using an auditory version of the paradigm. This was developed by using the technique used by Treiman and Danis (1988), who used an auditory paradigm in their study of speech errors. This move was taken in order to remove the spelling sound conversion that it was necessary for the participants to perform, and so make the paradigm a more accurate test of phonological encoding and not one of reading aloud. This auditory paradigm was also used in the exploration of the time-course and robustness of the learning.

Thus, this thesis has three main aims. Firstly, it is intended to further explore the acquisition of the phonotactics of a language. Secondly, it is intended to develop the paradigm of Dell et al. (2000) in to a useful tool for the collection and analysis of speech errors. And finally, the data collected from these studies will be used to propose changes to the currently proposed models of phonological encoding.

Chapter 2: General Experimental Method, Error Coding, and Analysis.

The basic experimental paradigm that was shared by all the experiments in this study is described in this chapter. The method section of the individual experiments will add any details that varied from the basic paradigm. This chapter also includes an overview of the analyses that were performed upon the data, which followed the methods used by Dell et al. (2000).

Participants

All participants were students from the School of Psychology, at the University of Wales, Bangor, and were recruited on the basis that English was their native language. None of the participants had been diagnosed as dyslexic, and they all had normal or corrected to normal vision. They received course credits for participation.

Stimuli

For each study two sets of 96 sequences were generated, each of which observed the following criteria. Each of the sequences contained four CVC syllables. Each syllable had /ɛ/ as its vowel, and the consonants /h/, /ŋ/, /m/, /n/, /g/, /k/, /f/, and /s/, each appeared once per sequence of 4 syllables (Therefore imposing the constraint that no syllable consisted of C_1VC_1). Within every sequence /h/ was always an onset and /ŋ/ was always a coda. Two of the remaining consonants were constrained in the same way (i.e. one always being an onset and one always a coda), with the constraints being reversed for the second of the two sequences. The pairs of constrained consonants were either /f/ and /s/ or /k/ and /g/.

These constraints allow for a vocabulary of 32 syllables in each condition. In each of the 96 sequences - a total of 384 (4 x 96) syllables - the syllable types appeared with the expected frequencies shown in Appendix 1, which represents all the syllable types in the condition where /f/ was always an onset and /s/ was always a coda. The frequencies with which the syllables appeared differed because each of the consonants had to appear once in each sequence, yet some of the consonants (/h/, /ŋ/, and the pair of constrained phonemes) were restricted to being either an onset or a coda. For example, the syllable *heng* occurs more frequently than *mek*, as the constituent phonemes of *mek* may also appear as *kem*, but the phonemes in *heng* can not appear in any other form. Each consonant appeared once per sequence so that the error opportunities amongst the consonants would be equal, and to allow for the evaluation of any syllable position constraints which may be followed by the unconstrained phonemes (/m/, /n/, and either /k/ and /g/, or /f/ and /s/).

The spelling of the syllables was straightforward except that /g/ when used as an onset was spelt "gh", in order to remove any hard-soft ambiguity associated with its pronunciation. The order of the syllables within the sequences was randomised for each sequence and each participant.

Procedure

The participants viewed the sequences one at a time. After the first sequence was made visible, an electric metronome (Model QM2 taktell, made by Wittner, Germany) was started at a rate of one beat per second, and the participant was asked to read 22 sequences of four syllables with each syllable coinciding with a beat. This slow rate of articulation was to ensure that the participants were pronouncing the syllables correctly, and to familiarize them with the sequences. This slower rate of production was not recorded. The metronome was then set to a rate of 2.53 beats per second and the participant produced the sets of 96 sequences at this faster rate. Each sequence remained visible until the participant had finished pronunciation. There was a brief pause (approximately 1 minute) between sets. Each session was recorded on mini-disc (Sony mini-disc recorder, model MZ-R500, Sony Corporation) using a condenser microphone (SoundLAB ALEM-106), which was amplified using a pre-amp made by the technical staff in the School of Psychology, University of Wales, Bangor. Participants in the informed conditions were told of the constraints placed upon the phonemes, e.g. "when you see an "f" then it will always be at the beginning of a syllable, and when you see an "s" then it will always be at the end of a syllable". Those in the uninformed conditions were told nothing about the distribution of the phonemes. At the end of the sessions, participants in the uninformed conditions were given a sheet of paper with the following instruction: " In the space below, please note any observations concerning the syllables that you were asked to pronounce in this experiment. For example: All of the syllables contained the vowel *e*."

Coding of Speech Errors

Recordings from all experiments were checked for speech errors. For each experiment a second coder, who was unaware of the experimental conditions, independently listened to the recordings from the first session of three randomly selected participants and identified errors. Both coders had access to the correct sequences. From the syllable productions that were examined by the coders the following figures were calculated: the numbers of syllables on which the coders agreed that there was an error and upon the nature of the error (E); and the number of syllable productions on which both coders agreed that there was no error (NE). The sum of these two was then divided by the total number of syllable productions examined (TS) to arrive at an overall rate of agreement (RA). Therefore, the overall rate of agreement was calculated as follows:

$$\frac{E + NE}{TS} = RA$$

The level of agreement across all experiments in this study was between 94 - 98% and so was acceptable. No alterations were made to the error codings as a result of these reliability checks.

Analyses

Analyses were not performed on the sequences produced in the practice sessions, which were only produced to allow the participants to familiarise themselves with the task.

The focus of the study was on errors in which one of the eight consonants was replaced by another of the consonants used in the sequence. All the consonant errors were tallied. This included cut-off errors, such as "kes" → "n...kes". Vowel errors or errors that involved the substitution of a consonant by a consonant not in the sequence, were rare and were not included in the error corpus.

The syllables produced at the faster rate were checked for errors, and the number of syllables that contained misordered phonemes was counted. From this the overall error rate was calculated. All of these erroneous syllables were represented in error matrices. The entries in these matrices indicated the number of times that a syllable had been produced in error. They were not confusion matrices, but represented the form taken by the erroneous syllable. The columns in the matrices represent the onset and the rows represent the codas. So, for example, in Table 1 the 4 in the first row indicates that the syllable /kɛg/ was produced in error on four occasions. The matrices do not indicate the source of the error, and so it may have been the onset, coda, or both phonemes that were produced erroneously.

The matrices also indicated whether the syllable broke any positional constraints, by labelling the errors as legal or illegal. Errors in the legal matrices involved movements in which the phonemes occupied the same syllabic position in the error as they did in the original sequence. For example, if the participant said /gɛŋ/ / nɛk/ /hɛk/ /sɛm/, instead of /gɛŋ/ / nɛk/ /hɛf/ /sɛm/, then this would count as a legal error since the /k/ occupied the same position in the error as in the source syllable. Table 1 gives an example of a legal matrix for the condition in which /f/ was constrained to be an onset and /s/ to be a coda.

Table 1: Example of a Legal Error Matrix (/f/ onset, /s/ coda).

	fes condition – legal outcomes							
	Onset							
Coda	g	k	m	n	f	s	h	ŋ
g		4	20	11	11		21	
k	3		8	6	12		9	
m	1	4		2	5		10	
n	9	19	10		16		27	
f								
s	8	9	12	12	8		23	
ŋ	29	7	10	11	17		28	
h								

Within the legal matrices the /ŋ/ onset and the /h/ coda are necessarily zero as these could never be legal errors in English. The onsets and codas for the constrained phonemes are also necessarily zero (depending upon the phonemes and the condition), as these could never be legal errors under the experimental constraints. Also note that outcomes in which the same phoneme occurs twice within the same syllable are also zero (e.g. /kɛk/), as one of the phonemes would always have broken a positional constraint.

The illegal matrices reflect errors in which a consonant is produced in a different syllable position than that which it occupied in the original sequence. Therefore, an illegal error involving /h/ or /ŋ/ would violate the language-wide constraints, an illegal error involving the artificially constrained phonemes would violate the experiment-wide constraints, and an illegal error involving any of the other consonants would violate the local positional constraints. Also note that four syllables within the illegal matrices are necessarily zero as these could only be legal errors. Table 2 gives an example of an illegal matrix in which /f/ was constrained to be an onset and /s/ a coda. Within the illegal matrix a single asterix (*) indicates the violation of an experiment-wide constraint, a double asterix (**) indicates the violation of a language-wide constraint, and the absence of an asterix are violations of local positional constraints only.

Table 2: Example of an Illegal Error Matrix (/f/ onset, /s/ coda).

	fes condition – illegal outcomes							
	Onset							
Coda	g	k	m	n	f	s*	h	ŋ**
g	10	3	11	13	11	-	21	-
k	4	12	4	-	5	-	1	-
m	1	4	3	8	9	-	1	-
n	7	3	7	5	7	-	7	-
f*	-	-	-	-	-	-	1	-
s	4	8	14	3		-		-
ŋ	11	2	14	1		1		-
h**	-	-	-	-	-	-	-	-

The total number of legal and illegal errors was then calculated for each phoneme. These were then combined into the appropriate groups to calculate the total number of legal and illegal errors for each of the constraints. From these, the percentage of errors that upheld the constraints was calculated. The percentages were calculated for the local positional constraints, language-wide constraints, and the experiment-wide constraints.

If a constraint was upheld, then the number of legal phoneme misorderings would be significantly higher than 50% (there being only two possible types of misorderings – legal or illegal). This being the case, a binomial test was required, and as the data was of nominal level of measurement (i.e. legal or illegal) then a sign-test was chosen. This would also allow for a direct comparison with the results obtained by Dell et al. (2000), as they performed the same analysis. The use of sign-tests also reduced the necessity for outliers to be removed, as sign-tests are extremely robust under these conditions. Therefore, outliers were not removed unless they constituted a large part of the data (i.e. more than one participant in a single experiment for this series of studies).

In the experiments where the variable was applied, the errors for the experiment-wide constraints were compared between the informed and uninformed participants. If the "learning" of the new constraints was not due to explicit knowledge then there would be no significant difference between the rates at which the constraints were followed for the informed and uninformed participants. However, if the "learning" was due to explicit knowledge, then the informed participants should make significantly less illegal errors on the experiment-wide constraints than those participants who were uninformed. This was again analysed using a sign-tests. The questionnaires completed by the uninformed participants were examined for any responses that indicated that they had acquired explicit knowledge of the artificial constraints placed upon the phonemes. If there was no significant difference between the number of legal and illegal experiment-wide errors made by the two groups, and no evidence of explicit knowledge, then this would be taken as evidence for the "learning" being implicit.

Chapter 3: Experiments 1 and 2 - Replication of Dell, Reed, Adams and Meyer (2000)

The first two experiments in this study set out to replicate the findings of Dell et al. (2000). Experiment 1 was an exact replication of the first experiment conducted by Dell and colleagues. The replication of Experiment 1 was conducted in order to test the reliability of the Dell et al. paradigm. In particular, I wanted to replicate the rate at which the syllable position effect had been upheld (68.2 – 77.5% in the Dell et al. study). Previous reports had placed this much higher, e.g. Ellis (1980) reported rates of between 84 - 96% for the syllable position effect. It was hypothesised that the results of Experiment 1 would be similar to those of the Dell et al. study.

If the rates reported were found to be replicable, then doubt would be cast upon models of phonological encoding that use separate representations of the onset and coda versions of the same phoneme (e.g. Dell, 1986, 1988). On the other hand, if such evidence were found, then dual-route models that encode phonemes and syllabic structure separately (e.g. Hartley & Houghton, 1996) offer a solution to the problem of modelling such data. A dual-route model would allow for an onset phoneme to be moved to the coda position, as the two elements (phonemes and syllabic structure) are not combined until the output layer, allowing for the error to occur.

Experiment 2 made alterations to the Dell et al. (2000) study, whilst seeking to achieve comparable results. These alterations included using computer software for the presentation of the stimuli, which allows for easier randomisation of the order of presentation, and removes errors involved with the manual presentation of the stimuli (e.g. errors in the handling of the stimuli sheets). There was also a reduction in the number of days over which the experiment was conducted. In line with the second experiment conducted by Dell et al., the experimentally constrained phonemes were also changed in order to check that the results were not an artefact of some property of the chosen phonemes. It was hypothesised that these alterations to Experiment 2 would not significantly alter the results from Experiment 1, but that the changes would allow for a more accurate, and quicker, version of the paradigm.

Experiment 1

Method

Participants

Eight participants were used for this experiment. They were all students from the School of Psychology, at the University of Wales, Bangor, and were recruited on the basis that English was their native language. None of them had been diagnosed as dyslexic, and they all had normal or corrected to normal vision. They received course credits for participation. They were randomly assigned to the conditions, i.e. two in each of the four conditions. The four conditions were: sef informed, fes informed, sef uninformed, and fes uninformed. Sef and fes refer to the experiment-wide constraints that were applied to /f/ and /s/ in each condition (e.g. fes refers to the condition where /f/ was always an onset and /s/ was always a coda). Informed and uninformed refer to whether the participants' were made aware of these constraints prior to beginning the experiment.

Stimuli and Materials

In this study, /f/ and /s/ were the experimentally constrained phonemes. For participants in the fes condition, /f/ was always an onset and /s/ was always a coda. For participants in the sef condition the reverse was true. The consonants /h/ and /ŋ/ were the language-wide constraints, and the consonants /m/, /n/, k, and /g/, each occurred once per sequence and were unconstrained. Each syllable had /ɛ/ as its vowel. The spelling of the syllables was straightforward except that /g/ when used as an onset was spelt "gh", in order to remove any hard-soft ambiguity associated with its pronunciation. The order of the syllables within the sequences was randomised for each sequence and each participant.

The sequences for each condition were printed in 16-point Helvetica lowercase typeface, one sequence per line and 11 lines per sheet of A4 paper. This was in line with Dell et al. (2000) who printed the stimuli rather than display them on a computer screen.

Procedure

Each participant produced 12 sets of 96 sequences, three a day on four separate days. The largest gap between sessions was 36 hours. The sequences were viewed through a cutout slot, thereby ensuring that only one sequence was visible at a time. After the first sequence was made visible, an electric metronome was started at a rate of one beat per second, and the participant was asked to read 22 sequences of four syllables with each syllable coinciding with a beat. The slower rate of articulation was only conducted at the first session as it was considered that the participants would be comfortable with the task once they had completed the first day's session. This slower rate produced very few errors, and these were not recorded.

Prior to continuing with the experiment, the participants in the informed condition were told about the constraints placed upon the phonemes /f/ and /s/, e.g. the participants in the informed "fes" condition were told, "when you see an "f" it will always be at the beginning of a syllable, and when you see an "s" it will always be at the end of a syllable". Those in the uninformed condition were told nothing about the distribution of the phonemes.

The metronome was then set to produce 2.53 beats per second, and the participant then produced the 96 sequences three times at this faster rate, with a brief pause (approximately 1 minute) between the trials. Each sequence remained visible until the participant had finished pronunciation. This process was repeated on the following 3 days.

At the end of the sessions, participants in the uninformed conditions were given a sheet of paper with the following instruction: " In the space below, please note any observations concerning the syllables that you were asked to pronounce in this experiment. For example: All of the syllables contained the vowel *e*."

Experiment1: Results

All consonant movements involving the misplacement of phonemes from the original sequence were tallied. Errors that involved the substitution of a consonant by a consonant not included in the sequence, or vowel errors, were rare and were not included in the error corpus (such errors constituted 0.12% of the data corpus).

Total Number of Errors

Of the 36,864 syllables produced by the participants at the faster rate (i.e. 4 syllables x 96 sequences x 3 repetitions x 4 days x 8 participants) a total of 1,313 syllables contained consonant misorderings. This gave an error rate of 3.56% (1,313 / 36,864). There were more errors on the first day than on any of the following days, with the final day having the fewest errors (Day 1, 473; Day 2, 305; Day 3, 311; Day 4, 224). Participants made significantly fewer errors on the final day than on the first, $p < .008$ (sign-test). All of the misordering errors are presented in Table 3, which follows the method used by Dell et al. (2000).

Table 3: Observed Errors in Experiment 1 Classified by Condition and Legality

fes condition – legal outcomes								
	Onset							
Coda	g	k	m	n	f	s	h	ŋ
g		4	20	11	11		21	
k	3		8	6	12		9	
m	1	4		2	5		10	
n	9	19	10		16		27	
f								
s	8	9	12	12	8		23	
ŋ	29	7	10	11	17		28	
h								

Table 3: continued

fes condition – illegal outcomes

	Onset							
Coda	g	k	m	n	f	s*	h	ŋ**
g	10	3	11	13	11	-	21	-
k	4	12	4	-	5	-	1	-
m	1	4	3	8	9	-	1	-
n	7	3	7	5	7	-	7	-
f*	-	-	-	-	-	-	1	-
s	4	8	14	3		-		-
ŋ	11	2	14	1		1		-
h**	-	-	-	-	-	-	-	-

sef condition – legal outcomes

	Onset							
Coda	g	k	m	n	f	s	h	ŋ
g		1	10	22		15	30	
k	8		10	12		13	26	
m	8	7		14		7	16	
n	13	15	4			6	18	
f	9	11	9	10		59	25	
s								
ŋ	10	8	6	9		24	5	
h								

sef condition – illegal outcomes

	Onset							
Coda	g	k	m	n	f*	s	h	ŋ**
g	12	13	7	4	2	14	19	-
k	9	14	5	3	-	6	5	-
m	2	6	5	3	-	9	5	-
n	10	13	10	13	-	23	8	-
f	4	2	7	7	-			-
s*	1	-	-	-	-	-	-	-
ŋ	11	8	6	8	1			-
h**	-	-	-	-	-	-	-	-

- * violates experiment-wide constraints.
- ** violates language-wide constraints.

Language-Wide Constraints

From Table 3 it can be seen that none of the errors involving the misplacement of the phonemes /h/ or /ŋ/, broke the language-wide constraints, and therefore they followed the phonotactic regularity effect, $p < .008$ (sign-test).

Local Positional Constraints

Turning to the local positional constraints, those involving a misplacement of /m/, /n/, /k/, or /g/, there was a total of 1,011 errors. Of these, 48.57% obeyed the local positional constraint, and there was therefore no significant difference between the number of legal and illegal errors, $p = 1.00$ (sign-test). Four of the participants produced more legal than illegal errors, with the reverse being true for the other four participants. Across participants, the range of legal errors was 39.23% to 66.15%.

Experiment-Wide Constraints

The experiment-wide constraints, those involving the misplacement of /f/ or /s/, were upheld on 96.55% of occasions (range across participants 81.81% - 100%). Of the 174 errors involving the misplacement of /f/ or /s/, 168 followed the constraints. This result was significant, $p < .008$ (sign-test). There was no evidence that learning had occurred with relation to the experiment-wide constraints, with there being no significant difference between the number of illegal experiment-wide errors on the first and the last days, $p = 1.00$ (sign-test).

A comparison was made between the rates at which the language-wide constraints and the experiment-wide constraints were upheld. It was found that the language-wide constraints were upheld at a higher rate (100%) than the experiment-wide constraints (96.55%), there was no significant difference between the two, $p > .125$ (sign-test).

Implicit or Explicit

All of the participants in the uninformed conditions produced at least one illegal error that broke the experiment-wide constraints, but none of the participants in the informed conditions made such illegal errors. However, this difference was not significant, $p > .063$ (sign-test). When participants were asked to report any knowledge about the syllables they had been pronouncing, none of them reported anything related to the experiment-wide constraints.

Experiment1: Discussion

In comparing the results of this study with those obtained by Dell et al. (2000), I find that there were some similarities but that there were also major differences. The rate at which the language-wide constraints were upheld was identical to that in the study by Dell et al., i.e. 100% of the errors involving the movement of /h/ or /ŋ/ upheld the constraint. However, with all the other constraints there were differences of varying magnitude.

The total number of errors was lower than that observed by Dell et al. (2000), with an error rate or 3.56% compared to 8.3% in the original study. It is not known why this decrease in error rate should have occurred, as the experiment was conducted in exactly the same way as the original Dell et al. study.

The local positional constraints (those involving the movement of /m/, /n/, /k/, or /g/) were upheld at a lower rate than in the Dell et al. (2000) study. In the original study, the constraints were upheld on 68.2% of occasions and this result was significantly greater than chance. And, therefore, it supported the syllable-position effect. However, in the present study, the local positional constraints were only upheld on 48.57% of occasions and were not significantly different from chance. This result questions whether movements are actually constrained by position. The fact that the constraints were upheld approximately 50% of the time, suggests that the position a phoneme moves to (onset or coda) is merely a matter of chance. It is hypothesised that the reason why the local-positional constraints were not upheld as strongly as expected, in both the study by Dell et al. and the present study, is that in naturalistic studies speech errors are normally taken from spoken sentences where context and semantics may also play a part in the movement of the syllables. Also, in the study by Ellis (1980), the syllable position effect was explored by looking at the differing rate of transpositions between CV and VC syllables, or transpositions between syllables that shared the same structure, CV or VC, rather than between syllables that had both onset and coda consonant (CVC). This may have affected the rate with which the local positional constraints were upheld due to the there being an increased number of possible transpositions between any two syllables.

These results, as well as those obtained by Dell et al. (2000), further call into question models of phonological encoding that do not allow for exceptions to the syllable-position effect. These models, such as that of Dell (1986, 1988), use separate nodes for the onset and coda representations of the same phoneme. However, there are models that do not have such limitations. For example, the dual route model developed by Hartley and Houghton (1996).

The experiment-wide constraints were upheld at a rate very close to that obtained from the Dell et al. (2000) study (96.55% compared to 97.70%

respectively). And, as in the Dell et al. study, all participants upheld the experiment-wide constraints at a significantly greater rate than they upheld the local positional constraints. This suggests that the rate with which the experiment-wide constraints are upheld is more than an artefact of the syllable position effect. However, this is the only way in which the data from the present study is comparable to that from the Dell et al. study. Dell and colleagues also reported that there was a decrease in the number of illegal experiment-wide errors between the first and last days. No such decrease was found in this study.

Dell et al. (2000) suggested that the experiment-wide effect was not due to verbalizable knowledge, as they found no significant difference between the experiment-wide errors made by participants, whether they were informed or uninformed of the distribution of the phonemes. However, in the present study, all of the uninformed participants made at least one error involving the experiment-wide constraints, whereas the informed participants made no such errors. However, none of the participants reported noticing the rules regarding the distribution of the constrained phonemes, and thereby showing explicit knowledge. However, this technique has been criticised as being an extremely weak test of implicitness (Dulany, Carlson, & Dewey, 1984).

If I now turn our attention to the claim by Dell et al. (2000), that the experiment-wide constraints represent the "middle-ground" of a continuum between the local-positional constraints and the language-wide constraints, then I find that this claim is not upheld by statistical analysis. The present study, and that of Dell et al., both found that the language-wide constraints were upheld 100%. And similarly, the rate with which the experiment-wide constraints were upheld is similar in both studies (96.55% for the present study and 97.70% for that by Dell et al.). However, when I look at the difference between the rates with which the language-wide and experiment-wide constraints were upheld in Experiment 1, I find that it is not significant, and so these constraints do not necessarily represent different populations. Therefore, I cannot support the claim made by Dell et al.

