A Distributed, Developmental Model of Word Recognition and Naming

Mark S. Seidenberg
McGill University
Montreal, Quebec, Canada

James L. McClelland
Carnegie Mellon University

A parallel distributed processing model of visual word recognition and pronunciation is described. The model consists of sets of orthographic and phonological units and an interlevel of hidden units. Weights on connections between units were modified during a training phase using the back-propagation learning algorithm. The model simulates many aspects of human performance, including (a) differences between words in terms of processing difficulty, (b) pronunciation of novel items, (c) differences between readers in terms of word recognition skill, (d) transitions from beginning to skilled reading, and (e) differences in performance on lexical decision and naming tasks. The model's behavior early in the learning phase corresponds to that of children acquiring word recognition skills. Training with a smaller number of hidden units produces output characteristic of many dyslexic readers. Naming is simulated without pronunciation rules, and lexical decisions are simulated without accessing word-level representations. The performance of the model is largely determined by three factors: the nature of the input, a significant fragment of written English; the learning rule, which encodes the implicit structure of the orthography in the weights on connections; and the architecture of the system, which influences the scope of what can be learned.

The recognition and pronunciation of words is one of the central topics in reading research and has been studied intensely in recent years (see articles in Besner, Waller, & MacKinnon, 1985, and Coltheart, 1987, for reviews). The topic is important primarily because of the immediate, "on-line" character of language comprehension (Marslen-Wilson, 1975), that is, the fact that text and discourse are interpreted essentially as the signal is perceived. Two aspects of lexical processing contribute to this characteristic of reading. First, words can be identified quickly; the rate for skilled readers typically exceeds five words per second (Rayner & Pollatsek, 1987). Second, identification of a word entails the activation of several types of associated information or codes, each of which contributes to the rapid interpretation of text. These codes include one or more meanings of

a word (Seidenberg, Tanenhaus, Leiman, & Bienkowski, 1982; Swinney, 1979), information related to its pronunciation or sound (Baron & Strawson, 1976; Gough, 1972; Tanenhaus, Flanigan, & Seidenberg, 1980), and information concerning the kinds of sentence structures in which the word participates (McClelland & Kawamoto, 1986; Tanenhaus & Carlson, 1989). Understanding the meanings of words is obviously an important part of text comprehension. The phonological code may be related to the retention of information in working memory, while other comprehension processes such as syntactic analyses or inferencing continue (Baddeley, 1979; Daneman & Carpenter, 1980). The third type of information facilitates the development of representations concerning syntactic and conceptual structures (Tanenhaus & Carlson, 1989). The picture that has emerged is one in which lexical processing yields access to several types of information in a rapid and efficient manner. Readers are typically aware of the results of lexical processing, not the manner in which it occurred. One of the goals of research on visual word recognition has been to use experimental methods to unpack these largely unconscious processes; in the model that we present in this article, we attempt to give an explicit, computational account of them.

Word recognition is also important because acquiring this skill is among the first tasks confronting the beginning reader; moreover, deficits at the level of word recognition are characteristic of children who fail to acquire age-appropriate reading skills (Perfetti, 1985; Stanovich, 1986). The model that we describe provides an account of the kinds of knowledge that are acquired, how they are used in performing different reading tasks, and the bases of some types of reading impairment. Specific deficits in word recognition are also observed as a consequence of brain injury; the study of these deficits has provided important information concerning the types of knowledge and processes involved in normal reading and clues to their neuro-

physiological bases (Patterson, Coltheart, & Marshall, 1985). Our model provides the basis for an account of some characteristics of pathological performance in terms of damage to the normal processing system; this aspect of the model is discussed in Patterson, Seidenberg, and McClelland (in press).

Finally, visual word recognition provides an interesting domain in which to explore general ideas concerning learning, the representation of knowledge, and skilled performance because it is a relatively mature area of inquiry. There has been an enormous amount of empirical research on the topic, and several models have already been proposed (Coltheart, 1978; Forster, 1976; LaBerge & Samuels, 1974; McClelland & Rumelhart, 1981; Morton, 1969). Our goal has been to develop an explicit, computational model that accounts for much of this extensive body of knowledge. At the same time, word recognition provides an interesting domain in which to explore the properties of the connectionist or parallel distributed processing approach to understanding perception, cognition, and learning (McClelland & Rumelhart, 1986b; Rumelhart & McClelland, 1986b) that we have used in this research. In particular, our model illustrates an important feature of this approach, the emergence of systematic, "rule-governed" behavior from a network of simple processing units.