In line with the study conducted by Dell et al. (2000), an attempt was made to replicate the study whilst placing the experiment-wide constraints on different phonemes than in Experiment 1. This was done in order to ensure that it was not some quality of the phonemes /f/ and /s/ that was producing the result, and so the phonemes /k/ and /g/ were used instead. However, several changes were made to the experiment. Firstly, it was decided to transfer the experiment to computer so that the presentation of the stimuli could more easily, and more precisely, be controlled. And secondly, the number of days on which the study was run, was reduced from four to one. This was done because no effect of learning was observed across the four sessions, and one session produced sufficient data for analysis purposes. It was hypothesised that these changes would not significantly alter the outcome of the experiment, i.e. that both Experiments 1 and 2 would support the findings of Dell et al. (2000).

Experiment 2

Method

Participants

Eight participants were used for this experiment. They were all students from the School of Psychology, at the University of Wales, Bangor, and were recruited on the basis that English was their native language, and that they had not participated in the previous experiment. None of them had been diagnosed as dyslexic, and they all had normal or corrected to normal vision. They received course credits for participation. They were randomly assigned to the conditions, i.e. two in each of the four conditions. The four conditions were: keg informed, gek informed, keg uninformed, and gek uninformed. Keg and gek refer to the experiment-wide constraints that were applied to /g/ and /k/ in each condition (e.g. gek refers to the condition where /g/ was always an onset and /k/ was always a coda). Informed and uninformed refer to whether the participants' were made aware of these constraints prior to beginning the experiment.

Stimuli and Equipment

Two sets of stimuli were prepared that fulfilled the same criteria as for Experiment 1, with the exception that /f/ and /s/ were now unconstrained, and that /k/ and /g/ were now restricted to being exclusively an onset or coda, depending on the condition. Changing the original stimuli in the following way produced this set of stimuli. All occurrences of the phoneme /f/ were replaced by a /k/, and all occurrences of /s/ were replaced by /g/. This meant that the new stimuli varied in as few dimensions as possible.

The sequences were displayed one at a time on a computer monitor (15 inch Panrix SVGA) using the program E-Prime (Psychology Software Tools Inc., Pittsburgh, PA, USA). The font was 22 point Helvetica lower case. The larger font was chosen to allow for the perceived size of the print on the screen. The participants regulated the speed of presentation, as they were required to press the spacebar on the keyboard when they had finished pronouncing the sequence.

Procedure

With the exception of the method of presentation (which now required the participants to press a button to bring up the next sequence of syllables), the alteration of the instructions for the informed participants to account for the new constrained syllables (/k/ and /g/ instead of /f/ and /s/), and the shortening of the experiment to one session, all aspects of the procedure were identical to that in Experiment 1.

Experiment 2: Results

Total Number of Errors

The errors were counted as per Experiment 1. A total of 745 syllables contained consonant movement errors, giving an 8.08% error rate for the 9,216 syllables. This is more than double the error rate for Experiment 1, however, it is not significantly different, $p > .07$ (sign-test). Table 4 gives the error matrices for Experiment 2. As in Experiment 1, a practice effect could be observed, with there being fewer errors in the final session than in the first session of the experiment, $p < .008$ (sign-test).

Table 4: Observed Errors in Experiment 2 Classified by Condition and Legality

keg condition – legal outcomes

Coda	Onset: g	k	m	n	f	s	h	ŋ
g		15	5	2	10	7	24	
k								
m		5		-	3	-	9	
n		20	8		13	9	16	
f		7	3	5		1	8	
s		4	5	2	3		5	
ŋ		15	9	1	7	9	13	
h								

keg condition – illegal outcomes

Coda	Onset: g*	k	m	n	f	s	h	ŋ**
g	-		23	3	4	7		-
k*	-	-	-	-	-	-	-	-
m	-	4	3	4	2	4	4	-
n	-	11	11	3	14	6	10	-
f	-	2	1	2	5	1	2	-
s	-	-	7	5	1	3	-	-
ŋ	-		10	1	1	-		-
h**	-	-	-	-	-	-	-	-

gek condition – legal outcomes

Coda	Onset: g	k	m	n	f	s	h	ŋ
g								
k	5		7	10	13	10	25	
m	2			9	1	1	7	
n	4		3		5	8	12	
f	3		-	5		4	7	
s	1		1	4	1		6	
ŋ	35		6	3	17	12	14	
h								

Table 4: continued

	gek condition – illegal outcomes							
	Onset							
Coda	g	k*	m	n	f	s	h	ŋ**
g*	3	-	1	2	6	1	4	-
k	[redacted]	-	1	5	6	-	[redacted]	-
m	2	-	-	6	-	-	-	-
n	7	-	15	4	3	3	5	-
f	2	-	1	3	3	1	-	-
s	2	-	1	1	-	1	-	-
ŋ	[redacted]	1	6	3	1	-	[redacted]	-
h**	-	-	-	-	-	-	-	-

- * violates experiment-wide constraints.
- ** violates language-wide constraints.

Language-Wide Constraints

In respect of the language-wide constraints, all movements of /h/ and /ŋ/ were phonotactically legal, $p < .008$ (sign-test). These results were not significantly different from those of Experiment 1, $p = 1.00$ (sign-test).

Local Positional Constraints

The local positional constraints, those involving /f/, /s/, /m/, or /n/, were upheld for 43.04% of the 481 errors. This was not significantly different from chance, $p > .289$ (sign-test). Across the participants the range was 34.55 – 51.61%, with only two participants showing a significant difference between legal and illegal movements of these phonemes. These both showed a greater number of illegal than legal errors, $p < .008$(sign-test). None of these results were not significantly different from those of Experiment 1, $p > .727$ (sign-test).

Experiment-Wide Constraints

The experiment-wide constraints were upheld for 88.82% of the 161 errors involving movements of /g/ or /k/ (range: 62.79 – 100%), $p < .008$ (sign-test). This is lower than that found for Experiment one (96.55%). However, there proved to be no significant difference between the rate with which the experiment-wide constraints were upheld in Experiment 1 and Experiment 2, $p = 1.00$ (sign-test).

There was no evidence of participants learning the experiment-wide constraints, as there was no significant difference between the number of illegal experiment-wide errors in the first session and in the final session, $p > .50$ (sign-test). In fact there were a greater number of such errors in the final session than in the first.

When the rates for the experiment-wide errors (88.82%) were compared to those for the language-wide errors (100%), no significant difference was found, $p > .25$ (sign-test). This result is the same as in Experiment 1.

Implicit or Explicit

There was no significant difference between the number of illegal experiment-wide errors produced by the informed and uninformed participants, $p = 1.00$ (sign-test). As in Experiment 1, the participants were asked to note anything they had noticed about the structure of the syllables. Participant 8 (uninformed gek condition) noted that /k/ always came at the end of a syllable. When questioned further, it appears that it was orthographic information that had led participant 8 to notice this. The participant said that *k* was always placed directly after *e*, something that would not happen in English, as there would usually be a *c* between the *e* and the *k* (as in *neck*). However, as there was only one participant who commented on the phonetic constraints, the lack of a statistical difference between the informed and uninformed participants supports the suggestion that the learning is implicit.

Experiment 2: Discussion

As can be seen from the results, Experiment 2 replicated the findings of Experiment 1 in all major respects. The language-wide constraints were upheld on 100% of occasions, which is identical to Experiment 1. The rate with which the local positional constraints were upheld was not significantly different from chance in either experiment. The experiment-wide constraints were upheld at a significant rate in both experiments. And, as in the first experiment, there was no significant difference between the language-wide and experiment-wide constraints. There was support for the implicit nature of the learning, although one participant commented on the experiment-wide constraints. However, as already mentioned, this may have been an artefact of the orthographic nature of the experiment.

From these studies there is evidence that phoneme movements are not necessarily constrained by position, and that it may be no more than chance that produces the position of the erroneous phoneme. It can also be seen that learning does take place, in as much as exposure to the experiment-wide constraints leads the participants to treat them in a similar way to the language-wide constraints.

However, despite the fact that the rate with which the experiment-wide constraints were upheld was lower than for the language-wide effects, there is no evidence that the experiment-wide effects constitute a "middle-ground" of a continuum, as claimed by Dell et al. (2000). It may be, therefore, that the error rates for these two constraints were from the same population.

Experiments 1 & 2: General Discussion

These experiments found partial support for the general findings of Dell et al. (2000), but differed in two respects. Firstly, no evidence was found for the experiment-wide constraints constituting a "middle-ground" of a breadth of constraint continuum, as there was no significant difference between the rate with which the language-wide constraints were upheld and that for the experiment-wide constraints, $p > .25$ (sign-test). Dell et al. (2000) did not perform this analysis, but had they done so, they may have found similar results. And secondly, it was found that rate with which the local positional constraints were upheld was at chance. In this case, there are implications for models of phonological encoding that use separate representations of the same phoneme for different syllabic positions (e.g. Dell et al. 1986, 1988). It is highly unlikely that such models could represent this data, as it is difficult to see how an onset representation of a phoneme could become a coda version of the same phoneme. But, as has already been said, dual-route models that encode phonemes and syllabic position separately (e.g. Hartley & Houghton, 1996), can represent such data. If an error could occur independently, then there is no reason why a phoneme could not be activated at the wrong time, so leading to it occupying the incorrect "slot" in the syllabic structure.

These experiments also showed that the paradigm could be successfully transferred to computer, and that the number of days over which the experiment was conducted could be reduced, whilst still producing comparable results. Therefore, the remaining experiments in this study will be conducted using software to control stimuli presentation, and a shortened version of the Dell et al. (2000) paradigm that requires only one session to complete.

Chapter 4: Experiment 3 - Are Language-wide Constraints Stronger in a Speaker's First- than in His or Her Second-language?

If, as suggested by Dell et al. (2000), it is the speaker's "lifetime" experience that makes the rate with which the language-wide constraints are upheld higher than for other constraints, then it should follow that the language-wide constraints would not be as strong in a speaker's second-language as in their first. Poulisse and van Lieshout (1997) studied the differences between slips of the tongue for first- and second-languages. They found that there were few differences between speech errors in the first- and second-languages of Dutch learners of English. But, interestingly for this study, they found that the errors that did exist decreased over time, with there being a decrease between 9th and 11th grade learners, and a further decrease with 2nd year university students. It would therefore seem that the number of errors produced in a speaker's second-language gradually reduces until the error rate is similar to that of their first-language.

The following experiment explored whether the difference between the speech errors of first- and second-language speakers described by Poulisse and van Lieshout (1997) could also be observed in the different constraints used in the Dell et al. (2000) paradigm. For this purpose, bilingual speakers of English and Welsh were recruited from the participant panel of the School of Psychology, at the University of Wales, Bangor. They were tested using two different stimuli sets. The first used the stimuli set from Experiment 1 (/f/ and /s/ as the experimentally constrained phonemes). The second set was modified so that it followed the phonotactics of Welsh, and contained constraints that were not applicable to English, although the experimentally constrained phonemes were the same. Therefore, it was hypothesised that there would be no significant difference between the rates at which first- and second-language experiment-wide constraints were upheld, as they were not normally constraints of either language, and they were the same for both sets of stimuli. However, the rate at which the local positional constraints and the language-wide constraints were upheld were expected to show a significant difference, as both involved knowledge of the phonotactics of that language.

Experiment 3

Method

Participants

Eight participants were used for this experiment. They were chosen from the participant panel of the School of Psychology, at the University of Wales, Bangor. They were paid £5 for participating in the study. The participants recruited were bilingual in English and Welsh, with four of them being Welsh first-language speakers, and four English first-language speakers. The second-language speakers of English began to learn English between the ages of 3 and 5 years of age (mean = 4 years). The second-language speakers of Welsh began to learn Welsh between the ages of 5 and 25 years of age (mean = 17.8 years). None of them had been diagnosed as dyslexic, and they had normal or corrected to normal vision. None of the participants had taken part in the previous experiments. They were randomly assigned to the conditions, i.e. two in each of the four conditions. The four conditions were: Welsh first-language fes; Welsh first-language sef; English first-language fes; and English first-language sef. Fes and sef refer to the experiment-wide constraints that were applied in each condition. In addition to this, half the participants viewed the Welsh syllables first, and half the English syllables first. This was done to counter-balance for order effects, and was not considered as part of the analysis.

Stimuli and Equipment

In this study both the English and Welsh syllables used the phonemes /f/ and /s/ as the experimentally constrained phonemes, as these phonemes can occur as both an onset and coda in both languages, the only differences being that the phoneme /f/ is graphemically represented as "ff" in Welsh. As in the previous experiments, the language-wide constraints for the English syllables were the phonemes /h/ and /ŋ/. The language-wide constraints for the Welsh syllables were /ɾ̥/ and /ð/ (represented by the graphemes "rh" and "dd" respectively). The phoneme /ɾ̥/ is always an onset in Welsh, but /ð/ occurs as an onset as well as a coda. However, if the occurrence of /ð/ in the word initial position is considered, then it has a much lower frequency. There are two words (*ddim – not*, and *ddoe – yesterday*) that use "dd" as an onset, but all other occurrences would be caused by the system of mutations used in Welsh, which would cause "d" to become "dd". The mutations are triggered by various factors, including the gender of the noun, or the preposition that preceded it. As the grapheme "d" was not used, and the syllables used in the Welsh section of the experiment were nonwords, a mutation would not be triggered. Therefore, the participants would only have experienced two examples of words with "dd" in the initial position. Therefore, it was felt that in the absence of a phoneme that was exclusively a coda, "dd" could be used as a close approximation of a language-wide constraint. In addition to this, the phoneme /k/ is graphemically represented as "c" in Welsh. All other phonemes were the same for both sets of syllables. The syllables for the Welsh sets were generated by replacing the grapheme in the English sets with the alternative Welsh grapheme, i.e. f → ff, h → rh, k → c, and ng → dd. The syllable sets were displayed to the participants in the same way, and using the same equipment, as for Experiment 2. All the experimental blocks were also recorded using the same equipment.

Procedure

The participants were first asked at what age they began to learn their second-language. Each participant then produced 6 sets of 96 sequences, three each of the English and Welsh sets. None of the participants were informed of the constraints placed upon the syllables, but they were told when they would be viewing Welsh or English syllables. Each participant was given a practice block of 22 strings of four syllables (fes or sef, and Welsh or English, depending upon the nature of the first experimental block too which they were assigned) and produced these with the metronome set to a rate of one beat per second. These were not recorded. The participants then produced the six experimental blocks with the metronome set to the rate of 2.53 beats per second. The first three blocks all followed the constraints of either English or Welsh, with this being reversed for the final three blocks.

Experiment 3: Results

Age that Second-Language Acquisition Began

The English second-language speakers began to learn between 3 and 5 years of age ($\overline{X}$ = 4 years). The Welsh second-language speakers began to learn between 5 and 25 years of age ($\overline{X}$ = 17.8 years). This difference was significant, t = 6.77, $p < .00001$.

Total Number of Errors

For the blocks in which participants were speaking their first-language (L1), errors were found on 337 of the 9,216 syllables, giving an error rate of 3.66%. For the blocks in which participants were speaking their second-language (L2), errors were found on 383 of the 9,216 syllables, giving an error rate of 4.16%. There was no significant difference between these error rates, p = 1.00 (sign-test). If the errors are divided into English and Welsh blocks, irrespective of the participants' first-language, then there was a significant difference between the errors produced in each language, with all participants producing more errors in the English than the Welsh blocks. English blocks contained 486 errors, giving an error rate of 5.27%; Welsh blocks contained 234 errors giving an error rate of 2.54%, $p < .008$ (sign-test). As in previous experiments a practice effect was observed, with there being fewer errors in the final session than in the first for both the L1 and L2 blocks. This effect was significant for both L1 and L2, $p < .008$ (sign-test).

Language-Wide Constraints

As can be seen from Tables 5 and 6, all errors obeyed the language-wide constraints for both L1 and L2 blocks, $p < .008$ (sign-test). The difference between the rate at which the language-wide constraints were upheld in L1 and L2 blocks was not significant, $p = 1.00$ (sign-test). Also, neither the L1 nor L2 blocks were significantly different from Experiment 1, $p > .289$ (sign-test).

Table 5: Observed Errors in Experiment 3 Classified by Condition and Legality (L1)

Ffes/fes (L1) condition – legal outcomes								
	Onset							
Coda	g	k	m	n	f	s	ɾ/h	ð/ŋ
g		1	5	3	5		-	
k	-		2	1	1		5	
m	2	2		-	1		-	
n	4	2	2		5		2	
f								
s	-	-	1	2	3		6	
ð/ŋ	4	3	-	3	7		7	
ɾ/h								

Ffes/fes (L1) condition – illegal outcomes								
	Onset							
Coda	g	k	m	n	f	s*	ɾ/h	ð/ŋ**
g	5	3	-	2	2	-	2	-
k	2	4	-	1	3	-	2	-
m	1	1	2	1	2	-	-	-
n	2	1	3	2	2	1	1	-
f*	1	-	-	-	-	-	-	-
s	4	4	-	1		1		-
ð/ŋ	1	2	4	1		-		-
ɾ/h**	-	-	-	-	-	-	-	-

Table 5: continued

Seff/sef (L1) condition – legal outcomes								
	Onset							
Coda	g	k	m	n	f	s	ɾ/h	ð/ŋ
g		4	6	7		10	13	
k	1		2	3		7	4	
m	1	3		-		2	3	
n	3	2	2			2	5	
f	7	8	2	2		10	13	
s								
ð/ŋ	3	2	1	1		8	9	
ɾ/h								

Seff/sef (L1) condition – illegal outcomes								
	Onset							
Coda	g	k	m	n	f*	s	ɾ/h	ð/ŋ**
g	5	1	3	5	-	11	8	-
k	4	2	1	1	-	-	2	-
m	-	4	-	1	-	-	-	-
n	3	2	5	1	-	1	-	-
f	3	2	1	1	-			-
s*	-	-	-	-	-	-	-	-
ð/ŋ	2	5	1	-	-			-
ɾ/h**	-	-	-	-	-	-	-	-

- * violates experiment-wide constraints.
- ** violates language-wide constraints.

Table 6: Observed Errors in Experiment 3 Classified by Condition and Legality (L2)

Ffes/fes condition (L2) – legal outcomes								
	Onset							
Coda	g	k	m	n	f	s	ɾ/h	ð/ŋ
g		4	2	1	2		4	
k	1		-	5	3		4	
m	1	-		2	1		11	
n	6	-	1		2		14	
f								
s	1	-	2	3	4		15	
ð/ŋ	7	5	1	8	1		10	
ɾ/h								

Ffes/fes condition (L2) – illegal outcomes								
	Onset							
Coda	g	k	m	n	f	s*	ɾ/h	ð/ŋ**
g	4	3	2	3	8	-	2	-
k	1	4	-	4	1	-	1	-
m	1	-	2	-	1	-	-	-
n	5	-	3	3	1	-	4	-
f*	-	-	-	-	-	-	-	-
s	3	3	-	5		-		-
ð/ŋ	2	2	1	9		-		-
ɾ/h**	-	-	-	-	-	-	-	-

Table 6: continued

Seff/sef (L2) condition – legal outcomes								
	Onset							
Coda	g	k	m	n	f	s	ɾ/h	ð/ŋ
g		1	1	3		-	5	
k	-		-	1		3	1	
m	1	2		4		3	2	
n	6	3	3			2	5	
f	2	6	-	2		6	3	
s								
ð/ŋ	15	3	2	1		10	6	
ɾ/h								

Seff/sef (L2)condition – illegal outcomes								
	Onset							
Coda	g	k	m	n	f*	s	ɾ/h	ð/ŋ**
g	3	4	2	4	-	4	4	-
k	-	3	2	2	-	3	2	-
m	2	1	-	-	-	1	-	-
n	6	2	5	1	-	2	2	-
f	5	4	-	2	-			-
s*	-	-	-	-	-	-	-	-
ð/ŋ	2	4	2	-	-			-
ɾ/h**	-	-	-	-	-	-	-	-

- * violates experiment-wide constraints.
- ** violates language-wide constraints.

Local Positional Constraints

The rate at which the local positional constraints were upheld was 43.55% for L1, and 46.82% for L2. Neither of these rates were significantly different from chance, $p = 1.00$ for both L1 and L2 (sign-test). The range across participants was 14.71 - 60.00% for L1 and 30.30 - 51.28% for L2. There was no significant difference between the local positional constraints for L1 and L2, $p > .289$ (sign-test).

Experiment-Wide Constraints

The rate at which the experiment-wide constraints were upheld was 95.08% for L1, $p < .031$ (sign-test); and for L2 the rate was 100%, $p < .008$ (sign-test). There was no significant difference between the rate at which the experiment-wide constraints were upheld for L1 and L2, $p > .727$ (sign-test). The range across participants was 50.00 - 100% for L1 and 100% for all participants in L2. All the errors that broke the experiment-wide constraints, a total of three, were made in the L1 blocks; two by a Welsh and one by an English speaker. There was no significant difference between the language-wide and experiment-wide constraints, $p > .289$ (sign-test).

Differences from Experiment 1

None of the results from this study were significantly different from those obtained from Experiment 1 (which used the same experiment-wide constraints for the English syllables) with, in all cases, $p > .289$ (sign-test).

Experiment 3: Discussion

To begin by examining the elements of this experiment that were present in the previous studies, it can be seen that there were no significant differences between either the English or Welsh blocks of the experiment and the results from Experiment 1. Therefore, it shows that the English blocks of the experiment produced the same results, and that the inclusion of Welsh phonemic constraints produced similar results to those found using the English constraints. A comparison between first- and second-languages can therefore legitimately be made.

When looking at the total number of errors made, it is interesting to note that all participants, irrespective of their first-language, made significantly more errors in the English blocks than in the Welsh. A possible explanation for this is that the pronunciation of Welsh is more predictable than English, and so any "uncertainty" concerning the pronunciation of the English nonsense syllables may have provided more error opportunities.

With regards to the local positional constraints, no significant differences could be found between first- and second-languages. However, the previous two studies have shown that these may not be constraints at all, but merely the result of chance (indicated by the 50-50 distribution of legal and illegal errors).

Turning to the experiment-wide constraints, no significant difference was expected, or found, due to the constraints being artificially imposed and unrelated to either first- or second-language.

When examining the hypotheses that participants would make significantly more errors involving local positional or language-wide constraints in their second-language than in their first, due to a more limited experience with the second-language, then no such differences can be found. However, as it has been seen in the previous two studies, errors that break the language-wide constraints are extremely rare, and so a larger sample may be needed in order to gather appropriate data.

A possible weakness in the experiment may lie in the ages at which the participants began to learn their second-language. The Welsh first-language speakers began to learn English at an early age and grew up in a bilingual society, and could therefore be said to have a similar amount of experience in both English and Welsh. This would negate any differences that may have been present between their first- and second-languages. The English first-language speakers did not begin to learn Welsh until much later (typically when they attended high school or university), and so only had a few years experience. The difference between the age at which acquisition began in English and Welsh L2 speakers was significant. However, the findings of Poulisse and van Lieshout (1997) may also point to these problems, as they found that any differences between first- and second-language speech errors reduced between the 9th grade of school and the 2nd year of university (age range of 13 to 19 years), and so the differences in the participants used for this study may already have diminished to a large extent.

It is felt that if a difference is to be found between the rates at which the language-wide constraints are upheld in a speakers first- and second-language, then it would be necessary to find participants who did not grow up in a bilingual society and who learned there second-language and a similar age (e.g. they all began to learn the second-language at high-school). Therefore, no conclusions can be made about the statement by Dell et al. (2000) - that the strength of the language-wide constraints are due to the speaker's life-time experience of their first-language – due to the inherent problems with participant selection in this experiment.

Chapter 5: Experiments 4 to 6 - Acquisition of Phonemic Constraints; Time-course and Durability.

Studies have been conducted that look at the acquisition of novel phonemic constraints in adults (e.g. Dell et al., 2000), and further studies have looked at this phenomenon in children (e.g. Chambers, Onishi, and Fisher, 2003; Saffran & Thiessen, 2003). However, it appears that no studies have looked at the robustness of this learning, and the time-course over which the learning is then lost. It would seem that the learning is not maintained after the end of the experiment, as the participants in the present study did not continue to use the experiment-wide constraints, possibly demonstrating that the learning is specific to the experimental setting. However, it is not known what happens when the participants are presented with new constraints that conflict with those previously learnt, e.g. if a phoneme that had previously always been an onset, was now always a coda.

In order to explore this, the participants were first trained using the technique described in Experiment 2, so that they had learnt the experimental constraints. They were then given a fourth block in which the experimental constraints were reversed, e.g. if the constraints had been /g/ as an onset and /k/ as a coda, then /k/ would now be the onset and /g/ would be the coda. It was hypothesised that the participants would have difficulty adapting to the new constraints and would continue to use the original constraints for some time. The length of time that this confusion continued would give an indication of the robustness and time-course of the learning. The exploration of the explicit-implicit nature of the learning was removed from Experiments 4, 5, and 6, as there were doubts about the efficacy of the method of measurement.

Experiment 4

Experiment 4 used the technique described above to explore the time-course and robustness of the learning involved in the acquisition of the phonotactic constraints. In this experiment the constraints were reversed at the beginning of the fourth block.

Method

Participants

Eight participants were used for this experiment, and were recruited on the basis that English was their native language. They were all students from the School of Psychology at the University of Wales, Bangor, and were given course credits for participation. None of the participants had been diagnosed as dyslexic, or had participated in the previous studies, and they all had normal or corrected to normal vision. They were randomly assigned to the two experimental conditions; initial constraint keg and initial constraint gek. Participants were not informed of the experimental constraints, as knowledge of these may have caused confusion when the constraints were reversed, that was not due to the learning process.

Stimuli and Equipment

In this study, /g/ and /k/ were the experimentally constrained phonemes. For participants in the gek condition, /g/ was always an onset and /k/ was always a coda. For participants in the keg condition the reverse was true. The consonants /h/ and /ŋ/ were the language-wide constraints, and the consonants /m/, /n/, k, and /g/, each occurred once per sequence and were unconstrained. Each syllable had /ɛ/ as its vowel. The spelling of the syllables was straightforward except that /g/ when used as an onset was spelt "gh", in order to remove any hard-soft ambiguity associated with its pronunciation. The order of the syllables within the sequences was randomised for each sequence and each participant.

This was the same set of stimuli that was used for Experiment 2, but the order of presentation differed in the following way. The practice session (produced at the rate of 1 beat per second) and the first three experimental blocks (produced at a rate of 2.53 beats per second) contained the same experiment-wide constraints (i.e. /k/ onset and /g/ coda, or vice-versa). The fourth, extra, experimental block reversed the experiment-wide constraints so that the coda became the onset and the onset became the coda. The blocks were identical in every other respect. The order of presentation was randomised within each block and for each participant. The equipment and software were the same as for Experiment 2.

Procedure

With the exception that no participants were informed of the constraints, the instructions given to the participants were identical to those for Experiments 1 and 2. Each participant was given a practice session of 22 strings of four syllables (keg or gek depending on the block to which they were assigned) and produced these with the metronome set to a rate of one beat per second. These were not recorded. The participants then produced four blocks of 96 sequences at the faster rate of 2.53 beats per second. The first three blocks all followed the same experiment-wide constraints (dependent upon the assigned condition). The fourth block reversed the experiment-wide constraints. All four blocks were recorded using the equipment described in Experiment 1.

Experiment 4: Results

Total Number of Errors

Considering the data in its entirety, errors were found on 383 of the 12,288 syllables produced by the participants. This gives an overall error rate of 3.12%, which is comparable to Experiment 1 and 2, $p = 1.00$ (sign-test). There was no evidence of a practice effect, as there was no significant difference between the number of errors in the first and final blocks of the initial constraints, $p = 1.00$ (sign-test).

Language-wide Constraints

As can be seen in Tables 7 and 8, all errors obeyed the language-wide constraints, $p < .008$ (sign-test). This is identical to Experiments 1 and 2.

Table 7: Observed Errors in Experiment 4 Classified by Condition and Legality (gek first condition).

gek condition – legal outcomes								
	Onset							
Coda	g	k	m	n	f	s	h	ŋ
g								
k	2		2	1	4	2	5	
m	1			2	-	1	1	
n	12		6		8	5	6	
f	4		-	4		3	-	
s	2		4	1	2		4	
ŋ	8		-	-	3	3	5	
h								

gek condition – illegal outcomes								
	Onset							
Coda	g	k*	m	n	f	s	h	ŋ**
g*	-	-	-	1	1	-	3	-
k		1	1	-	1	-		-
m	-	-	1	2	-	-	-	-
n	8	-	11	9	-	1	9	-
f	2	-	2	4	5	1	2	-
s	-	-	-	1	2	1	2	-
ŋ		-	2	6	-	-		-
h**	-	-	-	-	-	-	-	-

keg condition – legal outcomes								
	Onset							
Coda	g	k	m	n	f	s	h	ŋ
g		-	1	1	1	2	6	
k								
m		1		-	-	-	1	
n		-	2		-	2	3	
f		2	1	4		2	1	
s		-	1	1	2		-	
ŋ		3	1	-	1	-	6	
h								

Table 7: continued

keg condition – illegal outcomes								
	Onset							
Coda	g*	k	m	n	f	s	h	ŋ**
g	-		1	-	-	1		-
k*	-	-	-	1	1	-	-	-
m	-	-	-	-	-	-	-	-
n	-	2	4	2	-	1	-	-
f	-	-	-	2	1	-	-	-
s	-	-	-	-	-	-	-	-
ŋ	-		1	-	-	-		-
h**	-	-	-	-	-	-	-	-

- * violates experiment-wide constraints.
- ** violates language-wide constraints.

Table 8: Observed Errors in Experiment 4 Classified by Condition and Legality (keg first condition).

keg condition – legal outcomes								
	Onset							
Coda	g	k	m	n	f	s	h	ŋ
g		9	2	3	1	2	5	
k								
m		3		1	-	-	-	
n		4	5		-	3	8	
f		3	-	2		-	-	
s		-	-	-	1		2	
ŋ		10	2	2	4	-	9	
h								

keg condition – illegal outcomes								
	Onset							
Coda	g*	k	m	n	f	s	h	ŋ**
g	-		3	-	1	-		-
k*	-	1	-	1	2	-	-	-
m	-	-	-	2	-	-	-	-
n	-	3	5	3	2	-	1	-
f	-	1	-	-	3	-	-	-
s	-	-	-	1	2	-	-	-
ŋ	-		2	-	-	-		-
h**	-	-	-	-	-	-	-	-

gek condition – legal outcomes								
	Onset							
Coda	g	k	m	n	f	s	h	ŋ
g								
k	1		1	1	1	-	-	
m	-			-	1	-	-	
n	2		1		-	-	1	
f	-		-	-		1	-	
s	-		-	-	-		-	
ŋ	1		1	-	-	2	1	
h								

Table 8: continued

	gek condition – illegal outcomes							
	Onset							
Coda	g	k*	m	n	f	s	h	ŋ**
g*	-	-	-	-	-	-	2	-
k		1	-	-	-	-		-
m	-	1	-	1	-	-	-	-
n	-	1	-	-	-	-	-	-
f	-	-	-	-	-	-	-	-
s	1	2	-	-	-	-	-	-
ŋ		-	1	-	-	1		-
h**	-	-	-	-	-	-	-	-

- *violates experiment-wide constraints.
- ** violates language-wide constraints.

Local Positional Constraints

The local positional constraints were upheld 46.55%. This was not significantly different from chance, $p > .727$ (sign-test). This is also comparable to Experiments 1 and 2, $p = 1.00$ (sign-test). The range across participants for the local positional constraints was 30.88 – 60.00%.

Experiment-wide Constraints

If I consider the first three blocks of the experiment (those in which the constraints had been consistent), then the experiment-wide constraints were upheld 80.77%, which was not significant, $p > .125$ (sign-test). However, it is comparable to Experiment 2 (which had used the same experimental constraints), $p > .289$ (sign-test). The range across participants was 50.00 – 100%.

If the first block of the original constraints is compared to the final block (in which the constraints were reversed), then it can be seen that in the first block the experiment-wide constraints were upheld 71.43%, $p > .289$ (sign-test), and in the final block the rate was 60.87%, $p > .688$ (sign-test). There was no significant difference between these two rates, $p = 1.00$ (sign-test). The range across participants for the original block was 28.57 - 100%. The range was identical for the reversed block. The outliers (28.57% in both cases) were not removed for the reasons given in Chapter 2. There were a total of nine errors in the reversed block that broke the experiment-wide constraints, compared to eight in the original block. The distribution of these errors can be seen in Figures 8 and 9.

There was no significant difference between the rate at which the experiment-wide constraints were upheld and that for the language-wide constraints, $p > .289$ (sign-test).

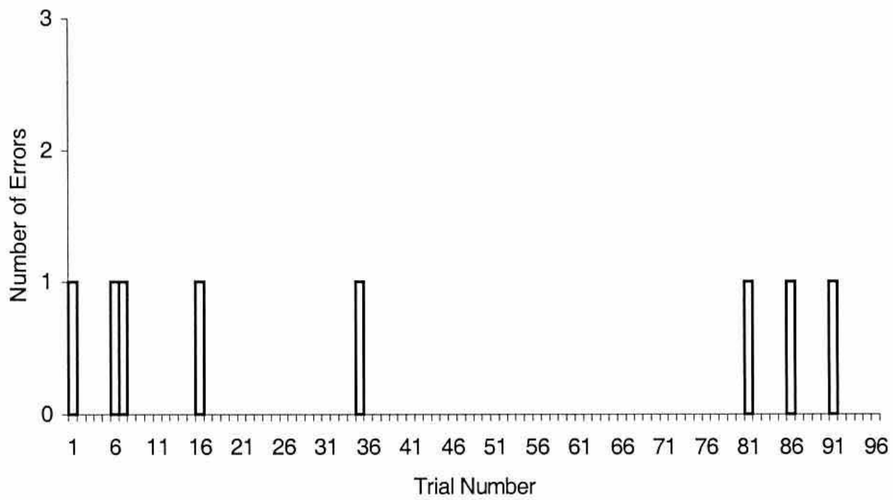

Figure 8: Distribution of Illegal Experiment-wide Errors in First Block of Experiment 4.

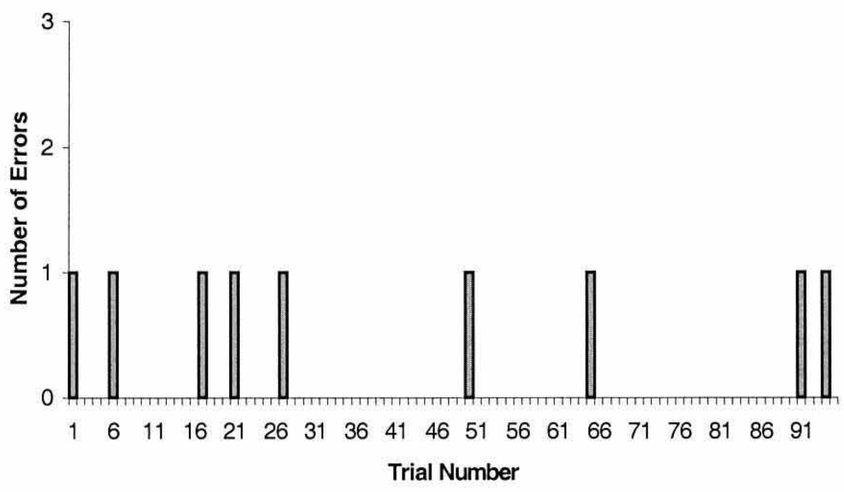

Figure 9: Distribution of Illegal Experiment-wide Errors in "Reversed" Block of Experiment 4.

Differences from Experiment 1 & 2

There were no significant differences between this experiment and Experiments 1 and 2, $p > .289$ (sign-test).

Experiment 4: Discussion

Experiment 4 replicated the findings of Experiments 1 and 2, with the exception that no practice effect was observed, but failed to find support for the hypothesis that the reversal of the constraints would cause difficulty for the participants that resulted in their continuing to use the original constraints. The distribution of illegal errors within the blocks also showed no obvious differences (see Figures 8 and 9), with errors occurring throughout most of each block.

There are two possible reasons why the hypothesis was not supported. Firstly, the low error rate in this experiment (3.12%), combined with the fact that only two individual blocks were being analysed, led to a paucity of data. For this reason it was decided to increase the number of participants to 16 in future studies. Secondly, there was a graphemic cue for the change of constraints for participants in the keg/gek condition. Participants in this condition began by seeing the set of syllables in which "k" was always an onset and "g" was always a coda. When the constraints were reversed "k" was now a coda, but the "g" onset was now spelt "gh" in order to avoid the hard/soft ambiguity which this grapheme may otherwise have caused. This acted as a visual cue for the change in constraints, and this was commented upon by two participants. To avoid this visual cue, the following experiments used the phonemes /f/ and /s/ as the experimental constraints as these have no ambiguities, and so would not provide graphemic cues.

Experiment 5

In the light of the possible difficulties encountered in Experiment 4, the follow-up study increased the number of participants to 16 in order to increase the data corpus, and reverted to the syllable set used in Experiment 1, i.e. using /s/ and /f/ as the experiment-wide constraints, in order to remove the apparent graphemic cues. Other than these changes, the experiment and hypotheses were unchanged, i.e. it was hypothesised that the participants would have difficulty adapting to the new constraints, and would continue to use the original constraints for some time. The length of time that this confusion continued would give an indication of the robustness and time-course of the learning.

Experiment 5: Method

Participants

Sixteen participants were used for this experiment. They were all students from the School of Psychology at the University of Wales, Bangor, and were given course credits for participation. None of the participants had been diagnosed as dyslexic, they all had normal or corrected to normal vision, and none of them had participated in the previous experiments. They all had English as their native language. They were randomly assigned to the two experimental conditions; initial constraint fes, and initial constraint sef. Participants were not informed of the experimental constraints as knowledge of these may have caused confusion when the constraints were reversed that was not due to the learning process.

Stimuli and Equipment

In this study, /f/ and /s/ were the experimentally constrained phonemes. For participants in the fes condition, /f/ was always an onset and /s/ was always a coda. For participants in the sef condition the reverse was true. The consonants /h/ and /ŋ/ were the language-wide constraints, and the consonants /m/, /n/, k, and /g/, each occurred once per sequence and were unconstrained. Each syllable had /ɛ/ as its vowel. The spelling of the syllables was straightforward except that /g/ when used as an onset was spelt "gh", in order to remove any hard-soft ambiguity associated with its pronunciation. The order of the syllables within the sequences was randomised for each sequence and each participant.

This was the same set of stimuli as used in Experiment 1, but the order of presentation differed in the following way. The practice session (produced at the rate of 1 beat per second) and the first three experimental blocks (produced at a rate of 2.53 beats per second) contained the same experiment-wide constraints (i.e. /f/ onset and /s/ coda, or vice-versa). The fourth, extra, experimental block reversed the experiment-wide constraints, so that the coda became the onset and the onset became the coda. The blocks were identical in every other respect. The equipment and software were the same as that used for Experiment 2 onwards.

Procedure

With the exception that no participants were informed of the constraints, the instructions given to the participants were identical to those for the previous experiments. Each participant was given a practice session of 22 strings of four syllables (fes or sef depending on the block to which they were assigned) and produced these with the metronome set to a rate of one beat per second. These were not recorded. The participants then produced four blocks of 96 sequences at the faster rate of 2.53 beats per minute. The first three blocks all followed the same experiment-wide constraints (dependent upon the assigned condition). The fourth block reversed these constraints. All four experimental blocks were recorded using the equipment described in Experiment 1.

Experiment 5: Results

Total Number of Errors

Considering the data in its entirety, errors were found on 1,795 of the 24,576 syllables produced by the participants. This gives an overall error rate of 7.88%, which is not significantly different from that for Experiment 1 and 2, $p >$ .07 (sign-test). There was no evidence of a practice effect, as there was no significant difference between the number of errors in the first and final blocks, p > .077 (sign-test).

Language-Wide Constraints

As can be seen in Tables 9 and 10, all errors obeyed the language-wide constraints, $p < .001$ (sign-test).

Table 9: Observed Errors in Experiment 5 Classified by Condition and Legality (fes first condition)

fes condition – legal outcomes								
	Onset							
Coda	g	k	m	n	f	s	h	ŋ
g		10	10	10	13		13	
k	1		5	6	12		11	
m	4	15		5	7		10	
n	9	18	5		16		19	
f								
s	3	10	2	6	13		17	
ŋ	24	28	9	3	23		21	
h								

Table 9: continued

fes condition – illegal outcomes								
	Onset							
Coda	g	k	m	n	f	s*	h	ŋ**
g	15	8	6	9	13	-	20	-
k	3	7	3	7	1	-	3	-
m	-	5	2	1	1	-	1	-
n	12	11	5	8	7	-	11	-
f*	-	-	-	-	-	-	-	-
s	5	9	4	-		2		-
ŋ	9	9	4	1		-		-
h**	-	-	-	-	-	-	-	-

sef condition – legal outcomes								
	Onset							
Coda	g	k	m	n	f	s	h	ŋ
g		4	7	2		4	7	
k	-		4	4		6	7	
m	1	5		2		1	1	
n	8	5	5			2	2	
f	2	2	3	2		11	4	
s								
ŋ	5	3	1	3		4	5	
h								

sef condition – illegal outcomes								
	Onset							
Coda	g	k	m	n	f*	s	h	ŋ**
g	5	1	1	1	-	1	2	-
k	-	2	-	1	-	3	2	-
m	-	-	-	1	1	-	-	-
n	2	3	1	-	-	3	1	-
f	2	-	-	1	-			-
s*	-	1	-	-	-	-	-	-
ŋ	4	-	-	-	-			-
h**	-	-	-	-	-	-	-	-

- * violates experiment-wide constraints.
- ** violates language-wide constraints.

Table 10: Observed Errors in Experiment 5 Classified by Condition and Legality (sef first condition).

sef condition – legal outcomes								
	Onset							
Coda	g	k	m	n	f	s	h	ŋ
g		17	23	26		31	26	
k	5		7	18		17	19	
m	6	5		4		4	8	
n	8	18	5			24	13	
f	14	19	5	19		34	37	
s								
ŋ	34	20	11	4		22	23	
h								

Table 10: continued

sef condition – illegal outcomes								
	Onset							
Coda	g	k	m	n	f*	s	h	ŋ**
g	13	12	11	10	-	19	31	-
k	10	16	5	8	-	5	4	-
m	-	6	4	4	-	1	1	-
n	15	15	15	14	-	10	13	-
f	7	10	5	9	-			-
s*	-	1	-	-	-	-	-	-
ŋ	14	15	2	1	-			-
h**	-	-	-	-	-	-	-	-

fes condition – legal outcomes								
	Onset							
Coda	g	k	m	n	f	s	h	ŋ
g		3	2	8	6		11	
k	2		1	2	7		4	
m	-	6		2	1		5	
n	5	5	9		11		5	
f								
s	-	4	5	4	5		5	
ŋ	10	5	4	2	5		4	
h								

fes condition – illegal outcomes								
	Onset							
Coda	g	k	m	n	f	s*	h	ŋ**
g	8	5	3	4	5	-	4	-
k	3	4	1	2	4	-	-	-
m	2	-	1	2	-	-	-	-
n	7	6	8	6	6	-	11	-
f*	-	-	-	-	-	-	-	-
s	3	1	1	-		-		-
ŋ	-	1	-	1		-		-
h**	-	-	-	-	-	-	-	-

- * violates experiment-wide constraints.
- ** violates language-wide constraints.

Local Positional Constraints

The local positional constraints were upheld 51.63%. This was not significantly different from chance, $p > .804$ (sign-test). This result is also comparable with Experiments 1 and 2, $p = 1.00$ (sign-test). The range across participants for the local positional constraint was 22.22% - 65.12%.

Experiment-Wide Constraints

The experiment-wide constraints were upheld at a rate similar to those for Experiment 1(which used the same experimental constraints), $p = 1.00$ (sign-test). In the initial blocks, the constraints were upheld 98.28%, $p < .001$ (sign-test). There was no evidence of the participants learning the constraints, as there was no significant difference between the number of illegal experiment-wide errors made in the first and last blocks of the "normal" constraints, $p > .625$ (sign-test). As in the previous experiments, the rate at which the experiment-wide constraints were upheld was not significantly different from that for the language-wide constraints, $p > .250$ (sign-test).

There was very little difference between the rate that the experiment-wide constraints were upheld in the final "reversed" block, which resulted in a rate of 96.43%, $p < .001$ (sign-test), and that in the first of the "original" blocks that produced an error rate of 97.87%, $p < .001$ (sign-test). There was no significant difference between the rates at which the constraints were upheld in these two blocks, $p > .625$ (sign-test). There were two illegal experiment-wide errors in both the "original" and the "reversed" blocks (errors occurring on trial 13 and 42 in the "original" block, and 23 and 67 in the "reversed").

Experiment 5: Discussion

The first three blocks of this experiment supported the findings of the first two experiments. The language-wide constraints were upheld 100% and the experiment-wide constraints 98.28% of the time. And, yet again, there was no significant difference between these two figures, providing further evidence that these two constraints may not be different parts of the same continuum as claimed by Dell et al. (2000). The local positional constraints were also upheld on approximately 50% of occasions, and this was not significantly different from either of the previous experiments.

However, the fourth "reversed" block of the experiment, showed no significant differences from the first block, and the participants did not show any initial difficulty when adapting to the reversed phonemes, nor did the errors conform to any pattern. This "instantaneous" switch to a new set of constraints would seem "impossible" at an intuitive level. It is therefore hypothesised that the participants were treating the fourth block as a completely new condition, and were not applying the rules acquired during the previous block. This would seem logical, as the first block of the experiment that they were exposed to showed no initial difficulty, and the errors were spread "randomly" throughout the block. The effectiveness of using separate blocks for this paradigm must, therefore, also be questioned. If the participants are treating each block as a separate case, then no learning of the experiment-wide constraints will be observed across the blocks. Further more, spreading the blocks over four days would reduce any effect even further. And, this is indeed what has been observed, with none of the experiments so far conducted showing any significant learning across the blocks/days for the experiment-wide constraints.

In order to further test this hypothesis, the next experiment reversed the constraints at the midpoint in the fourth block (i.e. after 48 of the 96 trials), so that participants would not have had the rest period between blocks and would, therefore, not treat the change as a separate condition. The reversed section could then be compared to the initial 48 trials of the first block.