Scope of the Problem

In acquiring word recognition skills, children must come to understand at least two basic characteristics of written English. First, there is the alphabetic principle (Rozin & Gleitman, 1977), the fact that in an alphabetic orthography there are systematic correspondences between the spoken and written forms of words. Beginning readers already possess large oral vocabularies; their initial problem is to learn how known spoken forms map onto unfamiliar written forms. The scope of this problem is determined by characteristics of the writing system. The alphabetic writing system for English is a code for representing spoken language; units in the writing system—letters and letter patterns—largely correspond to speech units such as phonemes. However, the correspondence between the written and spoken codes is notoriously complex; many correspondences are inconsistent (e.g., -AVE is usually pronounced as in GAVE, SAVE, and CAVE, but there is also HAVE) or wholly arbitrary (e.g., -OLO in COLONEL, -PS in CORPS).

These inconsistencies derive from several sources. One is the fact that the writing system also encodes morphological information. Chomsky and Halle (1968) argued that English orthography represents a solution to the problem of simultaneously representing information concerning phonology and morphology. According to their analysis, the writing system follows a general principle whereby phonological information is encoded only if it cannot be derived from rules that are conditioned by morphological structure. Thus, words with seemingly irregular pronunciations such as SIGN and BOMB preserve in their written forms information about morphological relations among words (SIGN–SIGNATURE; BOMB–BOMBARD); the correct pronunciations can be derived from a morphophonemic rule governing base and derived forms. Whatever the validity of Chomsky and Halle's approach (see Bybee, 1985, for an alternative), it is clear that some irregular correspondences between graphemes and phonemes are due to the competing demand that the writing system preserve morphological information.

Other inconsistencies derive from the fact that the spoken forms of words change over time, whereas the written forms are essentially fixed. In British English, for example, the word BEEN is a homophone of BEAN; in American English, it is a homophone of BIN. The American pronunciation has changed through a process of phonological reduction, resulting in an irregular spelling–sound correspondence. These diachronic changes in pronunciation are an important source of irregularities in spelling–sound correspondences. There are other sources as well, principally lexical borrowing from other languages, periodic spelling reforms, and historical accident. The net result is that the writing system encodes information related to pronunciation and sound, but the correspondence between written and spoken forms is not entirely regular or transparent. English is said to have a "deep" alphabetic orthography, in contrast to a "shallow" orthography such as that in Serbo-Croatian, which has more consistent spelling–sound correspondences (Katz & Feldman, 1981).

A second aspect of the writing system that the child must learn about concerns the distribution of letter patterns in the lexicon. Only some combinations of letters are possible, and the combinations differ in frequency. These facts about the distribution of letter patterns give written English its characteristic redundancy. Of the many possible combinations of 26 letters, only a small percentage yield letter strings that would be permissible words in English. An even smaller percentage are realized as actual entries in the lexicon. As Adams (1981) noted,

> From an alphabet of 26 letters, we could generate over 475,254 unique strings of 4 letters or less, or 12,376,630 of 5 letters or less. Alternatively, we could represent 823,543 unique strings with an alphabet of only 7 letters, or 16,777,216 with an alphabet of only 8. For comparison, the total number of entries in Webster's New Collegiate Dictionary is only 150,000. (p. 198)

Constraints on the forms of written words may play an important role in the recognition process. The reader must discriminate the input string from other vocabulary items, a task that might be facilitated by knowledge of the letter combinations that are permissible or realized. Many studies have provided evidence that skilled readers use this knowledge (see Henderson, 1982, for a review).

Orthographic redundancy also provides cues to other aspects of lexical structure, specifically, syllables and morphemes. For example, the written forms of words typically provide cues to their syllabic structure (Adams, 1981) for the following reason. Syllables derive from articulatory–motor properties of the spoken language; essentially, they reflect the opening and closing movements of the jaw cycle (Fowler, 1977; Seidenberg, 1989). Thus, the capacities of the articulatory–motor apparatus constrain the possible sequences of phonemes. Moreover, there are language-specific constraints on phoneme sequencing. Written English is largely a code for representing speech; hence, properties of speech such as syllables tend to be reflected in the orthography. For example, the fact that the letters GP never appear in word-initial position derives from a phonotactic constraint on the occurrence of the corresponding phonemes. These letters can appear at the division between two syllables (e.g., PIGPEN),