Experiment 6

For Experiment 6 the constraints were changed at the mid-point of the fourth block. It was hypothesised that in this situation participants would show signs of confusion when the constraints were reversed, by continuing to use the original constraints, and so the rate at which the constraints were upheld would temporarily drop, until the participants had learnt the new constraints. In order to examine this, the reversed block would be compared with the equivalent section of the first block (i.e. the first 48 trials of the experiment, where the participants would have had no experience of the experimental constraints).

Experiment 6: Method

Participants

Sixteen participants who had English as their native language were used for this experiment. They were all students from the School of Psychology at the University of Wales, Bangor, and were given course credits for participation. None of the participants had been diagnosed as dyslexic, they all had normal or corrected to normal vision, and none of them had participated in the previous experiments. They were randomly assigned to the two conditions: initial constraint fes; and initial constraint sef. As in the previous experiment, participants were not informed of the experiment-wide constraints prior to completing the experiment.

Stimuli and Equipment

The stimuli were the same as those used for Experiment 5, but the order of presentation differed in the following way. The practice session and the first three experimental blocks contained the same experiment-wide constraints (i.e. /f/ onset and /s/ coda, or vice-versa). The fourth experimental block began with the constraints continuing as in the first three blocks, but after 48 trials the constraints were reversed (i.e. /f/ onset /s/ coda was changed to /s/ onset and /f/ coda, and vice-versa). The order of presentation was randomised within and between each block. The equipment and software used was the same as for the previous experiment.

Procedure

With the exception that no participants were informed of the constraints, the instructions given to the participants were identical to those for the previous experiments. Each participant was given a practice session of 22 strings of four syllables (fes or sef depending on the block to which they were assigned), and produced these with the metronome set to a rate of one beat per second. These were not recorded. The participants then produced four blocks of 96 sequences at the faster rate of 2.53 beats per minute. The first three blocks all followed the same experiment-wide constraints (dependent upon the assigned condition). The fourth block continued with the same constraints as the first three blocks for the first 48 trials, and then reversed the constraints for the remaining trials. All four experimental blocks were recorded using the equipment described earlier.

Experiment 6: Results

Total Number of Errors

Considering the data in its entirety, errors were found on 2,248 of the 24,576 syllables produced by the participants. This gives an error rate of 9.15%. This is not significantly different from Experiments 1, 2 and 5, $p > .077$ (sign-test). However, there was evidence of a practice effect, as there were significantly fewer errors in the final block than in the first, $p < .035$ (sign-test).

Language-Wide Constraints

Of the 544 errors involving movement of /h/ and /ŋ/, 543 obeyed the language-wide constraint. Therefore, the language-wide constraints were upheld on 99.82% of occasions, $p < .001$ (sign-test). The errors observed in Experiment 6 can be seen in Tables 11 and 12. As can be seen in Table 11, the error that broke the language-wide constraints involved the movement of /ŋ/ to the onset position. This resulted in the production of the cut-off error /ŋ…ɛŋ/.

Table 11: Observed Errors in Experiment 6 Classified by Condition and Legality (fes first condition)

fes condition – legal outcomes

Coda	Onset: g	k	m	n	f	s	h	ŋ
g		20	28	30	20		52	
k	15		13	10	19		20	
m	4	8		9	11		17	
n	8	18	17		32		31	
f								
s	19	14	22	23	34		47	
ŋ	78	45	23	18	66		63	
h								

fes condition – illegal outcomes

Coda	Onset: g	k	m	n	f	s*	h	ŋ**
g	17	16	25	28	23	-	38	-
k	13	12	3	10	14	-	4	-
m	7	4	2	8	3	-	8	-
n	5	9	24	18	24	-	20	-
f*	1	1	-	1	1	-	-	-
s	7	9	9	12		-		-
ŋ	16	19	8	10		-		1
h**	-	-	-	-	-	-	-	-

sef condition – legal outcomes

Coda	Onset: g	k	m	n	f	s	h	ŋ
g		1	5	1		-	3	
k	2		1	1		1	1	
m	1	3		-		1	3	
n	2	2	-			1	1	
f	5	5	1	2		2	2	
s								
ŋ	13	3	1	1		5	10	
h								

sef condition – illegal outcomes

Coda	Onset: g	k	m	n	f*	s	h	ŋ**
g	1	2	1	2	1	2	1	-
k	1	-	-	1	2	1	4	-
m	-	1	-	-	-	-	-	-
n	1	3	3	2	2	2	4	-
f	1	2	-	1	-			-
s*	3	-	-	-	2	-	-	-
ŋ	2	-	-	2	-			-
h**	-	-	-	-	-	-	-	-

- *violates experiment-wide constraints.
- ** violates language-wide constraints.

Table 12: Observed Errors in Experiment 6 Classified by Condition and Legality (sef first condition)

sef condition – legal outcomes

Coda \ Onset	g	k	m	n	f	s	h	ŋ
g		13	22	28		17	33	
k	11		7	11		14	13	
m	3	6		6		4	16	
n	7	18	6			9	14	
f	18	12	9	23		44	41	
s								
ŋ	30	22	2	8		26	26	
h								

sef condition – illegal outcomes

Coda \ Onset	g	k	m	n	f*	s	h	ŋ**
g	12	13	12	26	-	20	32	-
k	2	16	4	1	-	5	2	-
m	-	4	1	3	-	1	3	-
n	13	15	12	10	2	10	11	-
f	7	8	2	5	-			-
s*	-	-	-	-	-	-	1	-
ŋ	9	4	2	4	-			-
h**	-	-	-	-	-	-	-	-

fes condition – legal outcomes

Coda \ Onset	g	k	m	n	f	s	h	ŋ
g		-	-	-	4		2	
k	-		-	1	-		1	
m	-	-		1	1		3	
n	1	2	1		2		2	
f								
s	1	-	2	2	1		4	
ŋ	3	-	2	2	1		5	
h								

fes condition – illegal outcomes

Coda \ Onset	g	k	m	n	f	s*	h	ŋ**
g	1	3	1	3	2	-	1	-
k	-	2	-	-	-	1	1	-
m	-	1	1	-	-	-	-	-
n	1	3	6	-	1	-	-	-
f*	-	-	-	-	-	-	-	-
s	-	-	-	-		-		-
ŋ	-	-	-	-		-		-
h**	-	-	-	-	-	-	-	-

- * violates experiment-wide constraints.
- ** violates language-wide constraints.

Local Positional Constraints

The rate at which the local positional constraints were upheld was 47.11%. This was not significantly different from chance, $p > .302$ (sign-test). The range across participants for the local positional constraints was 28.13 – 56.84%. This is not significantly different from Experiments 1, 2 and 5, $p > .077$ (sign-test).

Experiment-Wide Constraints

If the first part of the experiment is examined (that in which the constraints had been consistent), then the rate at which the constraints were upheld was 97.72%, $p < .001$ (sign-test). The range across participants was 81.25 – 100%. In the final section (in which the constraints were reversed) the rate at which the constraints were upheld was 63.89, $p > .125$ (sign-test). The range across participants in the "reversed" block was 20.00 – 100%. If the "reversed" section of the final block is compared to the equivalent section of the first block, then it can be seen that the constraints were upheld at a far higher rate (95.83%), $p < .001$ (sign-test). However, the difference between these two rates was not significant, $p > .125$ (sign-test). There were a total of thirteen errors in the final "reversed" block that broke the experiment-wide constraints, compared to two such errors in the first 48 trials of the first block. Figures 10 and 11 show the distribution of these errors throughout the two sections.

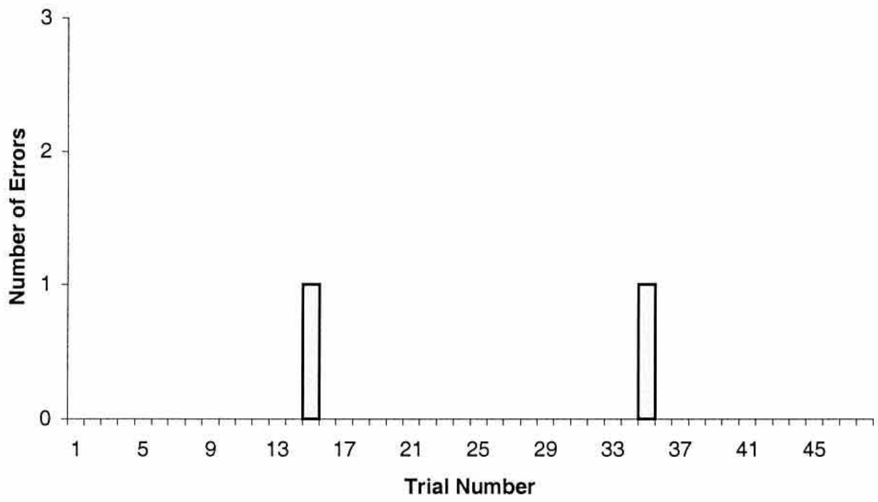

Figure 10: Distribution of Illegal Experiment-wide Errors in Initial Section of Experiment 6.

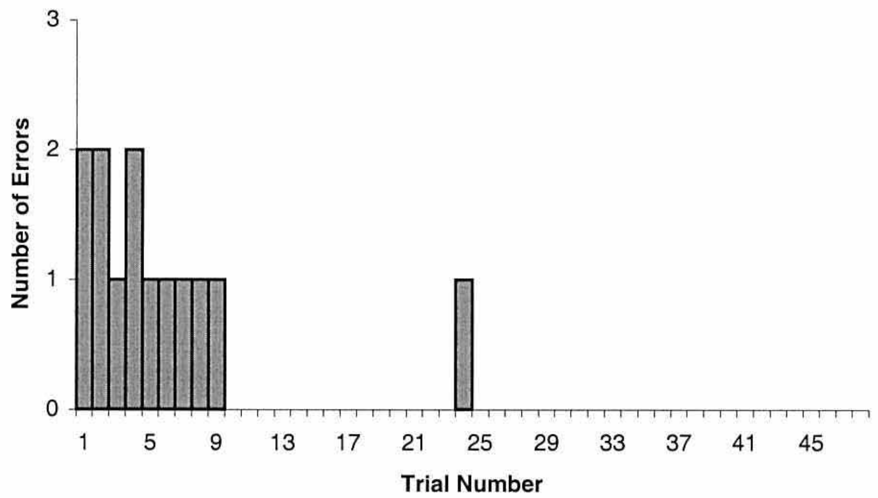

Figure 11: Distribution of Illegal Experiment-wide Errors in "Reversed" Section of Experiment 6.

There was a significant correlation between the number of illegal experiment-wide errors made throughout all the normal trials and the number made in the reversed trials, $r = .480$, $p < .030$ (Spearman's Rho one-tailed). Table 13 shows the distribution of these errors throughout the experiment.

Table 13. Distribution of Illegal Experiment-wide Errors throughout Experiment 6.

Participant	"Normal" Trials	"Reversed" Trials
3	78	1, 2, 4, 6
7	35, 179	1-5, 7-9
8	249	-
14	15, 285, 286	-
15	-	24

* Numbers relate to the number of the trial on which the error occurred.
** Participants not listed made no illegal errors in either condition.

The rate at which the experiment-wide constraints were upheld was not significantly different from that for the language-wide constraints, $p > .375$ (sign-test). There was no evidence or learning with there being no significant decrease in the number of illegal experiment-wide errors between the first and last blocks of the initial "normal" constraints, $p > .625$ (sign-test).

Differences between Experiments

With the exception of the effects caused by the reversal of the constraints in the final experiment, there were no significant differences between any of the three experiments conducted in this study, $p > .05$ (sign-test).

Experiment 6: Discussion

The results of this study provide general support for the findings of all the previous studies, with the exception that there was evidence of learning, i.e. there were fewer illegal experiment-wide errors in the last of the "normal" blocks than in the first. This casts doubt upon the idea that individual blocks were treated separately by the participants for the purpose of applying constraints.

However, moving the point at which the constraints were reversed did have an effect in line with that predicted in the hypothesis (see below). Therefore, it may not be the case that participants treat each block separately, but that it is only the suddenness of the change that causes the difficulty in adapting to the new constraints.

The violation of the language-wide constraints by one participant, cast further doubt upon the claim by Dell et al. (2000) that the experiment-wide constraints constituted the middle ground of a breadth of constraint continuum, as this brings the two levels of constraint closer together. The error involved the production of /ŋ…ɛŋ/ (which was independently agreed upon by both coders). This would appear to be an attempt to begin a syllable with /ŋ/ that was partially corrected by producing the correct vowel and coda of the target syllable. This error shows that language-wide constraints for an individual's first language can be violated.

The main results of this experiment relate to the number and spacing of the illegal experiment-wide errors. There were thirteen errors in the "reversed" section that broke the experiment-wide constraints, and twelve of these occurred within the first nine trials. This is compared with two such errors throughout the first section of the "normal" trials, and these were randomly spread across all the blocks. During the nine trial period in which some participants experienced difficulty, the rate at which one participant upheld the constraints dropped to 11.01% (from a previous rate of over 99% for the initial constraints), but rose to 100% for the remaining 37 trials. During the first nine trials of the first block of the initial constraints, the experiment-wide constraints were upheld 100% for all participants, showing that it was not due to an initial learning period.

As eleven of the sixteen participants produced no illegal experiment-wide errors in either part of the experiment, and the two participants who showed difficulty with the switch produced errors in both parts, it would seem that the ease with which the new constraints were adapted to, is related to the "difficulty" that the participant had with learning the initial constraints, and this was supported by the significant correlation coefficient.

Experiments 4 to 6: General Discussion

This series of experiments has raised several questions relating to the original experiment and paradigm. It has also shed light on the underlying processes involved in the learning of the constraints.

Experiment 4 showed that the artificially constrained phonemes must be chosen carefully, if unintended cues are not to be given to the participants. Therefore, the use of /k/ and /g/ is not recommended, due to the need to use the graphemically represent /g/ as "gh" in order to avoid the hard soft ambiguity.

Experiment 5 raised doubts about the necessity to run the experiment over several days, as it suggested that participants treated each block as a separate condition and did not carry the learning over to the next block. This was evidenced by the lack of confusion when the constraints were changed for the final block, and by the lack of learning shown across the first three blocks. If this is the case, then the need to run the experiment over more than one block must be questioned.

Experiment 6 casts doubt upon this idea, by showing that the participants experienced difficulty in adjusting to new constraints when the change happened within a single block, but also showing that learning had occurred during the first three blocks. Therefore, it may purely be the suddenness of the change that leads to this effect, and not that participants are treating the blocks differently. Although there was no significant difference between the rates for the initial and reversed blocks, the majority of illegal experiment-wide errors that occurred after the change in constraints happened in the first nine trials. However, there may be a problem with the analysis, due to the relatively few data points gathered from the 48 "reversed" trials produced by the participants.

In the light of these studies it was decided that the paradigm should be altered so that a). graphemic information was removed, and b). the task was more difficult for the participant in order to increase the error corpus. To this end, the experiment was adapted for the auditory modality.

Chapter 6: Experiments 7 and 8 - Can Novel Phonemic Constraints be acquired through the Auditory Modality?

These experiments set out to establish whether similar results to those of the Dell et al. (2000) study could be achieved through a version of the paradigm that was adapted for the auditory modality. The alteration of the paradigm for the auditory modality was done in an attempt to rectify the problems encountered in earlier experiments, i.e. cues given by graphemic information (the spelling of /ɛk/ as "ek", which is unusual for English), and the relative paucity of the data. This would also make the experiment closer to a test of phonological encoding and move it away from being a test of reading aloud.

Experiments involving phonological encoding that use an auditory paradigm are not unusual, and have been successfully implemented in previous studies. For example, Treiman and Danis (1988) explored whether syllables are coded in terms of onset and rhyme units, and then whether the rhyme is coded in terms of subunits. Participants in this study were played lists of nonsense syllables through headphones, after which they were they were asked to repeat the syllables in the order that they heard them. Responses were recorded and checked for errors, and then classified and counted in accordance with the requirements of the experiment. This technique could easily be adapted to the paradigm developed by Dell et al. (2000).

Onishi et al. (2002) also used the auditory modality in their extension of the Dell et al. (2000) paradigm. This study did show that phonotactic constraints could be acquired through the auditory modality, as participants responded more quickly to those items that followed the phonotactics of the syllables in the original session. Again, this shows that the auditory modality can successfully be used in the study of speech errors.

Therefore, the Dell et al. (2000) paradigm was modified for the auditory modality. This effectively removed all orthographic information from the stimuli. And, as the syllables were no longer visible while the participant repeated them, the task was more difficult, relying upon the participants' short-term memory for repetition of the stimuli. It was hypothesised that this increase in difficulty would result in an increase in the error corpus.

Experiment 7

Experiment 7 attempted to replicate the results of Experiment 1, using the same experiment-wide constraints, but with the shorter (1 session of 3 blocks) version of the experiment. However, the number of participants was increased to enlarge the size of the data corpus. The use of the ask-tell technique as a measure of the implicitness of the learning was also reinstated in order to make a comparison with Experiment 1. It was hypothesised that the results of Experiment 7 would not be significantly different from Experiment 1(i.e. that it would replicate the findings of Dell et al., 2000), with the exception that there would be an increase the total number of errors.

Experiment7: Method

Participants

Sixteen participants were used for this experiment. They were all students from the School of Psychology, at the University of Wales, Bangor, and were recruited on the basis that English was their native language, and that they had not participated in the previous experiment. None of them had been diagnosed as dyslexic, and they all had normal hearing. The participants received course credits for participation. They were randomly assigned to the two conditions, i.e. two in each of the four conditions. The four conditions were: sef informed, fes informed, sef uninformed, and fes uninformed. Sef and fes refer to the experiment-wide constraints that were applied to /f/ and /s/ in each condition. Informed and uninformed refer to whether the participants' were made aware of these constraints prior to beginning the experiment.

Stimuli and Equipment

This study used the same phonemic constraints as Experiment 1, i.e. /f/ and /s/ were the experimentally constrained phonemes. For participants in the fes condition, /f/ was always an onset and /s/ was always a coda. For participants in the sef condition the reverse was true. The consonants /h/ and /ŋ/ were the language-wide constraints, and the consonants /m/, /n/, /k/, and /g/, each occurred once per sequence and were unconstrained. Each syllable had /ɛ/ as its vowel. A female volunteer, whose voice had been generally agreed upon to be clear and free from any strong regional accent, spoke the syllables. These were recorded using the equipment mentioned earlier. These were transferred to computer and converted to .wav files using Cool Edit Pro 2.0 (Syntrillium Software Corporation, Scottsdale, Arizona, USA), and were then spliced together to form the required sequences of four syllables.

The presentation of the sequences was controlled using the program E-prime. The sequences were played to the participants through headphones, with the participants controlling the rate of presentation by pressing the spacebar when they had finished repeating a sequence.

Procedure

Each participant produced 3 sets of 96 sequences. The participant was first asked to repeat 22 sequences of four syllables. This slower rate produced very few errors, and these were not recorded. Prior to continuing with the experiment, the participants in the informed condition were told about the constraints placed upon the phonemes /f/ and /s/, e.g. the participants in the informed "fes" condition were told, "when you hear an "f" it will always be at the beginning of a syllable, and when you hear an "s" it will always be at the end of a syllable". Those in the uninformed condition were told nothing about the distribution of the phonemes. The participant then produced the 96 sequences three times, with a brief pause (approximately 1 minute) between the trials.

At the end of the session, participants in the uninformed conditions were given a sheet of paper with the following instruction: " In the space below, please note any observations concerning the syllables that you were asked to pronounce in this experiment. For example: All of the syllables contained the vowel *e*."

Experiment 7: Results

Total Number of Errors

A total of 7,037 syllables contained consonant movement errors, giving an error rate of 38.17% for the 18,432 syllables produced by the participants. This is an increase on the error rate for such consonant movements in Experiment 1 (the equivalent visual version of the experiment), which was 3.56% for the 36,864 syllables produced by the participants.

In the visual versions of the experiment, errors involving the substitution of a phoneme for one that was not in the original sequence were rare, and so were not included in the error corpus. However, in the auditory version of the experiment the number of this type of error was larger, and including them significantly increased the error rate from 38.17% to 43.51%, $p < .001$ (sign-test). Recounting the errors from Experiment 1 showed that including this type of error did not significantly increase the error rate (3.56% to 3.68%), $p > .727$ (sign-test). As in Experiment 1 there was a practice effect with there being fewer errors in the final session than in the first, $p < .019$ (sign-test). All the phoneme misordering errors are shown in the matrices in Table 14.

Table 14: Observed Errors in Experiment 7 Classified by Condition and Legality

fes condition – legal outcomes

Coda	Onset: g	k	m	n	f	s	h	ŋ	?
g		22	18	29	32		51		
k	22		24	68	80		53		
m	6	14		26	34		65		
n	31	42	39		51		75		
f									
s	23	56	50	79	175		223		
h									
ŋ	75	61	37	27	114		104		
?									

fes condition – illegal outcomes

Coda	Onset: g	k	m	n	f	s*	h	ŋ**	?
g	14	3	4	31	26	-	19	-	9
k	14	16	13	87	53	2	28	-	33
m	5	11	3	21	36	-	56	-	71
n	28	41	63	68	37	2	57	-	46
f*	-	-	-	-	1	-	2	-	-
s	8	11	24	42		2		-	70
h**	-	-	-	-	-	-	-	-	-
ŋ	17	19	35	8		-		-	283
?	46	-	1	1	9	-	5	-	2

sef condition – legal outcomes

Coda	Onset: g	k	m	n	f	s	h	ŋ	?
g		22	38	38		64	64		
k	25		33	39		41	55		
m	9	27		31		30	58		
n	44	46	47			74	41		
f	55	92	34	80		219	219		
s									
h									
ŋ	44	40	64	31		151	154		
?									

sef condition – illegal outcomes

Coda	Onset: g	k	m	n	f*	s	h	ŋ**	?
g	6	12	29	15	-	48	32	-	34
k	22	13	27	34	-	19	19	-	47
m	6	25	1	26	-	21	67	-	40
n	43	24	66	52	-	52	38	-	66
f	34	39	12	76	-			-	54
s*	8	20	10	28	-	-	9	-	-
h**	-	-	-	-	-	-	-	-	-
ŋ	21	24	54	22	2			-	214
?	92	23	38	77	-	90	56	-	64

- * violates experiment-wide constraints.
- ** violates language-wide constraints.
- ? Indicates the use of a consonant not in the original target stream.

Language-Wide Constraints

From Table 14 it can be seen that none of the errors involving the misplacement of /h/ or /ŋ/ broke the language-wide constraints, and so significantly upheld the phonotactic regularity effect, $p < .001$ (sign-test). This is identical to that found in Experiment 1, which used the same experimental constraints.

Local Positional Constraints

There were a total of 4,194 errors involving a misplacement of /m/, /n/, /k/, or /g/, and of these 2,121 broke the local positional constraints, giving a 49.43% rate at which these constraints were upheld, which was not significantly different from chance, $p > .804$ (sign-test). The rate at which the local positional constraints were upheld in this experiment is similar to that found in Experiment 1 (48.57%). Across participants the range with which the local positional constraints were upheld was 35.71 - 63.09%.

Experiment-Wide Constraints

There were 1,527 errors involving the misplacement of /f/ or /s/, and of these 1,442 obeyed the constraints giving a rate of 94.43%, which was significant, $p < .001$ (sign-test). This is comparable to Experiment 1, which had a rate of 96.55%. Across participants the rate at which the constraints were upheld in Experiment 7 was 60.00 – 100%. There was no evidence that learning of the constraints had occurred, with the final session not having significantly fewer errors than the first, $p > .453$ (sign-test).

As in previous experiments the rate at which the experiment-wide constraints were upheld (94.43%) was not significantly different from that for the language-wide constraints (100%), $p > .125$ (sign-test).

Implicit or Explicit

Participants in the informed condition made more errors that broke the experiment-wide constraints (89.96% were upheld), than those in the uninformed condition (99.32%), which is the opposite of that which would be expected if the learning were due to verbalisable knowledge. However, this difference was not significant, $p > .727$ (sign-test).

When questioned about the syllables that they had been repeating, none of the participants in the uninformed condition reported anything connected to the experiment-wide constraints.

Experiment 7: Discussion

This study aimed to achieve results similar to those achieved in Experiment 1, and those of Dell et al. (2000). With the exception of the large increase in the overall error rate (which was a desired effect), this was successful. The rate at which the language-wide constraints were upheld was identical to that of Experiment 1 and the results of the Dell et al. study. The rate at which the local positional constraints were upheld was, yet again, not significantly different from chance, and the experiment-wide constraints were upheld at a rate below, but not significantly, that of the language-wide constraints. Therefore, this study shows that phonemic constraints can be learned through the auditory modality. The increase in the number of errors involving consonants that were not in the original target sequence cannot be explained purely by the overall increase in errors (there was a significant increase in the number of such errors in Experiment 7, but not in Experiment 1). However, this may be explained by the removal of the graphemic information from the experiment, which acted as a visual reinforcement for the consonants that were in the original sequence.

This result shows that an auditory version of the experiment can be used to increase the data corpus, while still achieving the same results, thereby providing a solution for one of the potential problems with the visual version of the experiment. These results also demonstrate that the effects found in the visual versions of the experiment are true production effects and do not display any artefacts of the reading component. However, it is still unclear whether it offers a solution to the problems associated with the graphemic cues that were being given by the visual version. In order to verify that it was information gained from the graphemes that led to participants noticing the change of constraints in the keg-gek version, it was necessary to replicate this study using this auditory technique.

Experiment 8

Experiment 8 was a replication of Experiment 7, with the exception that the number of participants was reduced from 16 to 8, due to the larger than expected increase in the error corpus from the auditory version of the paradigm (i.e. the number of participants could be decreased while still obtaining a far larger error corpus than in visual version of the paradigm), and the change of the experimentally constrained phonemes to /k/ and /g/. It was hypothesised that the results of Experiment 8 would not be significantly different from Experiment 1(i.e. that it would replicate the findings of Dell et al., 2000), with the exception that there would be an increase the total number of errors.

Method

Participants

Eight participants were used for this experiment. They were all students from the School of Psychology, at the University of Wales, Bangor, and were recruited on the basis that English was their native language, and that they had not participated in the previous experiment. None of them had been diagnosed as dyslexic, and they all had normal hearing. They received course credits for participation. They were randomly assigned to the two conditions, i.e. two in each of the four conditions. The four conditions were: keg informed, gek informed, keg uninformed, and gek uninformed. Keg and gek refer to the experiment-wide constraints that were applied to /k/ and /g/ in each condition.

Informed and uninformed refer to whether the participants' were made aware of these constraints prior to beginning the experiment.

Stimuli and Equipment

This study used the same phonemic constraints as Experiment 2, i.e. /k/ and /g/ were the experimentally constrained phonemes. For participants in the keg condition, /k/ was always an onset and /g/ was always a coda. For participants in the gek condition the reverse was true. The consonants /h/ and /ŋ/ were the language-wide constraints, and the consonants /m/, /n/, /f/, and /s/, each occurred once per sequence and were unconstrained. Each syllable had /ɛ/ as its vowel. The sequences were prepared in the same way as Experiment 7.

The presentation of the sequences was controlled using the program E-prime. The sequences were played to the participants through headphones, with the participants controlling the rate of presentation by pressing the spacebar when they had finished repeating a sequence.

Procedure

Each participant produced 3 sets of 96 sequences. The participant was first asked to repeat 22 sequences of four syllables. This slower rate produced very few errors, and these were not recorded. Prior to continuing with the experiment, the participants in the informed condition were informed about the constraints placed upon the phonemes /k/ and /g/, e.g. the participants in the informed "keg" condition were told, "when you hear a "k" it will always be at the beginning of a syllable, and when you hear a "g" it will always be at the end of a syllable". Those in the uninformed condition were told nothing about the distribution of the phonemes. The participants then produced the 96 sequences three times, with a brief pause (approximately 1 minute) between the trials.

At the end of the session, participants in the uninformed conditions were given a sheet of paper with the following instruction: " In the space below, please note any observations concerning the syllables that you were asked to pronounce in this experiment. For example: All of the syllables contained the vowel *e*."

Experiment 8: Results

Total Number of Errors

A total 3,824 syllables contained consonant movement errors, giving an error rate of 39.70% for the 9,216 syllable productions. This is not significantly different to that for Experiment 7, $p = 1.00$ (sign-test). There was no evidence of a practice effect, with there being no significant decrease in the number of errors between the first and last sessions, $p > .289$ (sign-test). The number of syllables on which there were errors increased to 41.49% when errors involving new consonants were included, although this increase was not significant, $p > .125$ (sign-test). All the consonant misorederings are shown in Table 15.

Table 15: Observed Errors in Experiment 8 Classified by Condition and Legality

keg condition – legal outcomes

Coda	Onset: g	k	m	n	f	s	h	ŋ	?
g		31	17	35	21	32	71		
k									
m		17		39	8	5	25		
n		28	20		22	19	20		
f		32	4	20		19	47		
s		49	11	24	25		48		
h									
ŋ		159	21	24	77	67	104		
?									

keg condition – illegal outcomes

Coda	Onset: g*	k	m	n	f	s	h	ŋ**	?
g	-		4	11	6	8		-	2
k*	-	-	1	2	-	-	-	-	-
m	-	8	3	50	2	8	34	-	9
n	-	19	37	39	13	19	10	-	2
f	-	14	2	27	2	13	40	-	-
s	-	30	11	30	14	46	19	-	1
h**	-	-	-	-	-	-	-	-	-
ŋ	-		15	16	8	18		-	2
?	-	14	23	33	11	36	14	-	1

gek condition – legal outcomes

Coda	Onset: g	k	m	n	f	s	h	ŋ	?
g									
k	73		32	47	30	42	94		
m	11			10	5	6	29		
n	15		23		30	26	17		
f	32		11	32		45	62		
s	51		13	23	36		33		
h									
ŋ	115		47	16	49	48	107		
?									

gek condition – illegal outcomes

Coda	Onset: g	k*	m	n	f	s	h	ŋ**	?
g*	18	2	20	22	15	5	25	-	-
k		-	12	20	23	17		-	5
m	5	-	2	13	8	3	21	-	14
n	18	-	37	7	24	23	18	-	4
f	12	2	7	21	4	12	32	-	-
s	14	1	16	15	23	17	11	-	1
h**	-	-	-	-	-	-	-	-	-
ŋ		1	25	12	19	12		-	7
?	3	-	11	6	4	16	19	-	-

- * violates experiment-wide constraints.
- ** violates language-wide constraints.
- ? Indicates the use of a consonant not in the original target stream.

Language-Wide Constraints

All of the errors involving misplacements of /h/ or /ŋ/ obeyed the language-wide constraints, $p < .008$ (sign-test). This was not significantly different from Experiment 7, $p = 1.00$ (sign-test).

Local Positional Constraints

There were a total of 2,542 errors involving the misplacement of /m/, /n/, /f/, or /s/, and of these 1,231 broke the local positional constraint, giving a rate of 51.57% at which the constraints were upheld. This was not significantly different from chance, $p > .727$ (sign-test). The range across participants was 45.81 – 58.29%. The rate at which the local positional constraints were upheld was not significantly different from the rate for Experiment 7, $p > .727$ (sign-test).

Experiment-Wide Constraints

A total of 788 errors involved the misplacement of /k/ or /g/, and of these 113 broke the experiment-wide constraints, giving a rate of 85.66% at which the constraints were upheld (range across participants 38.24 – 100%), but this was not significant, $p > .07$ (sign-test). This was lower than the previous study, but this can be explained by the data from Participant 2, who upheld the experiment-wide constraints on 38.24% of occasions. With Participant 2 removed from the data, the number of errors decreases to 736, with 29 of these breaking the experiment-wide constraints. This gives a rate of 95.55% at which the constraints were upheld (range across participants 91.30 – 100%), which is significant, $p < .016$ (sign-test).

There was no evidence that learning had occurred, with there being no fewer illegal experiment-wide errors in the final session than in the first (Participant 2 included or excluded), $p = 1.00$ (sign-test).

The experiment-wide errors were significantly different from the language-wide errors with participant 2 included, $p < .031$ (sign-test); but was not significant with Participant 2 excluded, $p > .063$ (sign-test).

These results were not significantly different from Experiment 7, $p = 1.00$ (sign-test).

Implicit or Explicit

Participants in the informed condition upheld the experiment-wide constraints on 93.68% of occasions, and those in the uninformed condition on 78.19% (98.16% with Participant 2 excluded). There was no significant difference between the two rates, whether participant 2 was excluded or not, $p = 1.00$ (sign-test).

None of the participants reported anything connected with the experiment-wide constraints, when questioned at the end of the sessions.

Experiment 8: Discussion

Insofar as there were no significant differences between the two experiments, the aims of this study were met. It has been shown that the use of the auditory modality has effectively removed the confusion that had been attributed to the graphemic information present in the visual version of the experiment. And so, using the auditory version should allow for the use of any pair of phonemes as the experiment-wide constraints in future studies.

The unusual results produced by one participant in the uninformed condition may seem a cause for concern at first sight, but on further investigation it was decided not to exclude them from the data. The participant in question may represent an extreme score for the experiment-wide constraints, but there was no evidence that the participant had difficulty with the "normal" constraints that applied to the English language. The participant upheld the local positional constraints and language-wide constraints at a level that was within one standard deviation of the mean, so there would seem to be no reason to believe that they suffered from an undiagnosed language related condition or that they had hearing difficulties. Therefore, it is surmised that they failed to acquire knowledge of the constraints and, as the rate at which the experiment-wide constraints were upheld (38.24%) was closer to that that has been observed across experiments for the local positional constraints, it is believed that this participant treated the phonemes /k/ and /g/ as such. It is not known why the participant should have failed to acquire the constraints.

Experiment 7 and 8: General Discussion

Experiments 7 and 8 have shown that an auditory version of the experiment can produce the same results as the visual version, whilst increasing the error corpus. It is hoped that this increase in data will improve the sensitivity of the experiment, and so shed further light upon other aspects of the acquisition of phonemic constraints. For that purpose, the "time-course and robustness" experiments were replicated, in the hope that they would be more revealing using the new auditory version of the experiment.

Chapter 7: Experiments 9 and 10 – An Investigation of the Robustness and Time-course of Phonemic Constraints acquired through the Auditory Modality.

Two experiments were conducted using the auditory version of the paradigm. These replicated Experiments 5 and 6 that were conducted using the visual version of the paradigm, i.e. in the first experiment the constraints were reversed at the onset of the fourth block, and in the second experiment the constraints were reversed at the mid-point of the fourth block. The constrained phonemes for both these experiments were /k/ and /g/, as it had been shown that the choice of phonemes was not a consideration when using the auditory modality.

Experiment 9

Experiment 9 replicated Experiment 5, but adapted it to the new auditory version of the paradigm. Using the auditory modality, it was hypothesised that when the constraints were switched (at the onset of the fourth block), the participants would have difficulty in adapting to the change, and would continue to use the original constraints for some time. During this period, the rate at which the experiment-wide constraints were upheld would drop, and there would be a significant difference between the rate in the final block and that in the initial block.

Experiment 9: Method

Participants

Eight participants were used for this experiment. They were all students from the School of Psychology at the University of Wales, Bangor, and were given course credits for participation. None of the participants had been diagnosed as dyslexic, they all had normal hearing, and none of them had participated in the previous experiments. They all had English as their native language. The participants were randomly assigned to the two experimental conditions, initial constraint keg and initial constraint gek. Participants were not informed of the experimental constraints, as knowledge of these may have caused confusion when the constraints were reversed, that was not due to the learning process.

Stimuli and Equipment

This study used the same phonemic constraints as Experiment 2 and 8, i.e. /k/ and /g/ were the experimentally constrained phonemes. For participants in the keg condition, /k/ was always an onset and /g/ was always a coda. For participants in the gek condition the reverse was true. The consonants /h/ and /ŋ/ were the language-wide constraints, and the consonants /m/, /n/, /f/, and /s/, each occurred once per sequence and were unconstrained. Each syllable had /ɛ/ as its vowel. The sequences were prepared in the same way as the previous auditory experiments.

The presentation of the sequences was controlled using the program E-prime. The sequences were played to the participants through headphones, with the participants controlling the rate of presentation by pressing the spacebar when they had finished repeating a sequence.

Procedure

Each participant was given a practice session of 22 strings of four syllables (keg or gek depending on the block to which they were assigned), which were not recorded. The participants then produced four blocks of 96 sequences. The first three blocks all followed the same experiment-wide constraints (dependent upon the assigned condition). The fourth block reversed the experiment-wide constraints. There was a brief pause of approximately 1 minute between each block. All four blocks were recorded using the equipment described earlier.

Experiment 9: Results

Total Number of Errors

Of the 12,288 syllables produced by the participants, there were errors involving phoneme misplacements on 4,165, giving an error rate of 33.89%. If the errors that involved the production of a phoneme not in the original sequence are included, then there were a total of 4,379 errors, and the error rate increases to 35.64%. This is not significantly different from Experiment 8 (which used the same experimental constraints and the auditory modality), $p > .289$ (sign-test).

The practice effect observed in some experiments was not seen in this experiment, and did not achieve significance, $p > .07$ (sign-test). All the phoneme misorderings are represented in Tables 16 and 17.

Table 16: Observed Errors in Experiment 9 Classified by Condition and Legality (gek first condition).

Initial gek condition – legal outcomes									
	Onset								
Coda	g	k	m	n	f	s	h	ng	?
g									
k	24		25	49	33	31	44		
m	20			32	6	16	52		
n	50		41		17	41	42		
f	12		6	10		8	39		
s	15		8	10	20		38		
h									
ng	30		14	11	30	38	63		
?									

Initial gek condition – illegal outcomes									
	Onset								
Coda	g	k*	m	n	f	s	h	ng**	?
g*	1	-	2	-	-	-	1	-	2
k		-	5	29	9	33		-	64
m	14	-	-	28	3	11	38	-	41
n	41	1	41	26	28	41	19	-	33
f	3	-	1	13	1	6	9	-	8
s	3	-	8	14	4	41	5	-	34
h**	-	-	-	-	-	-	-	-	-
ng		-	6	3	6	11		-	11
?	3	-	1	2	-	8	4	-	-

Reversed keg condition – legal outcomes									
	Onset								
Coda	g	k	m	n	f	s	h	ng	?
g		15	5	6	7	9	21		
k									
m		8		17	-	4	13		
n		9	16		7	12	15		
f		9	1	8		9	16		
s		11	5	5	11		14		
h									
ng		19	6	6	6	13	28		
?									

Table 16: continued

Reversed keg condition – illegal outcomes									
				Onset					
Coda	g*	k	m	n	f	s	h	ng**	?
g	2		6	3	-	5		-	-
k*	2	1	2	4	11	3	7	-	-
m	-	3	-	6	-	4	6	-	-
n	8	11	22	7	5	13	12	-	-
f	13	2	1	6	2	12	8	-	-
s	2	7	2	2	-	9	4	-	-
h**	-	-	-	-	-	-	-	-	-
ng	2		1	5	-	1		-	-
?	-	3	2	-	1	2	1	-	-

- * Indicates a violation of the experiment-wide constraints.
- ** Indicates a violation of the language-wide constraints.
- ? Indicates the use of a consonant not in the original target stream.

Table 17: Observed Errors in Experiment 9 Classified by Condition and Legality (keg first condition).

Initial keg condition – legal outcomes									
				Onset					
Coda	g	k	m	n	f	s	h	ng	?
g		88	32	48	33	36	96		
k									
m		35		42	23	16	57		
n		51	31		33	13	23		
f		31	3	31		21	41		
s		45	10	16	28		34		
h									
ng		64	24	15	41	28	83		
?									

Initial keg condition – illegal outcomes									
				Onset					
Coda	g*	k	m	n	f	s	h	ng**	?
g	1		9	10	14	16		-	-
k*	-	1	5	2	2	1	2	-	-
m	-	18	3	25	9	6	36	-	2
n	-	48	22	24	18	11	19	-	1
f	-	16	4	25	1	16	24	-	-
s	-	14	5	14	13	9	8	-	-
h**	-	-	-	-	-	-	-	-	-
ng	-		6	12	4	8		-	-
?	-	7	24	8	3	14	7	-	-

Table 17: Continued

	Reversed gek condition – legal outcomes								
	Onset								
Coda	g	k	m	n	f	s	h	ng	?
g									
k	16		5	11	10	2	16		
m	7			16	1	2	12		
n	5		12		7	5	10		
f	8		3	11		4	13		
s	4		5	3	2		12		
h									
ng	18		2	7	10	12	14		
?									

	Reversed gek condition – illegal outcomes								
	Onset								
Coda	g	k*	m	n	f	s	h	ng**	?
g*	5	1	5	11	3	9	7	-	-
k		-	4	6	4	5		-	-
m	4	-	-	7	2	-	5	-	-
n	6	2	11	11	8	2	8	-	1
f	4	1	-	1	-	2	3	-	-
s	2	-	1	3	5	-	1	-	-
h**	-	-	-	-	-	-	-	-	-
ng		-	1	5	1	3		-	-
?	2	1	6	8	2	6	-	-	-

- * Indicates a violation of the experiment-wide constraints.
- ** Indicates a violation of the language-wide constraints.
- ? Indicates the use of a consonant not in the original target stream.

Language-Wide Constraints

As can be seen in Tables 16 and 17, all errors obeyed the language-wide constraints, $p < .008$ (sign-test). This result was identical to Experiment 8.

Local Positional Constraints

The rate at which the local positional constraints were upheld was 54.46%, and was not significantly different from chance, $p > .07$ (sign-test). The range across participants was 46.14 - 65.08%. This was not significantly different from Experiment 8, $p > .289$ (sign-test).

Experiment-Wide Constraints

The experiment-wide constraints were upheld at a rate not significantly different to Experiment 8, $p = 1.00$ (sign-test). In the initial blocks the constraint was upheld 96.54%, $p < .008$ (sign-test). The range across participants was 57.14 – 100%. There was no evidence of learning, as there was no significant difference between the number of errors in the first and last session of the initial constraints, $p > .625$ (sign-test). The difference between the experiment-wide and language-wide constraints did not achieve significance, $p > .063$ (sign-test).

In the final section (in which the constraints were reversed) the rate at which the experiment-wide constraints were upheld was 59.70%, $p > .727$ (sign-test). The range across participants in the reversed block was 26.67% - 100%. If this rate is compared to that in the equivalent section of the initial constraints (i.e. the first block), then it is found that the constraints were upheld in the first block was 94.69%, $p < .008$ (sign-test). This was not significantly different from the final "reversed" block, $p > .289$ (sign-test). The distribution of these errors throughout these blocks is represented in Figures 12 and 13.

There is a rise in the number of illegal errors after the constraints were reversed, which lasts for approximately seven trials. After this there is only one occasion (Trial 52) on which the number of errors on a single trial rose above three. This is similar to the pattern found in Experiment 6.

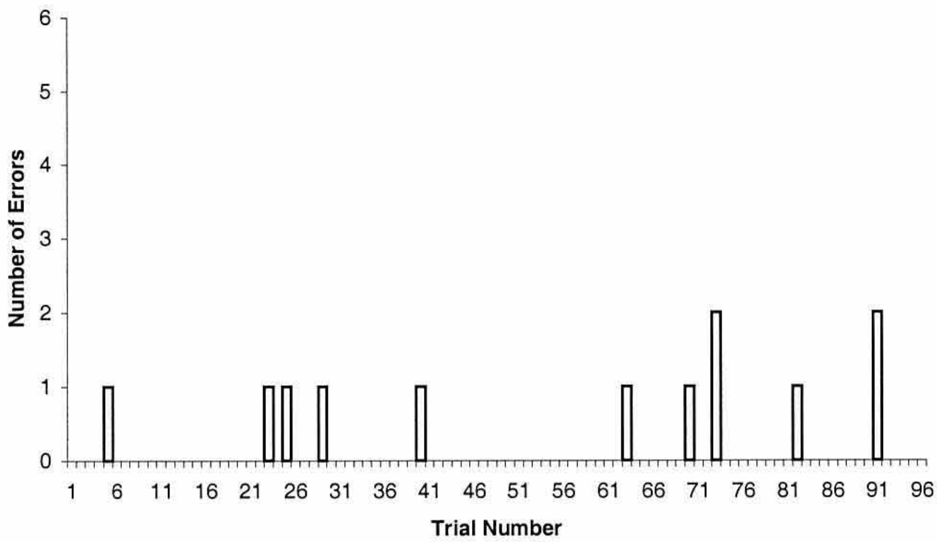

Figure 12: Distribution of Illegal Experiment-wide Errors in Initial Section of Experiment 9.

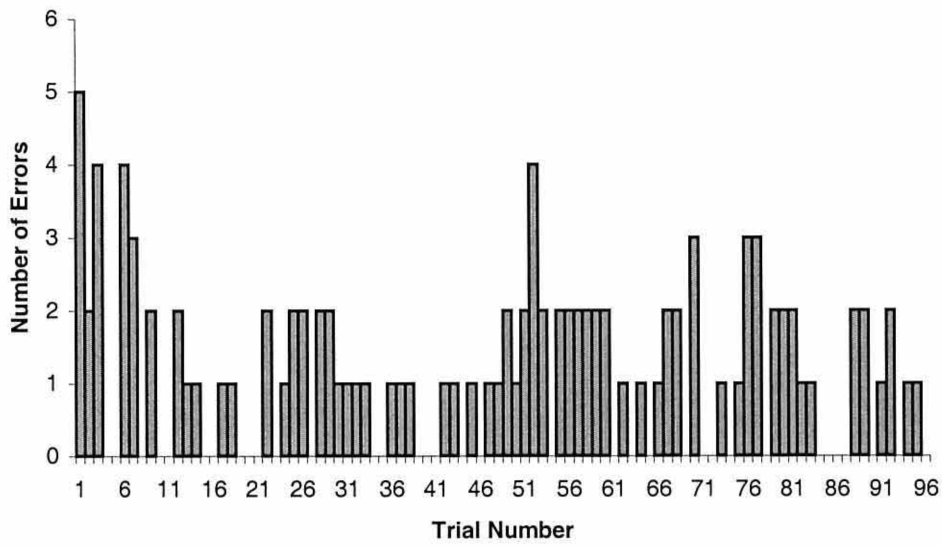

Figure 13: Distribution of Illegal Experiment-wide Errors in "Reversed" Section of Experiment 9.