reflecting the fact that there are more constraints on the sequencing of phonemes within syllables than between. As a result, the letter patterns at syllable boundaries tend to be lower in frequency than the letter patterns that occur intrasyllabically (Adams, 1981; Seidenberg, 1987). Thus, facts about the distribution of phonemes characteristic of spoken syllables are reflected in the distribution of letter patterns in their written realizations. As in the case of grapheme–phoneme correspondences, however, the realizations of syllables in the orthography are not entirely consistent, as illustrated by minimal pairs such as WAIVE–NAIVE, BAKED–NAKED, and DIES–DIET, which are similar in orthography but differ in syllabic structure. Thus, written English provides cues to syllabic structure, but these cues are not entirely reliable.

The situation is similar when we turn to the level of morphology, which concerns the organization of sublexical units that contribute to meaning. The meaning of a word is often a compositional function of the meanings of its morphemes; consider prefixed words such as PREVIEW and DECODE. That English is systematic in this regard is seen in the coining of new words such as PRECOMPILE or DEBUG. That it is inconsistent is illustrated by words such as PRETENSE (unrelated to TENSE) or DELIVER (unrelated to LIVER). Again, written English encodes information related to morphological structure, but not in a regular or consistent manner.

In sum, the English orthography partially encodes several types of information simultaneously. The reader's knowledge of the orthography can be construed as an elaborate matrix of correlations among letter patterns, phonemes, syllables, and morphemes. Written English is an example of what we call a *quasiregular system*—a body of knowledge that is systematic but admits many irregularities. In such systems, the relations among entities are statistical rather than categorical. Many other types of knowledge may have this character as well.

The child's problem, then, is to acquire knowledge of this quasiregular system. The task of reading English might be facilitated by the systematic aspects of the writing system such as the constraints on possible letter sequences and the correspondences between spelling and sound. However, there are barriers to using these types of information. Facts about orthographic redundancy cannot be used until the child is familiar with a large number of words. Acquiring useful generalizations about spelling–sound correspondences is inhibited by the fact that many words have irregular correspondences, and these words are overrepresented among the items the child learns to read first (e.g., GIVE, HAVE, SOME, DOES, GONE). The child must, nonetheless, learn to use knowledge of the orthography in a manner that supports the recognition of words within a fraction of a second.

Our model addresses the acquisition and use of knowledge concerning orthographic redundancy and orthographic–phonological correspondences. We focus on these types of information because they are sufficient to account for phenomena related to the processing of monosyllabic words, which is our model's domain of application. In the general discussion we return to issues concerning syllabic and morphological knowledge and the processing of more complex words. Our goal has been to determine how well the basic phenomena of word naming and recognition might be accounted for by a minimal model of

lexical processing, in which as little as possible of the solution of the problem is built in and as much as possible is left to the mechanisms of learning. The model is realized within the connectionist framework being applied to many problems in perception and cognition (McClelland & Rumelhart, 1986b; Rumelhart & McClelland, 1986b). The model provides an account of how these types of knowledge are acquired and used in performing simple reading tasks such as naming words aloud and making lexical decisions. One of the main points of the model is that, because of the quasiregular character of written English, it is felicitous to represent these types of knowledge in terms of the weights on connections between simple processing units in a distributed memory network. Learning then involves modifying the weights through experience in reading and pronouncing words. Thus, the connectionist approach is ideally suited to accounting for word recognition because of the nature of the task, which is largely determined by these characteristics of the orthography.

A key feature of the model we propose is the assumption that there is a single, uniform procedure for computing a phonological representation from an orthographic representation that is applicable to irregular words and nonwords as well as regular words. A central dogma of many earlier models (e.g., the dual-route accounts of Coltheart, 1978; Marshall & Newcombe, 1973; Meyer, Schvaneveldt, & Ruddy, 1974) is that irregular words and nonwords require separate mechanisms for their pronunciation: Irregular words require lexical lookup because they cannot be pronounced by rule, whereas nonwords require a system of rules because their pronunciations cannot be looked up (see Seidenberg, 1985b, 1988, for discussion). Whether, in fact, two mechanisms are required, and whether they are the mechanisms postulated in dual-route models, are among the main issues that our model addresses. The model does not entail a lookup mechanism because it does not contain a lexicon in which there are entries corresponding to individual words. Nor does it contain a set of pronunciation rules. Instead, it replaces both by a single mechanism that learns to process regular words, irregular words, nonwords, and other types of letter strings through experience with the spelling–sound correspondences implicit in the set of words from which it learns.