Experiment 9: Discussion

The core results of this experiment replicated the findings of the previous studies. However, the key results of this study relate to the difficulty that the participants experienced when the experiment-wide constraints were reversed at the beginning of the fourth block. Figures 12 and 13 show that there was interference with the participants' ability to produce the syllables accurately. It also showed, as did Experiment 6, that the first trial after the change showed the greatest number of errors (in both Experiment 6 and 9 there were 5 errors on this trial), and that this initial disruption continued for approximately 10 trials (Mean of 10.76 seconds for Experiment 6 and 20.47 seconds for Experiment 9 – the increase in time for Experiment 9 is due to the necessity to listen to the stimuli before repeating them, in the auditory version of the paradigm). However, in Experiment 9 the rate at which the constraints returned to normal was less clear, as there were errors throughout the rest of the block, although they never returned to the level shown on the first trial.

The difference between these two Experiments was that the reversal of the constraints in Experiment 6 was at the midpoint of the final block, rather than at the beginning. And, therefore, it was decided to re-run Experiment 9, but with the reversal of the constraints happening at the mid-point in the final block.

Experiment 10

For Experiment 10, the reversal of constraints happened at the midpoint of the final block. It was hypothesised that changing the constraints at this point would cause further difficulty for the participants, and so highlight the effect that was seen in Experiment 9. As in Experiment 6, the reversed block would be compared with the equivalent section of the first block (i.e. the first 48 trials of the experiment, where the participants would have had no experience of the experimental constraints).

Method

Participants

Eight participants were used for this experiment. They were all students from the School of Psychology at the University of Wales, Bangor, and were given course credits for participation. None of the participants had been diagnosed as dyslexic, they all had normal hearing, and none of them had participated in the previous experiments. They all had English as their native language. The participants were randomly assigned to the two experimental conditions, initial constraint keg and initial constraint gek. Participants were not informed of the experimental constraints, as knowledge of these may have caused confusion when the constraints were reversed, that was not due to the learning process.

Stimuli and Equipment

This study used the same phonemic constraints as Experiment 2 and 8, i.e. /k/ and /g/ were the experimentally constrained phonemes. For participants in the keg condition, /k/ was always an onset and /g/ was always a coda. For participants in the gek condition the reverse was true. The consonants /h/ and /ŋ/ were the language-wide constraints, and the consonants /m/, /n/, /f/, and /s/, each occurred once per sequence, and were unconstrained. Each syllable had /ɛ/ as its vowel. The sequences were prepared in the same way as the previous auditory experiments.

The presentation of the sequences was controlled using the program E-prime. The sequences were played to the participants through headphones, with the participants controlling the rate of presentation by pressing the spacebar when they had finished repeating a sequence.

Procedure

Each participant was given a practice session of 22 strings of four syllables (keg or gek depending on the block to which they were assigned), which were not recorded. The participants then produced four blocks of 96 sequences. The first three blocks all followed the same experiment-wide constraints (dependent upon the assigned condition). The fourth block reversed the experiment-wide constraints at the mid-point of the block (the 49th trial). There was a brief pause of approximately 1 minute between each block. All four blocks were recorded using the equipment described in earlier.

Experiment 10: Results

Total Number of Errors

There were errors that involved the misplacement of consonants on 4,341 of the 12,288 syllables produced by the participants, giving an error rate of 35.33%. If errors that involved the production of a syllable not in the original sequence are included, then the error rate increases to 37.87%. These error rates are not significantly different from Experiment 9 (the previous auditory time-course experiment), $p > .529$ (sign-test). There was evidence of a practice effect, with there being significantly fewer errors in the final session than in the first, $p < .008$ (sign-test). All of the phoneme misorderings are represented in Tables 18 and 19.

Table 18: Observed Errors In Experiment 10 Classified by Condition and Legality (gek first condition).

Initial gek condition – legal outcomes									
	Onset								
Coda	g	k	m	n	f	s	h	ng	?
g									
k	38		32	45	33	29	75		
m	17			19	10	8	58		
n	21		19		33	46	30		
f	17		6	24		13	49		
s	30		18	16	23		55		
h									
ng	60		15	22	34	27	66		
?									

Initial gek condition – illegal outcomes									
	Onset								
Coda	g	k*	m	n	f	s	h	ng**	?
g*	5	-	5	2	5	2	4	-	-
k		2	7	22	24	10		-	76
m	7	-	1	19	8	5	30	-	60
n	19	2	21	21	23	26	15	-	45
f	2	-	4	12	7	3	5	-	48
s	11	-	4	13	25	6	3	-	41
h**	-	-	-	-	-	-	-	-	-
ng		1	27	7	18	5		-	64
?	-	-	9	1	-	8	1	-	-

Table 18: Continued

Reversed keg condition – legal outcomes

Coda \ Onset	g	k	m	n	f	s	h	ng	?
g		6	1	4	6	4	14		
k									
m		1		2	3	-	9		
n		6	1		7	3	4		
f		1	1	3		1	4		
s		5	-	-	7		8		
h									
ng		15	2	3	12	9	14		
?									

Reversed keg condition – illegal outcomes

Coda \ Onset	g*	k	m	n	f	s	h	ng**	?
g	1		3	1	4	-		-	1
k*	-	-	-	-	3	2	-	-	-
m	1	1	1	1	-	-	4	-	2
n	1	1	3	2	7	2	1	-	1
f	1	1	-	1	1	3	1	-	-
s	-	2	1	-	1	2	2	-	-
h**	-	-	-	-	-	-	-	-	-
ng	2		3	1	4	-		-	-
?	-	1	-	-	1	3	-	-	-

- *Indicates a violation of the experiment-wide constraints.
- ** Indicates a violation of the language-wide constraints.
- ? Indicates the use of a consonant not in the original target stream.

Table 19: Observed Errors In Experiment 10 Classified by Condition and Legality (keg first condition).

Initial keg condition – legal outcomes

Coda \ Onset	g	k	m	n	f	s	h	ng	?
g		68	41	35	36	37	103		
k									
m		28		33	8	15	49		
n		36	35		34	23	17		
f		64	15	53		40	67		
s		77	31	30	26		66		
h									
ng		188	43	18	61	89	176		
?									

Table 19: Continued

				Initial keg condition – illegal outcomes				
				Onset					
Coda	g*	k	m	n	f	s	h	ng**	?
g	-		24	12	12	31		-	-
k*	-	-	-	-	-	1	-	-	-
m	-	24	1	20	3	15	30	-	4
n	4	24	33	15	12	22	18	-	2
f	4	38	11	44	1	40	47	-	-
s	3	31	15	15	10	31	16	-	-
h**	-	-	-	-	-	-	-	-	-
ng	5		42	8	10	45		-	-
?	-	1	6	1	1	5	1	-	-

				Reversed gek condition – legal outcomes					
				Onset					
Coda	g	k	m	n	f	s	h	ng	?
g									
k	10		5	4	2	4	10		
m	3			9	1	4	8		
n	3		5		7	1	2		
f	9		2	5		10	10		
s	3		1	1	2		8		
h									
ng	21		5	2	6	17	19		
?									

				Reversed gek condition – illegal outcomes					
				Onset					
Coda	g	k*	m	n	f	s	h	ng**	?
g*	3	-	1	1	-	-	2	-	-
k		1	2	3	2	6		-	-
m	2	1	-	5	-	2	4	-	1
n	1	-	5	1	2	3	1	-	-
f	4	-	3	4	-	4	7	-	-
s	1	-	-	1	-	6	3	-	-
h**	-	-	-	-	-	-	-	-	-
ng		4	3	5	-	3		-	-
?	-	-	-	-	-	2	-	-	-

- *Indicates a violation of the experiment-wide constraints.
- ** Indicates a violation of the language-wide constraints.
- ? Indicates the use of a consonant not in the original target stream.

Language-Wide Errors

As can be seen in Tables 18 and 19, all errors obeyed the language-wide constraints, $p < .008$ (sign-test). This was identical to Experiment 9.

Local Positional Constraints

The rate at which the local positional constraints were upheld was 54.47%, which was not significantly different from chance, $p > .289$ (sign-test). The range across participants was 45.01 - 62.37%. This was not significantly different from Experiment 9, $p > .248$ (sign-test).

Experiment-Wide Constraints

In the initial blocks the experiment-wide constraints were upheld at a rate of 94.35%, which was significant, $p < .008$ (sign-test). This result was not significantly different from Experiment 9, $p > .399$ (sign-test). The range across participants during these blocks was 79.59 – 100%. There was no evidence of the participants learning the constraints, as there was no significant difference between the number of illegal experiment-wide errors in the first and last of the initial blocks, $p > .689$ (sign-test). There was evidence that the experiment-wide constraints constituted a different population to the language-wide constraints, as there was a significant difference between these two rates, $p < .016$ (sign-test).

The rate at which the constraints were upheld in the final section of the fourth block, in which the experimental constraints were reversed, was 79.83% (range across participants 62.50 – 100%), and this was significant, $p < .008$ (sign-test). When comparing this to the equivalent section of the initial constraints (the first 48 trials of the first block), there was a significant difference (the rate in the first 48 trails was 97.48%), $p < .031$ (sign-test). The distribution of the errors throughout these sections is shown in Figures 14 and 15.

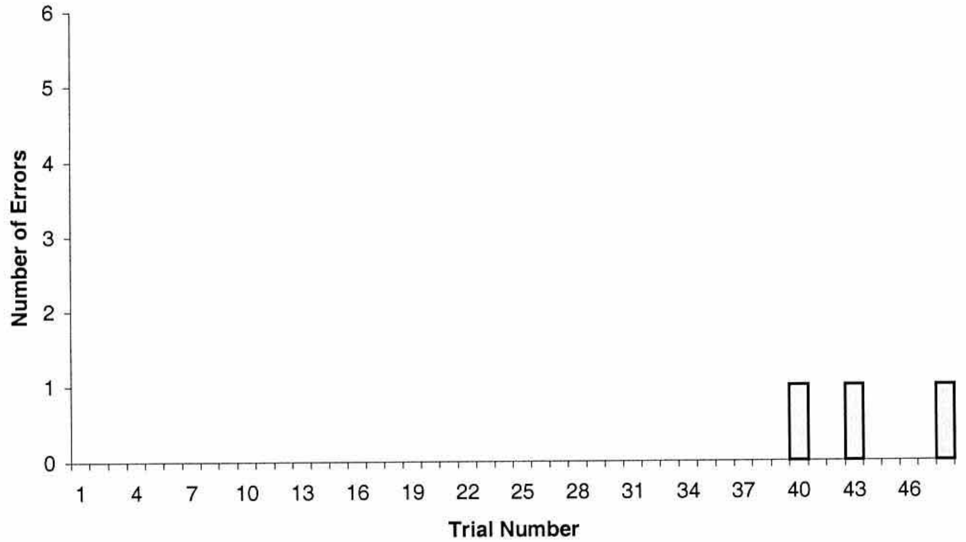

Figure 14: Distribution of Illegal Experiment-wide Errors in Initial Section of Experiment 9.

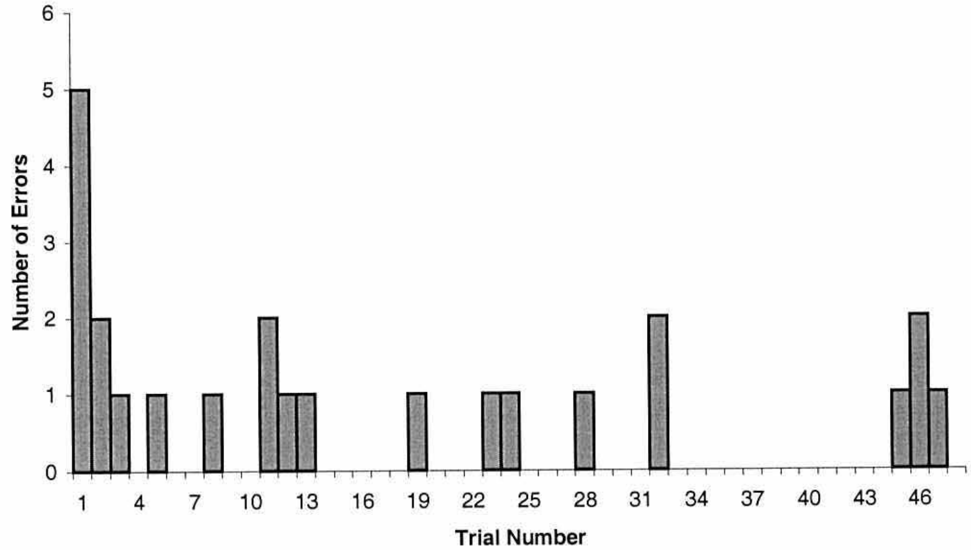

Figure 15: Distribution of Illegal Experiment-wide Errors in "Reversed" Section of Experiment 9.

As can be seen from Figure 14, there were only three errors in the first 48 trials of the initial constraints. However, there is a rise in the number of illegal experiment-wide errors immediately after the reversal of the constraints in the final block (Figure 15). This increase subsides immediately after the first trial, but continues throughout most of the 48 trials of the reversed section. This is a similar pattern to those seen in the previous time-course experiments.

Experiment 10: Discussion

This experiment replicated the main findings of the original studies, with the exception that a significant difference was found between the rate at which the language-wide constraints were upheld, and that for the experiment-wide constraints. This, therefore, supports the claim made by Dell et al. (2000) that the experiment-wide constraints represent the middle ground of a breadth of constraint continuum.

The key result from this experiment was the significant difference between the rate at which the experiment-wide constraints were upheld in the initial 48 trials of the original constraints, and that in the final 48 trials (as per Experiment 6), in which the experiment-wide constraints had been reversed. The pattern of errors following the reversal was similar to that seen in previous time-course experiments, and although the drop in the rate at which the constraints were upheld was lower than had previously been seen, a higher proportion of the participants now showed the reduction. However, the previous experiment had shown a disruption to the rate at which the constraints were upheld that lasted for approximately ten trials after the reversal, and then continued for the rest of the block. In Experiment 10 it appears that the disruption recedes after the first trial, although it does continue for the remainder of the block.

Experiments 9 and 10: General Discussion

Experiment 9 and 10 both showed the pattern of disruption following the reversal of the experiment-wide constraints. However, the effect does not appear as clear as it did in Experiment 6 which showed a period of nine trials following the reversal during which there were errors on every trial, there being only one further error later in the block. It may be that the increased difficulty of the auditory version of the experiment, and the resultant increase in the error rate, has obscured the underlying pattern. Experiment 10 did show, for the first time, a significant reduction in the rate at which the experiment-wide errors were upheld. It is argued that these results show a reduction in the overall rate at which the experiment-wide constraints were upheld, that has a preliminary effect lasting for up to 10 trials, and a secondary effect lasting for an indeterminate time beyond this. Any model that employs a learning algorithm (e.g. Dell et al., 1986; Elman, 1990; Hartley & Houghton, 1996; Jordan, 1996) could represent these data. The rate at which the model "re-learnt" the constraints could easily be modified to account for these data, and would not cause a serious problem.

Chapter 8: Discussion

This series of experiments has addressed a number of different areas, and I will examine each of these in turn, before moving on to look at the implications that these findings have for models of phonological encoding. I will then look at the potential weaknesses of this kind of study. Finally, I shall look at proposals for furthering this area of research.

Local Positional Constraints

One of the aims of this series of experiments was to replicate the findings of Dell et al. (2000), particularly in relation to the local positional constraints. Dell et al. had reported that 68.2 – 77.5 % of local positional constraints had been upheld, which was lower than expected. For example, Ellis (1980) reported a rate of 84 – 96% for the local positional constraints. If the lower rate found by Dell et al. could be shown to be replicable, then there would be implications for models of phonological encoding that represent onset and coda phonemes as separate entities. Such models could not represent such data, as it would be impossible for an onset phoneme to become a coda version of the same phoneme or vice-versa.

All of the experiments in this study showed that the local positional constraint was upheld at a rate of approximately 50% (range 43.04 - 54.47%), and was, therefore, lower than any of the previously reported results. As a result, the validity of these models must be questioned, and alternatives for representing these data must be considered. As previously mentioned, dual-route models (e.g.

Hartley & Houghton, 1996) can represent these data, as could the OSCAR model of Vousden et al. (2000), and these models will be looked at in more detail later in this chapter.

Language-wide Constraints

In all but one of the studies, the language-wide constraints were upheld on 100% of occasions. The one exception involved the production of the cut-off error /ŋ…ɛŋ/. The production of this utterance may have the same implication as the results for the local positional constraints, i.e. models of phonological encoding must be capable of producing any phoneme in either the onset or coda position. However, the production of an utterance that breaks the phonotactics of a language, shows that models of phonological encoding may also have to be capable of producing any phoneme in any position, even if it breaks the phonotactics of a given language.

The production of the cut-off error /ŋ…ɛŋ/ also has consequences for the claim by Dell et al. (2000) that the experiment-wide constraints constituted the middle ground of a breadth of constraint continuum, with language-wide constraints at one end and local positional constraints at the other. This will be discussed in the next section.

Although the production of /ŋ…ɛŋ/ would seem to represent a violation of the language-wide constraints, it may have alternative explanations (e.g. the participant was a second-language speaker of Welsh, and so may have had exposure to /ŋ/ as an onset). Therefore, until further examples of such errors can be found, this result should be viewed with some caution.

Experiment-wide Constraints

The experiment-wide constraints were upheld approximately 90% of the time (range across experiments 80.77 – 100%). And while this is certainly comparable to the results of the studies conducted by Dell et al. (2000), the main findings of this series of studies relate to the claim by Dell and colleagues that the rate at which the experiment-wide constraints were upheld represent the middle ground of a breadth of constraint continuum. Dell et al. claimed to have found a case of the "middle ground" on the basis that the rate at which the experiment-wide constraints were upheld was above that for the local positional constraints, but lower than that for the language-wide constraints. While this is numerically true, no statistical analysis was provided by Dell et al. to support the claim.

The same was true for all but one of the studies in this thesis; the exception being a study in which the rate at which the constraints were upheld was 100% for both the experiment-wide and language-wide constraints. There was concern, however, due to the fact that the rate at which these two constraints were upheld seemed very close to each other; typically the language-wide constraints were upheld on 100% of occasions, and the experiment-wide constraints were upheld at a rate of 90% or higher. Statistical analysis was performed on these data, and on only one occasion was there found to be a significant difference between the language-wide and experiment-wide constraints. Therefore, it would appear that the rates at which these two constraints were upheld are from the same population. And, given the results of the Dell et al. (2000) study, it would therefore seem possible that the same is true for their results.

Two conclusions can be drawn from this result. Firstly, given that there was one occasion on which the language-wide constraints were broken, there may be a ceiling effect for this constraint, and further occurrences of syllables that break the language-wide constraints may be seen if the data corpus were larger. Secondly, if the experiment-wide constraints are from the same population as the language-wide constraints, it would seem that the new constraints are learnt very quickly (they take effect almost immediately, possibly after one exposure), and attain a strength which leads to them being upheld at a similar rate to the language-wide constraints.

Modelling these results should not cause too much difficulty, as any model that uses a learning algorithm (e.g. Dell et al., 1986; Elman, 1990; Hartley & Houghton, 1996; Jordan, 1996) could learn the new constraints, and there are classes of models that can learn in a single exposure (i.e. a single trial serial-order learning mechanism) such as the model developed by Hartley and Houghton (1996).

Bilingual Study

This study (Experiment 3) aimed to support the claim made by Dell et al. (2000), that the strength of the language-wide constraints was due to the participants' lifetime experience of the language. This was also driven by the study carried out by Poulisse and van Lieshout (1997), who showed that the type of errors committed by a learner of a second language were not qualitatively different from those that occurred in the speakers first language, and that the number of errors decreased as the speaker became more competent. Therefore, it was hypothesised that, using the Dell et al. paradigm, a speaker would produce more errors that broke the language-wide constraints when speaking their second language than they produced when speaking their first language.

The results of the study did not support the hypothesis, as there was no significant difference between the language-wide errors produced in the first and second languages. However, this may have been due to the choice of English and Welsh as the two languages. First-language speakers of Welsh grow up in a bilingual society, so are exposed to both languages from birth. This would negate any differences that may have been present. Also, the first-language English speakers had been speaking Welsh for a number of years, so may already have reached a high degree of competency. This would concur with the findings of Poulisse and van Lieshout (1997), who found that the differences in speech errors disappeared over a period of approximately six years.

It is believed that the hypothesis is still valid, but that two factors need to be taken into consideration. Firstly, the languages chosen should not represent a bilingual culture (i.e. one in which both languages are spoken from an early age). And secondly, the participants should have been learning the second language for a relatively short length of time.

Auditory Studies

The results of the auditory studies, using a modified version of the paradigm, show that the experiment can be transferred to the auditory modality whilst still obtaining the same results. Changing the paradigm to the auditory modality also had a number of benefits. Firstly, it produced a larger data corpus thereby increasing the sensitivity of the experiment. Secondly, it removed the orthographic component of the experiment, which was a problem in Experiment 2 (a participant had gained phonotactic information about the experiment-wide constraints, due to the graphemic information contained in the visual version of the paradigm). And, the removal of the graphemic information also removed the necessity for the participants to perform grapheme-to-phoneme conversion. It was felt that grapheme-to-phoneme conversion added error opportunities that would not be present in spontaneous speech and, therefore, made the paradigm closer to an investigation of reading aloud.

Time-course and Robustness

The time-course and robustness experiments aimed to shed light upon the learning process that allowed the participants to acquire the experiment-wide constraints. The modified version of the experiment reversed the experiment-wide constraints, after the participants had learnt the original constraints, and the pattern of errors in the reversed block was compared to the equivalent section of the original constraints. Five experiments were conducted. Experiments 4 and 5 used the visual version of the paradigm and reversed the constraints at the beginning of the fourth block, with different experimentally constrained phonemes being used for each experiment. Experiment 6 also used the visual version of the paradigm, but reversed the constraints at the mid-point of the fourth block. Experiments 9 and 10 used the auditory paradigm, with one experiment reversing the constraints at the beginning of the fourth block, and the other reversing them at the mid-point of the fourth block.

Experiment 4 failed to find any significant difference between the reversed block and the equivalent section of the original constraints. However, this may have been due to flaws within the experimental design (i.e. the choice of phonemes led to the reversal of the constraints being immediately noticeable). Experiment 5 made alterations to the design in order to correct these problems, but still found no significant differences between the two blocks. Both of these experiments reversed the constraints at the beginning of the fourth block and, therefore, there was a rest period between the blocks of approximately two minutes. This may have led to the participants treating the fourth block as a new condition, and so did not apply the rules that they had learnt in the previous block. This idea is supported by the fact that few of the experiments showed evidence of the constraints being learnt across the separate blocks or days, and that the errors in the first block were "randomly" spread throughout the block.

In the light of the findings Experiments 4 and 5, Experiment 6 reversed the constraints at the mid-point of the fourth block, thereby removing the pause between the last trial of the original constraints and the first trial of the reversed constraints. This experiment did show signs of "confusion" from the participants when the constraints were reversed, although there was no significant difference between the blocks. In the original block there were only two errors that broke the experiment-wide constraints, and these were "randomly" spread through the block. However, when the constraints were reversed there was an increase in the number of errors, with twelve of the thirteen errors occurring within the first nine trials. This would seem to add support to the idea that the participants were treating the final block as a separate condition. However, in this experiment, learning was observed across the first three trials (there were significantly less experiment-wide errors in the third block than in the first), and so it may be that it was the "suddenness" of the reversal that caused the confusion, and not that each block is being treated separately. If this is the case, then there will be a maximum length of time between the last occurrence of the original constraints and the first occurrence of the reversed constraints, which will still produce the confusion displayed by the participants in Experiment 6. And, it is felt that this would be a fruitful area for future research.

Moving on to the auditory versions of the time-course and robustness experiments (Experiments 9 and 10), it would seem that the removal of the "reinforcement" provided by the visual representation of the syllables, and the increased difficulty of the task (due to it now requiring the use of the short-term phonological memory), allowed for a longer time period between the last occurrence of the original constraints and the first occurrence of the reversed constraints, whilst still causing confusion among the participants. This was evidenced by the fact that the auditory version of the paradigm produced confusion in the participants whether the reversal occurred at the beginning or in the middle of a block. Experiment 9 demonstrated confusion throughout the block, with the highest number of errors for a single trial occurring on trial one. Experiment 10 showed a similar pattern, with the most errors on a single trial also occurring on the first trial, and with the confusion continuing throughout most of the block, as it had in Experiment 9. It would seem that although the increased difficulty increased the sensitivity of the experiment and, therefore, showed the affect of the constraint reversal wherever it took place, it also extended the period over which the "confusion" could be observed.

Although this series of experiments demonstrated that the learning acquired during the initial phase did have some robustness, it is difficult to quantify this with any degree of accuracy. In the experiments that showed a pattern of confusion, it was repeatedly shown that the highest number of errors occurred on the first trial. However, the length of time that this confusion continued varied as a function of the difficulty of the task and the suddenness with which the reversal of the constraints occurred. Therefore, it would be necessary to conduct more tests before further statements could be made about the time-course and robustness of the learning.

Implicit or Explicit?

When the rates with which the experiment-wide errors were upheld were compared for the informed and uninformed participants, no significant differences were found. This supports the idea that the learning of the constraints is implicit. However, in Experiment 2, one participant correctly commented upon the experiment-wide constraints, stating that the phoneme /k/ always occurred at the end of a syllable. The participant said that they had noticed this due to the *k* always occurring immediately after the *e*, something that would not normally happen in English. Although this may seem to go against the idea that the learning is implicit, it is thought that this occurrence was due to the orthographic information present in the visual version of the paradigm, and the choice of phonemes used in Experiment 2. Therefore, these experiments generally support the conclusion of Dell et al. (2000), that the learning of the constraints is implicit.

Implications of the Results for Models of Phonological encoding

From the results of these studies, it can be seen that any model of phonological encoding must be capable of representing several aspects of the data. Firstly, models cannot represent language-wide constraints in a way that does not allow for the constraints to be broken, as the results have shown that language-wide constraints may possibly be broken. This causes no major problems for models of phonological encoding, as the only restriction is that the knowledge of the constraints cannot be "hard-wired" into the model. This is a particular problem for models that use separate representations for the onset and coda versions of the same phoneme, as such models may only include the appropriate version of the phoneme for the language being modelled. And, if both representations were included, then it would be impossible for an onset representation of a phoneme to become a coda representation. The OSCAR model of Vousden et al. (2000) would have no difficulty with this requirement, as there is no long-term representation of phonetic constraints within the model. The Hartley and Houghton (1996) model does contain a syllable template that contains knowledge of the phonotactics of a language, but, as there is also a direct route from the phonemes to the syllable nodes, the model could represent such data.

Secondly, the models must be capable of acquiring new experiment-wide constraints very quickly, possibly after one exposure. The models developed by Hartley and Houghton (1996) and Vousden et al. (2000) are both capable of learning and recalling a sequence of phonemes in a single trial, and so this aspect causes no problem for them. However, learning the constraints implies that a long-term representation has been acquired. For the Hartley and Houghton model, this learning would involve the strengthening of the connections between specific phoneme nodes and the onset or rhyme node of each syllable. This would lead to the phoneme being more likely to occur in that position. If the model were exposed to the paradigm of Dell et al. (2000), the consistent appearance of a phoneme in the same syllabic position would strengthen the connection to that node. However, the decay would have little effect, as the connection would be constantly reinforced. The non-constrained phonemes would appear in both positions and, therefore, there would be little difference between the connections to the onset or rhyme node. The Vousden et al. model contains no long-term representation of phonotactics, beyond the span of one sequence of phonemes, and so could not acquire new rules.

Thirdly, once the new constraints have been acquired, the model must show an increased error rate for the experiment-wide constraints when they are reversed. This may pose a problem, but not an insurmountable one. As the model must be capable of learning the initial constraints very quickly (possibly in a single trial), the model must develop a long-term representation of the experiment-wide constraints (i.e. an increase in the weights on the connections between the phonemes concerned and the appropriate slot in the syllable frame), otherwise the reversed constraints would also be learnt in a single trial. This requirement could be met by the Hartley and Houghton (1996) model, as the constraints could be learnt in the manner described above. However, the Vousden et al. (2000) model could not do this due to its lack of a long-term representation of phonotactics.

Finally, and most significantly for models that use separate onset and coda representations of the same phoneme, the model must be able to move an onset representation to a coda representation of the same phoneme. This was due to the fact that the rates at which the local positional constraints were upheld were not significantly greater than chance (i.e. 50%). Models that use separate representations for the onset and coda version of the same phoneme cannot represent such data, as all the activation will be associated with either the onset or coda representation. Therefore, it would be impossible for the opposite representation to become active at the recall stage. This does not cause a problem for the models by Hartley and Houghton (1996), or the OSCAR model by Vousden et al. (2000), as neither of these includes hard-wired representations of phonetic constraints.

Therefore, it would seem that the Hartley and Houghton (1996) model offers the most promising explanation of the data. The model would need to be implemented using the Dell et al. (2000) paradigm, to confirm that it would reproduce the data from this series of studies.

Wider Implications of the Data

The Syllable Position Effect in Spontaneous Speech Production-

The rate with which the syllable position constraint was upheld was lower in this series of experiment than in the data previously recorded (e.g. Ellis, 1980). And, the lower rate for the syllable position effect is also at odds with the apparent dominance of the syllable position effect in spontaneous speech production. It is therefore necessary to ask why the syllable position effect should have been virtually absent in this series of studies?

There are two possible explanations for the lower rate seen in the experimental data from this series of studies. Firstly, the rates with which the syllable position effect is upheld may be lower in spontaneous speech production than is actually thought. The reason for this is that listeners can be biased against hearing sounds in unexpected positions (Cole, 1973; Marslen-Wilson & Welsh, 1978). And, even when a speech error has been made, not all listeners notice the mistake, or may report a different mistake (Cutler, 1981). However, as this is also a weakness of speech error data in general, this will be discussed further in the section looking at the weaknesses of the study. Secondly, it may be that there are other factors in spontaneous speech production that increase the size of the syllable position effect. For example, as most collections of errors from spontaneous speech contain more anticipations than perseverations or transpositions (Nooteboom, 1969), then there will be more occurrences that uphold the syllable position effect (i.e. perseverations, by their nature, have to uphold the syllable position effect). However, the data collected for the present series of studies does not allow for any conclusions to be drawn.

The Rapid Acquisition of Novel Phonotactic Constraints –

The data collected in this series of studies support the implicit learning of phonotactic constraints. The error patterns were unaffected by informing the participants of the distribution of the experimentally constrained phonemes. And, with the exception that one participant gained partial knowledge of the constraints due to the graphemic information in the stimuli, none of the uninformed participants reported any knowledge of the distribution of the phonemes. Finally, as the error data were not intended responses by the participants, it would seem unlikely that the error pattern was the result of the speakers' conscious intentions.

Of particular interest is the speed at which participants acquired the novel phonemic constraints. In Experiment 1, which was conducted across four sessions on four separate days, the learning could clearly be seen on the first day. For example, on the first day of Experiment 1, only 3 of the 49 misplacements of /f/ or /s/ were violations of the experiment-wide constraints, and these 3 errors were randomly distributed across the trials. This indicates that learning was extremely rapid. In offering an explanation for the speed of this learning, I find that I must agree with the mechanism proposed by Dell et al. (2000), in that the learning is acquired through a form of phoneme repetition priming. The mechanism proposed by Dell et al. suggests that the utterance of a syllable tunes the speech production system in favour of the production of that, and similar, syllables. The effect of this tuning lasts for more than a single trial, and it accumulates along with the tuning associated with further utterances. This results in the production system being adapted to recent experience.

It would seem that the syllable position effect is, at least partially, due to recent experience. For example, you may say "plone" instead of "phone" because you have recently said words beginning with /pl/. However, you would not say "lphone". Dell et al. (2000) suggest that this may be because the phonology of the speaker's language does not allow for the consonant cluster /lp/. But, they also suggest that it may merely be that you have not had recent experience of saying /lp/. There are phonological theories that see the phonology of a language as being projected from its lexicon (Broe & Pierrehumbert, 1999). That is, there is no abstract representation of phonological patterns that is independent from the lexicon, or, that an abstraction is computed from the lexicon. Dell et al. add to this idea that, if the abstraction is projected, then it is preferentially projected from the most accessible parts of the lexicon, i.e. recently experienced sound forms.

I will now move on to discuss the limitations of this study, and others that use similar techniques for the collection and analysis of speech errors.

Limitations of the Study

Data Analysis –

The data analysis performed in this series of studies was, as mentioned in the general method section, chosen due to the low level of the data (i.e. speech error frequency), and so that the results could be compared with those of Dell et al. (2000), who also used sign tests. There are disadvantages and advantages to using nonparametric statistics. For example, when compared to the *t* test, the sign test makes less than maximum use of the data, in that it substitutes ranks for raw scores (the rank being the result of subtracting the score for one condition from the score for a second condition), thus losing some of the subtle differences among the data points. The result of this loss of data can be seen in Table 20. Although this is an exaggerated example, it does illustrate the potential problems. In Test 1, although half the participants scored far higher in Condition B, the other half scored slightly lower and so there were equal numbers of positive and negative scores. Therefore, there was no significant difference between Conditions A and B in Test 1. However, in Test 2 all the participants scored slightly lower in Condition B than in Condition A, although only by 1 point, and so there was a significant difference.

Table 20: Example of the Loss of Differences among Data Points when using Sign Tests

Test 1		Test 2	
Condition A	Condition B	Condition A	Condition B
53	43	50	49
45	42	58	57
23	21	34	33
34	33	57	56
56	1000	46	45
45	1034	49	48
51	1102	39	38
38	997	25	24
Test 1 $p = 1.00$		Test 2 $p < .008$	

Although, when the assumptions for the *t* test can be met, the sign test has less power, when the assumptions for the *t* test cannot be met, the sign test has considerably more power (Howell, 2002). And, as already mentioned, the data in this series of studies did not meet the assumptions for a *t* test (e.g. the sample size was small, and there were outlying scores).

Reliability of Speech Error Data –

There are issues related to the reliability of speech error data. These issues affect the coding of the speech errors, and the production of the errors by the participant if they are performing a speech repetition task (i.e. shadowing).

Cutler (1981) discussed these issues at length, and it is from this paper that I will be drawing the main points.

The focus of the paper by Cutler (1981) was on the problem of detectability. That is, the problems that may be encountered when looking for speech errors. These problems were considered in relation to the collection of naturalistic speech errors, but some of the problems could equally apply to laboratory error recording. Cutler divides the problems into four areas, and I will look at each of these in turn and consider the implications for the series of studies in this thesis.

1). Slips of the Ear –

Hearing errors are demonstrated less frequently than speech errors, as the listener has to admit to the error for it to be recognised. However, these errors do occur in everyday life. For example:

Spoken: On the eve of the motor show she'll officially open…

Perceived: On the eve of the motor show Sheila Fishley open…

(Cutler, 1981).

Such errors suggest that listeners attempt to make sense of the misheard sentence, and this has also been shown in laboratory experiments. For Example, Warren (1970) replaced single sounds in an utterance with a white noise. It was found that the participants reported hearing a "cough-like" sound that occurred simultaneously with the speech, rather than instead of the speech. This demonstrates that listeners are extremely efficient at constructing a meaningful message, even when the acoustic information is degraded.

With regards to the present series of studies, the tendency for misheard sounds to be reinterpreted as meaningful could be problematic for the coders of speech errors made on nonwords. If the speech error was not particularly clear, or if the coder simply misheard the error, it may be that the coder would be more likely to hear the error as a meaningful sound (i.e. a word). A similar problem could occur for participants in the auditory version of the paradigm. Mishearing the nonwords in this situation may result in them producing a word (lexicalisation). There were occurrences of this in the error data, but it is not known whether this was due to the participants mishearing the stimuli or whether it was due to chance.

In addition to the tendency to interpret misheard utterances as meaningful phrases, it has also been shown that certain elements of words are more likely to be misheard. For example, Browman (1980) showed that consonants are more likely to be misheard than vowels, and Cutler (1980) suggests that the end of a word is less likely to be correctly heard than the beginning.

These tendencies could also be problematic for the studies in this thesis. The data of interest in the present studies were errors involving consonants, which are more likely to be heard as errors and, as all the nonwords were CVC syllables, there could be a bias towards mishearing the coda. Again, this applies to both the coding process and the participants' error production during the auditory version of the paradigm.

2). Shadowing and Mispronunciation Detection –

Several studies have looked at the detection of mispronunciations. These studies present participants with speech that contains errors, and ask them to repeat the speech back as quickly as possible (shadow the text), or make a response as soon as an error is detected. These studies have shown that participants miss the mispronunciation of single sounds very frequently, and that this is particularly true if the mispronunciation differs from the intended sound by only a single feature (e.g. /k/ for /g/), or if the mispronunciation is near the end of a word (Cole, 1973; Marslen-Wilson & Welsh, 1978). Lackner (1980) found that the presentation rate also increased the chance of errors being overlooked.

These problems are perhaps the most serious for the present studies. The phonemes used for the experiment-wide constraints were very close to each other (especially /k/ and /g/) and, as shown by Cole (1973), and Marslen-Wilson and Welsh (1978), this makes them more likely to be overlooked. As the stimuli in the Dell et al. (2000) paradigm are produced in time to the beats of a metronome (2.53 syllables per second), then, according to Lackner (1980), this also increases the chance of a speech error being missed.

3). Perceptual Confusions –

Studies have been conducted that look at the likelihood with which certain sounds are confused with others. These studies usually involve the presentation of isolated syllables, sometimes with the addition of noise masking, and the participants are asked to report what they hear (Miller & Nicely, 1955; Peterson & Barney, 1952; Wang & Bilger, 1973).

Of particular interest to the studies in this thesis are the findings related to the identification of consonants. Miller and Nicely (1955) showed that whether or not a consonant is a nasal is likely to be perceived correctly, as is the difference between voiced and unvoiced consonants. However, the place of articulation is more likely to be mistaken, as is whether or not a consonant is a fricative. Thus, /b/ is more likely to be misheard as /d/, /g/, or /v/, as these differ on place of articulation or frication only, than as /m/ or /p/ that involve a change in nasality and voice respectively.

Findings such as these are important in any attempt to interpret the frequency of confusions between sounds in speech errors (e.g. Shattuck-Hufnagel & Klatt, 1979, 1980; van den Broecke & Goldstein, 1980). The evidence on mispronunciation detection suggests that sound errors are more likely to be overlooked if they differ in only one feature, especially if the altered feature is one that is easily confused (e.g. place of articulation).

According to Cutler (1981) the detectability of certain sounds may also be influenced by the size of the response set. For example, as the nasality value is highly likely to be perceived correctly, there is a greater chance that a listener would correctly hear a mispronounced nasal than they would a mispronounced non-nasal. That is to say, that there seems to be a greater likelihood of a /m/ mispronounced as /n/ being heard as /m/, than for a /b/ mispronounced as /d/ being heard as /b/. This would manifest itself in speech error data as a greater chance for errors to occur in non-nasal consonants.

Goldstein (1980) demonstrated that response bias in perception confusion data for consonants correlates with lexical frequency (i.e. the number of words that contain the sound in question, not the absolute frequency of the sound) and with phonological naturalness (the probability with which a sound appears across all languages). Goldstein points out that although perceptual confusion experiments do show an asymmetry as a result of bias (e.g. /b/ is more likely to be reported as /p/ than vice versa, Wang & Bilger, 1980), no such asymmetry can be seen in the speech error data. That is to say, that for any pair of sounds each one is as equally likely to be substituted for the other (Shattuck-Hufnagel & Klatt, 1980). However, there are indications that phoneme error data may not be as contaminated as the above findings would suggest, as the most commonly reported sound substitutions are between sounds that are most like each other (Cutler, 1981), and these are exactly the kind of errors that, the evidence discussed here suggests, are the most difficult to detect.

The most serious problem for the present studies would seem to be the suggestion that consonants that differ in only one feature are the most likely to be overlooked (leaving aside the fact that these are also the most commonly reported substitutions). As previously mentioned, this could be particularly problematical for the experiment-wide constraints, especially /k/ and /g/ that differ in voicing only.

4). Relative Salience of Beginnings and Ends of Words –

In speech errors there is evidence that the beginning of words are particularly important. Fay and Cutler (1977) demonstrated that form-related word substitution errors resemble their targets very strongly in the initial segments. And, although there are similarities in later segments (Hurford, 1981), these are significantly weaker than earlier ones (Cutler & Fay, 1982).

The implication for the study of speech errors is that early segments are more likely to be noticed than later ones. And, it has been noted that error collections tend to contain more examples of errors from initial position, than for final position (Cohen, 1966; Garret, 1980; Goldstein, 1980; van den Broecke & Goldstein, 1980).

General Problems of Speech Error Collection and Analysis –

The above evidence has raised various questions about the reliability of speech error data. There are problems with the reporting of hearing errors (or rather, knowing when a hearing error has occurred), various difficulties with the frequency with which sounds in speech errors are confused with each other, and an increase likelihood of errors being noticed in the initial position when compared to later positions. Therefore, not all speech errors are equally detectable, and so all collections of speech errors will be confounded, to some degree, by these problems.

As mentioned earlier, Cutler (1981) was discussing the collection of speech errors from everyday spontaneous speech, so not all the criticisms can be made of laboratory based experiments. Indeed, Cutler acknowledges that many of the problems can be avoided, or reduced, by the use of laboratory techniques. The recording of speech errors, the ability to listen to the errors repeatedly, and for the errors to be coded several times by different people, all reduce the chance for error. And, it is believed that the use of such techniques has reduced, although not eradicated, the chance for such errors in this thesis. This is supported by the high rate of agreement between the coders in this thesis (94% - 98%).

Plans for Further Research

The expansion of the Dell et al. (2000) paradigm to include the auditory modality, offers further possibilities for the study of phonological encoding. The removal of the orthographic elements from the task, results in a paradigm that is a more accurate examination of normal phonological encoding. Also, increasing the data corpus, while maintaining the error patterns of the visual paradigm, gives a far more sensitive tool for examining the phonological encoding system.

Possible areas for investigation include the following additional studies:

1. Tracking the course of the acquisition of the constraints by varying the point at which the constraint "reversal" takes place, e.g., after 1 block, 2 blocks, 3 blocks etc. Comparisons could then be made of the number of errors following the switch, as a function of length of experience. If constraints were picked up rapidly (suggestive of a phonological repetition priming mechanism), then I would expect that errors caused by switching constraints would appear early, and return to the previous levels after a few exposures. If a more protracted learning mechanism is involved (changes to longer term representations), then longer exposure before a switch should increase errors.

2. The present results, based on errors following a switch, suggest that the learning effects are fairly short lived. However, this may be an undue interpretation of null effects (participants naturally make an effort to avoid errors). It would, therefore, be fruitful to investigate alternative measures of learning. For example, following a "training" period, where participants repeat strings of nonwords conforming to local phonotactic constraints, they will be presented with a series of isolated stimuli to repeat. Some stimuli will conform to the local constraints embodied in the "training set", and others not (though no stimulus used in the training set will be repeated). The dependent measure will be response latency. If some form of lasting sub-syllabic learning is taking place, then this may be detectable in longer latencies for stimuli that go against previously experienced implicitly learnt constraints.

These methods will allow the further examination of both the acquisition processes, and the robustness involved in the learning. On the theoretical side, it would be useful to frame more specific hypotheses regarding the nature and specific locus of the learning effects. This is best done through (quantitative) computer simulation with an existing model capable of learning, e.g. a modified version of the Hartley and Houghton (1996) model. This will, it is believed, be an extremely fruitful area for exploration.

References

Aslin, R. N., Saffran, J. R., & Newport, E. L. (1998). Computation of conditional probability statistics by 8-month-old infants. *Psychological Science, 9*, 321-324.

Badecker, W., Miozzo, M., & Zanuttini, R. (1995). The two-stage model of lexical retrieval: Evidence from a case of anomia with selective preservation of grammatical gender. *Cognition, 57*, 193-216.

Baddeley, A. D., & Hitch, G. (1974). Working Memory. In G. Bower (Ed.), *The psychology of learning and motivation.* New York: Associated Press.

Baron, J., & Strawson, C. (1976). Use of orthographic and word-specific knowledge in reading words aloud. *Journal of Experimental Psychology: Human Perception and Performance,* 2, 386-392.

Berg, T. (1988). *Die Abbildung des Sprachproduktionprozess in einem Aktivationsflussmodell* [A spreading activation model of language production]. Tübingen, Germany: Max Niemeyer.

Berg, T., & Schade, U. (1992). The role of inhibition in a spreading-activation model of language production: I. The psycholinguistic perspective. *Journal of Psycholinguistic Research, 22,* 405-434.

Besner, D., & Smith, M. C. (1992). Basic processes in reading: Is the orthographic depth hypothesis sinking? In R. Frost & L. Katz (Eds.) *Orthography, phonology, morphology, and meaning* (pp. 45-66). Amsterdam: Elsevier.

Bisiacchi, P. S., Cipolottis, L., & Denes, G. (1989). Impairments in processing meaningless verbal material in several modalities: The relationship between short-term memory and phonological skills. *Quarterly Journal of Experimental Psychology, 41A*, 292-320.

Bock, K., & Loebell, H. (1990). Framing sentences. *Cognition, 35,* 1-39.

Boomer, D. S., & Laver, J. D. M. (1968). Slips of the tongue. *British Journal of Disorders of Communication, 3*, 1-12.

Bredart, S., & Valentine, T. (1992). From Monroe to Moreau: An analysis of face naming errors. *Cognition, 45,* 187-223.

Broe, M, & Pierrehumbert, J. (Eds.). (1999). *Papers in laboratory phonology V: Acquisition and the lexicon.* Cambridge, England: Cambridge University Press.

Browman, C. P. (1980). Perceptual processing: Evidence from slips of the ear. In V. A. Fromkin (Ed.), *Errors in Linguistic Performance: Slips of the Tongue, Ear, Pen, and Hand.* N. Y., London: Academic Press.

Brown, G. D. A., Preece, T., & Hulme, C. (2000). An oscillator-based model of memory for serial order. *Psychological Review, 107*, 127-181.

Brown, R. & McNeil, D. (1966). The 'tip of the tongue' phenomenon. *Journal of Verbal learning and Verbal Behaviour, 5*, 325-337.

Burgess, N., & Hitch, G. (1992). Toward a network model of the articulatory loop. *Journal of Memory and Language, 31 (4)*, 429-460.

Burgess, N., & Hitch, G. J. (1996). A connectionist model of STM for serial order. In S. E. Gathercole (Ed.), *Models of short-term-memory* (pp. 51-72). UK: Erlbaum.

Burgess, N., & Hitch, G. (1999). Memory for serial order: A network model of the phonological loop and its timing. *Psychological Review, 106(3)*, 551-581.

Butterworth, B. (1989). Lexical access in speech production. In W. Marslen-Wilson (Ed.), *Lexical representation and process* (pp. 108-135). Cambridge, MA: MIT Press.

Campbell, R., & Butterworth, B. (1985). Phonological dyslexia and dysgraphia in a highly literate subject: A developmental case with associated deficits of phonemic processing and awareness. *Quarterly Journal of Experimental Psychology, 37A*, 435-475.

Caramazza, A. (1997). How many levels of processing are there in lexical access? *Cognitive Neuropsychology, 14(1)*, 177-208.

Caramazza, A., & Hillis, A. E. (1990). Where do semantic errors come from? *Cortex, 26*, 95-122.

Caramazza, A., Miceli, G., & Villa, G. (1986). The role of (output) phonology in reading writing and repetition. *Cognitive Neuropsychology, 3*, 37-76.

Chambers, K. E., Onishi, K. H., & Fisher, C. (2003). Infants learn phonotactic regularities from brief auditory experience. *Cognition, 87 (2)*, 69-77.

Cohen, A. (1966). Errors of speech and their implication for understanding the strategy of language users. *Zeitschrift für Phonetik, 21*, 177-181.

Cole, R. A. (1973). Listening for mispronunciations: a measure of what we hear during speech. *Perception and Psychophysics, 11*, 153-156.

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review, 82*, 407-428.

Coltheart, M. (1978). Lexical access in simple reading tasks. In G. Underwood (Ed.), *Strategies of information processing.* (pp. 151-216). London: Academic Press.

Coltheart, M. (1985). Cognitive neuropsychology and the study of reading. In M. I. Posner & O. S. M. Marin (Eds.), *Attention and Performance XI* (pp. 3-37). Hillsdale, NJ: Erlbaum.

Cutler, A. (1981). The reliability of speech error data. *Linguistics, 19*, 561-582.

Cutler, A., & Fay, D. A. (1982). One mental lexicon, phonologically arranged: comments on Hurford's comments. *Linguistic Inquiry, 13*, 107-113.

Dell, G. S. (1985). Positive feedback in hierarchical connectionist models: Applications to language production. *Cognitive Science, 9*, 3-23.

Dell, G. S. (1986). A spreading activation theory of retrieval in sentence production. *Psychological Review, 93*, 283-321.

Dell, G. S. (1988). The retrieval of phonological forms in production: Tests of predictions from a connectionist model. *Journal of Memory and Language, 27*, 124-142.

Dell, G. S. (1989). The retrieval of phonological forms in production: Tests of predictions from a connectionist model. In W. Marslen-Wilson (Ed.), *Lexical representation and process* (pp. 136-165). Cambridge, MA: MIT Press.

Dell, G. S., Burger, L. K., & Svec, W. R. (1997). Language production and serial order: A functional analysis and a model. *Psychological Review, 104 (1)*, 123-147.

Dell, G. S., Juliano, C., & Govindjee, A. (1993). Structure and content in language production: A theory of frame constraints in phonological speech errors. *Cognitive Science, 3*, 128-138.

Dell, G. S., & O'Seaghdha, P. G. (1991). Mediated and convergent lexical priming in language production: A comment on Levelt et al. *Psychological Review, 98*, 606-614.

Dell, G. S., Reed, K. D., Adams, D. R., & Meyer, A. S. (2000). Speech errors, phonotactic constraints, and implicit learning: A study of the role of experience in language production. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26(6)*, 1355-1367.

Dell, G. S., & Reich, P. A. (1981). Stages in sentence production: An analysis of speech error data. *Journal of Verbal Learning and Verbal Behaviour, 20*, 611-629.

Dell, G. S., Schwartz, M. F., Martin, N., Saffran, E. M., & Gagnon, D. A. (1997). Lexical Access in Aphasic and Nonaphasic Speakers. *Psychological Review, 104 (4)*, 801-838.

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

del Viso, S., Igoa, J. M., & Garcia-Albert, J. E. (1991). On the autonomy of phonological encoding: Evidence from slips of the tongue in Spanish. *Journal of Psycholinguistic Research, 20,* 161-185.

Eikmeyer, H. J., & Schade, U. (1991). Sequentialization in connectionist language-production models. *Cognitive Systems, 3,* 128-138.

Ellis, A. W. (1980). Errors in speech and short-term memory: The effects of phonemic similarity and syllable position. *Journal of Verbal Learning and Verbal Behaviour, 19,* 624-634.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science, 14*, 179-211.

Estes, W. K. (1972). An associative basis for coding and organisation in memory. In *Coding processes in human memory* (pp. 161-190). Washington, DC: Winson.

Fay, D., & Cutler, A. (1977). Malapropisms and the structure of the mental lexicon. *Linguistic Inquiry, 8,* 505-520.

Feldman, J. A., & Ballard, D. H. (1982). Connectionist models and their properties. *Cognitive Science, 6,* 205-254.

Ferreira, F. (1993). The creation of prosody during sentence production. *Psychological Review, 100,* 233-253.

Fodor, J. A. (1976). *The language of thought.* Hassocks: Harvester Press.

Fodor, J. D., Fodor, J. A., & Garrett, M. F. (1975). The psychological unreality of semantic representations. *Linguistic Inquiry, 6*, 515-531.

Fodor, J. A., Garrett, M. F., Walker, E. C. T., & Parkes, C. H. (1980). Against definitions. *Cognition, 8*, 263-367.

Fromkin, V. A. (1971). The non-anomalous nature of anomalous utterances. *Language, 47,* 27-52.

Fudge, E. C. (1969). Syllables. *Journal of Linguistics, 6,* 193-320.

Garrett, M. F. (1975). The analysis of sentence production. In G. Bower (Ed.), *Psychology of learning and motivation* (pp. 133-177). New York: Academic Press.

Garrett, M. F. (1980). Levels of processing in sentence production. In B. Butterworth (Ed.), *Language production* (pp. 177-220). New York: Academic Press.

Garrett, M. F. (1982). Production of speech: Observations from normal and pathological language use. In A. W. Ellis (Ed.), *Normality and pathology in cognitive functions* (pp. 19-76). London: Academic Press.

Garrett, M. F. (1984). The organization of processing structure of language production. In B. Butterworth (Ed.), *Language production* (pp. 177-220). New York: Academic Press.

Gathercole, S. E., & Adams, A. (1993). Phonological working memory in very young children. *Developmental Psychology, 29(4)*, 770-778.

Gathercole, S. E., & Baddeley, A. D. (1989). Evaluation of phonological STM in the development of vocabulary in children: A longitudinal study. *Journal of Memory and Language, 28*, 200-213.

Gathercole, S. E., Willis, C. S., Baddeley, A. D., & Emslie, H. (1994). The children's test of nonwords repetition: A test of phonological working memory. *Memory. Special Issue: Memory tests and techniques, 2(2)*, 103-127.

Glaser, W. R., & Glaser, M. O. (1989). Context effects in Stroop-like word and picture processing. *Journal of Experimental Psychology: General, 118*, 13-42.

Goldstein, L. (1980). Bias and asymmetry in speech perception. In V. A. Fromkin (Ed.), *Errors in Linguistic Performance: Slips of the Tongue, Ear, Pen, and Hand.* N. Y., London: Academic Press.

Goswami, U., & Bryant, P. (1990). *Phonological skills and learning to read.* London: Erlbaum.

Grossberg, S. (1982). *Studies of mind and brain: Neural principles of learning, perception, development, cognition, and motor control.* Boston: Reidel.

Harley, T. A. (1984). A critique of top-down independent levels models of speech production: Evidence from non-plan-internal speech errors. *Cognitive Science, 15*, 49-87.

Harley, T. A. (1990). Paragrammatisms: Syntactic disturbance or break-down of control? *Cognition, 34*, 85-91.

Hartley, T., & Houghton, G. (1996). A linguistically Constrained Model of Short-Term Memory for Nonwords. *Journal of Memory and Language, 35*, 1-31.

Henson, R. N. A., & Burgess, N. (1998). Representations of serial order. In J. A. Bullinaria, D. W. Glasspool, & G. Houghton (Eds.), *Proceedings of the Fourth Neural Computation and Psychology Workshop: Connectionist representations* (pp. 283-300). London: Springer-Verlag.

Henson, R. N. A., Norris, D. G., Page, M. P. A., & Baddeley, A. D. (1996). Unchained memory: Error patterns rule out chaining models of immediate serial recall. *The Quarterly Journal of Experimental Psychology A, 49*, 80-115.

Houghton, G. (1990). The problem of serial order: A neural network memory of sequence learning and recall. In R. Dale, C. Mellish, & M. Zock (Eds.), *Current research in natural language generation.* London: Academic Press.

Houghton, G., & Hartley, T. (1996). Parallel models of serial behaviour: Lashley revisited. *Psyche: An Interdisciplinary Journal of Research on Consciousness, 2(25)*, pp. NP.

Howell, D. C. (2002). *Statistical methods for psychology.* Pacific Grove, CA: Duxbury.

Hulme, C., Maughn, S., & Brown, G. D. A. (1991). Memory for familiar and unfamiliar words: Evidence for a long-term memory contribution to short-term memory span. *Journal of Memory and Language, 30*, 685-701.

Hulme, C., & Snowling, M. (1992). Deficits in output phonology: An explanation of reading failure? *Cognitive Neuropsychology, 9(1)*, 47-72.

Hurford, J. (1981). Malapropisms, left-to-right listing and lexicalism. *Linguistic Inquiry, 12*, 419-423.

Jones, G. V., & Langford, S. (1987). Phonological blocking in the tip of the tongue state. *Cognition, 26*, 115-122.

Jordan, M. I. (1986). *Serial order: A parallel distributed processing approach.* (Tech. Rep. 8604). La Jolla: University of California, San Diego.

Jusczyk, P. W., Friederici, A. D., Wessels, J. M. I., Svenkerud, V. Y., & Jusczyk, A. M. (1993). Infants' sensitivity to sound patterns of native language words. *Journal of Memory and Language, 32*, 402-420.

Kempen, G., & Huijbers, P. (1983). The lexicalization process in sentence production and naming: Indirect election of words. *Cognition, 14*, 185-209.

Kintsch, W. (1974). *The representation of meaning in memory.* Hillsdale, NJ: Erlbaum.

Lackner, J. R. (1980). Speech production: Correction of semantic and grammatical errors during speech shadowing. In V. A. Fromkin (Ed.), *Errors in Linguistic Performance: Slips of the Tongue, Ear, Pen, and Hand.* N. Y., London: Academic Press.

La Heij, W. (1988). Components of Stroop-like interference in picture naming. *Memory and Cognition, 16,* 400-410.

Lapointe, S. G. & Dell, G. S. (1989). A synthesis of some recent work in sentence production. In G. N. Carlson & M. K. Tananhaus (Eds.), *Linguistic structure in language processing* (pp. 107-156). Dordrecht, The Netherlands: Kluwer.

Lashley, K. S. (1951). The problem of serial order in behaviour. In L. A. Jefress (Ed.). *Cerebral mechanisms in behaviour* (pp. 112-136). New York: Wiley.

Levelt, W. J. M. (1989). *Speaking: From intention to articulation.* Cambridge, MA: MIT Press.

Levelt, W.J.M. (1992). Accessing words in speech production: Stages, processes and representations. *Cognition, 42,* 1-22.

Levelt, W. J. M., Schriefers, H., Vorberg, D., Meyer, A. S., Pechmann, T., & Havinga, J. (1991). The time course of lexical access in speech production: A study of picture naming. *Psychological Review, 98,* 122-142.

MacKay, D. G. (1970). Spoonerisms: The structure of errors in the serial order of speech. *Neuropsychologia, 8,* 323-350.

MacKay, D. G. (1972). The structure of words and syllables: Evidence from errors in speech. *Cognitive Psychology, 3,* 210-227.

MacKay, D. G. (1982). The problem of flexibility, fluency, and speed-accuracy trade-off in skilled behaviors. *Psychological Review, 89,* 483-506.

MacKay, D. G. (1987). *The organization of perception and action: A theory for language and other cognitive skills.* New York: Springer-Verlag.

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

Marslen-Wilson, W. D., & Welsh, A. (1978). Processing interactions and lexical access during word recognition in continuous speech. *Cognitive Psychology, 10,* 29-63.

Martin, N., Dell, D. S., Saffran, E. M., & Schwartz, M. F. (1994). Origins of paraphasia in deep dysphasia: Testing the consequences of a decay impairment to an interactive spreading activation model of lexical retrieval. *Brain & Language, 47,* 609-660.

Martin, N., Weisberg, R. W., & Saffran, E. M. (1989). Variables influencing the occurrence of naming errors: Implications for models of lexical retrieval. *Journal of Memory and Language, 28,* 462-485.

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

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

Meijer, P. J. A. (1994). *Phonological encoding: The role of suprasegmental structures.* Unpublished doctoral dissertation, Nijmegen University, Nijmegen, The Netherlands.

Meyer, A. S. (1991). The time course of phonological encoding in language production: Phonological encoding inside a syllable. *Journal of Memory and Language, 30,* 68-89.

Meyer, A. S. (1994). Timing in sentence production. *Journal of Memory and Language, 33,* 471-492.

Meyer, A. S., & Bock, K. (1992). The tip-of-the-tongue phenomenon: Blocking or partial activation? *Memory & Cognition, 20,* 715-726.

Miceli, G., Benvegnu, B., Capassos, R., & Caramazza, A. (1997). The independence of phonological and orthographic lexical forms: Evidence from aphasia. *Cognitive Neuropsychology, 14(1)*, 35-69.

Miller, G. A., & Nicely, P. E. (1955). An analysis of perceptual confusions among some English consonants. *Journal of the Acoustical Society of America, 27,* 338-352.

Morton, J., & Patterson, K. E. (1980). A new attempt at an interpretation, or, an attempt at a new interpretation. In M. Coltheart, K. E. Patterson, & J. C. Marshall (Eds.), *Deep dyslexia* (pp. 91-118). London: Routledge & Kegan Paul.

Motley, M. T., & Baars, B. J. (1975). Encoding sensitivities to phonological markedness and transition probability: Evidence from spoonerisms. *Human Communication Research, 2*, 351-361.

Nooteboom, S. G. (1969). The tongue slips into patterns. In A. G. Sciarone, A. J. van Essen, & A. A. van Raad (Eds.), *Leyden studies in linguistics and phonetics* (pp. 114-132). The Hague: Mouton.

Norris, D., McQueen, J. M., Cutler, A., & Butterfield, S. (1997). The possible-word constraint in the segmentation of continuous speech. *Cognitive Psychology, 34*, 191-243.

Onishi, K. H., Chambers, K. E., & Fisher, C. (2002). Learning phonotactic constraints from brief auditory experience. *Cognition, 83*, 13-23.

Paap, K. R., & Noel, R. W. (1991). Dual route models of print to sound: Still a good horse race. *Psychological Research, 53*, 13-24.

Patterson, K., Graham, N., & Hodges, J. R. (1994). The impact of semantic memory loss on phonological representations. *Journal of Cognitive Neuroscience, 6(1)*, 57-69.

Patterson, K. E., & Morton, J. (1985). From orthography to phonology: An attempt at an old interpretation. In K. E. Patterson, J. C. Marshall, & M. Coltheart (Eds.), *Surface dyslexia: Neuropsychological and cognitive studies of phonological reading* (pp. 335-359). London: Erlbaum.

Peterson, G. E., & Barney, H. L. (1952). Control methods used in a study of the vowels. *Journal of the Acoustical Society of America, 24*, 175-184.

Peterson, R. R., & Savoy, (1998). Lexical selection and phonological encoding during language production: Evidence for cascading processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 539-557.

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

Plaut, D. C., & McClelland, J. L. (1993). Generalization with componential attractors: Word and nonwords reading in an attractor network. In *Proceedings of the 5th Annual Conference of the Cognitive Science Society* (pp. 824-829). Hillsdale, NJ: Erlbaum.

Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review, 103*, 56-115.

Plaut, D. C., & Shalice, T. (1993). Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology, 10,* 377-500.

Potter, M., & Faulconer, B. A. (1975). Time to understand pictures and words. *Nature, 253,* 437-438.

Poulisse, N., & van Lieshout, A. (1997). Second language acquisition: some evidence from L2 learners' slips of the tongue. *Communication & Cognition, 30 (1)*, 55-73.

Rapp, B., Benzing, L., & Caramazza, A. (1997). The autonomy of lexical orthography. *Cognitive Neurosychology, 14(1)*, 71-104.

Roelofs, A. (1992). A spreading activation theory of lemma retrieval in speaking. *Cognition, 42,* 107-142.

Roelofs, A. (1996). Serial order in planning the production of successive morphemes of a word. *Journal of memory and Language, 35*, 854-876.

Roelofs, A. (1997). *The WEAVER model of word-form encoding in speech production.* Manuscript submitted for publication.

Romani, C. A. (1992). *The representation of prosodic and syllabic structure in speech production.* Unpublished doctoral dissertation, Johns Hopkins University, Baltimore.

Rumelhart, D. E., & Norman, D. (1982). Simulating a skilled typist: A study of skilled motor performance. *Cognitive Science, 6*, 1-36.

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

Saffran, J. R., & Thiessen, E. D. (2003). Pattern induction by infant language learners. *Developmental Psychology, 39 (3)*, 484-494.

Schriefers, H. (1992). Lexical access in the production of noun phrases. *Cognition, 45,* 33-54.

Schriefers, H., Meyer, A. S., & Levelt, W. J. M. (1990). Exploring the time-course of lexical access in production: Picture-word interference studies. *Journal of Memory and Language, 29,* 86-102.

Sevald, C. A., Dell, G. S., & Cole, J. (1995). Syllable structure in speech production: Are syllables chunks or schemas? *Journal of Memory and Language, 34,* 807-820.

Shattuck-Hufnagel, S. (1979). Speech errors as evidence for a serial-ordering mechanism in sentence production. In W. E. Cooper & E. C. T. Walker (Eds.) *Sentence processing: Psycholinguistic studies presented to Merrill Garrett* (pp. 295-342). Hillsdale, NJ: Erlbaum.

Shattuck-Hufnagel, S. (1992). The role of word structure in segmental serial ordering. *Cognition, 42,* 213-259.

Shattuck-Hufnagel, S., & Klatt, D. (1979). The limited use of distinctive features and markedness in speech production: Evidence from speech error data. *Journal of Verbal Learning and Verbal Behavior, 18*, 41-55.

Shattuck-Hufnagel, S., & Klatt, D. (1980). How single phoneme error data rule out two models of error generation. In V. A. Fromkin (Ed.), *Errors in Linguistic Performance: Slips of the Tongue, Ear, Pen, and Hand.* N. Y., London: Academic Press.

Stemberger, J. P. (1982). *The lexicon in a model of language production.* Unpublished doctoral dissertation. University of California, San Diego, La Jolla.

Stemberger, J. P. (1983). *Speech errors and theoretical phonology: A review.* Bloomington: Indiana University Linguistics Club.

Stemberger, J. P. (1985). An interactive activation model of language production. In A. W. Ellis (Ed.), *Progress in the psychology of language* (Vol. 1, pp. 143-186). Hillsdale, NJ: Erlbaum.

Stemberger, J. P. (1990). Wordshape errors in language production. *Cognition, 35,* 123-157.

Stemberger, J. P. (1991). Apparent anti-frequency effects in language production: The addition bias and phonological underspecification. *Journal of Memory and Language, 30,* 161-185.

Theios, J., & Amrhein, P. C. (1989). Theoretical analysis of the cognitive processing of lexical and pictorial stimuli: Reading, naming, and visual and conceptual comparisons. *Psychological Review, 96,* 5-24.

Treiman, R. (1986). The division between onsets and rhymes in English syllables. *Journal of Memory and Language, 25,* 476-491.

Treiman, R., & Danis, C. (1988). Short-term memory errors for spoken syllables are affected by the linguistic structure of the syllables. *Journal of Experimental Psychology: Learning, Memory and Cognition, 14(1),* 145-152.

Treisman, M., Cook, N., Naish, P. L. N., & MacCrone, J. K. (1994). The internal clock: Electroencephalographic evidence for oscillatory processes underlying time perception. *The Quarterly Journal of Experimental Psychology A, 47,* 241-289.

Treisman, M., Faulkner, A., & Naish, P. L. N. (1992). On the relation between time perception and the timing of motor action: Evidence for a temporal oscillator controlling the timing of movement. *The Quarterly Journal of Experimental Psychology A, 45,* 235-263.

Van den Broecke, M. P. R., Goldstein, L. (1980). Consonant features in speech errors. In V. A. Fromkin (Ed.), *Errors in Linguistic Performance: Slips of the Tongue, Ear, Pen, and Hand.* N. Y., London: Academic Press.

Vigliocco, G., Antonini, T., & Garrett, M. F. (1997). Grammatical gender is on the tip of Italian tongues. *Psychological Sciences, 8(4),* 314-317.

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

Vousden, J. I., Brown, G. D. A., & Harley, T. A. (2000). Serial Control of phonology in speech production: A hierarchical model. *Cognitive Psychology, 41*, 101-175.

Wang, M. D., & Bilger, R. C. (1973). Consonant confusions in noise: a study of perceptual features. *Journal of the Acoustical Society of America, 54*, 1248-1266.

Warren, R. M. (1970). Perceptual restoration of missing speech sounds. *Science, 167*, 392-393.

Wells, R. (1951). Predicting slips of the tongue. *Yale Scientific Magazine, 3,* 9-30.

Wickelgren, W. A., (1969). Context-sensitive coding, associative memory, and serial order in (speech) behaviour. *Psychological Review, 76*, 1-15.

Zorzi, M., Houghton, G., & Butterworth, B. (1998). Two routes or one in reading aloud? A connectionist dual-process model. *Journal of Experimental Psychology: Human Perception and Performance, 24(4)*, 1131- 1161.

Appendix 1

Syllable Types in the fes Condition in each Set of 96 Sequences		
Two restricted sounds (expected frequency: 24)	One restricted sound (expected frequency: 12)	No restricted sounds (expected frequency:8)
heng	keng	kem
hes	gheng	ken
fes	meng	keg
feng	neng	mek
	hek	men
	heg	meg
	hem	ghem
	hen	ghek
	fek	ghen
	feg	nem
	fen	nek
	fem	neg
	kes	
	ghes	
	nes	
	mes	