The model gives a detailed account of a range of empirical phenomena that have been of continuing interest to reading researchers, including (a) differences between words in terms of processing difficulty, (b) differences between readers in terms of word recognition skill, (c) transitions from beginning to skilled reading, and (d) differences between silent reading and reading aloud. The model also provides an account of certain forms of dyslexia that are observed developmentally and as a consequence of brain injury.

Description of the Model

Precursors

Before we turn to the model itself, it is important to acknowledge several precursors of this work. In some ways, this model can be seen as an application of many of the principles embodied in the interactive activation model of word perception (McClelland & Rumelhart, 1981) to a more distributed model

of the kind used by Rumelhart and McClelland (1986a) in their simulation of the acquisition of past tense morphology. This work draws heavily on insights into distributed representation due primarily to Hinton (1984; Hinton, McClelland, & Rumelhart, 1986) and exists only because of Rumelhart, Hinton, and Williams's (1986) discovery of a learning procedure for multilayer networks. In applying many of these ideas to the task of reading, we follow in the footsteps of Sejnowski and Rosenberg's (1986) NETtalk model, which was the first application of the Rumelhart et al. algorithm to the problem of learning the spelling–sound correspondences of English. Sejnowski and Rosenberg recognized that this knowledge could be represented by a parallel distributed network rather than a set of pronunciation rules. Our goal was to explore the adequacy of this approach by developing a model that could be related to a broad range of phenomena concerning human performance.

Several previous models of visual word recognition also influenced the development of the somewhat different account presented here. Among them are Morton's (1969) logogen model, the dual-route model of Coltheart (1978), and Glushko's (1979) lexical analogy model. Later in the text we show how our model relates to these precursors. Finally, our account of lexical decision is similar to ones proposed by Gordon (1983) and Balota and Chumbley (1984).

The Larger Framework

As we have noted, the model was developed with the goal of using a minimal architecture in which the learning aspect played a dominant role. Some minimal structural assumptions were required, however. A second goal was to keep things as simple as possible; therefore, the model we have implemented is a simplification of the larger, somewhat richer processing system that surely is required to account for aspects of single word processing outside our primary concerns. We begin by describing the larger framework of which the model we have implemented is a part; we then describe the simplifications and detailed assumptions of the implementation.

The larger framework assumes that reading words involves the computation of three types of codes: orthographic, phonological, and semantic. Other codes are probably also computed (concerning, e.g., the syntactic and thematic functions of words), but we have not included them in the present model because they probably are more relevant to comprehension processes than to the recognition and pronunciation of monosyllabic words. Each of these codes is assumed to be a distributed representation; that is, to be a pattern of activation distributed over a number of primitive representational units. Each processing unit has an activation value that in our model ranges from 0 to 1. The representations of different entities are encoded as different patterns of activity over these units.

Processing in the model is assumed to be interactive (Marslen-Wilson, 1975; McClelland, 1987; McClelland & Rumelhart, 1981; Rumelhart, 1977). That is, we assume that the process of building a representation at each of the three levels both influences, and is influenced by, the construction of representations at each of the other levels. We also assume, in keeping with this inherently interactive view, that word processing can be influenced by contextual factors arising from syntactic, semantic,


Figure 1. General framework for lexical processing. (The implemented model is in boldface type.)

and pragmatic constraints, although the scope and locus of these effects is a matter of current debate (see McClelland, 1987; Rumelhart, 1977; Tanenhaus, Dell & Carlson, 1987, for discussion). We assume that at least some of these types of information constrain the construction of the representation at the semantic level and, thus, indirectly influence construction of representations at the other levels, and conversely, that the construction of a representation of the context is influenced by activation at the semantic level.

As in other connectionist models, processing is mediated by connections among the units. However, it is well known that there are limits on the processing capabilities inherent in networks in which there are only direct connections between units at different representational levels (Hinton et al., 1986; Minsky & Papert, 1969). In view of these limits, it is crucial that there be a set of so-called *hidden units,* mediating between the pools of representational units.

The assumptions described thus far are captured in Figure 1, in which each pool of units—both hidden units and representational units—is represented by an ellipse. Connections between units on different levels are represented by arrows. These arrows always run in both directions, in keeping with the assumption of interactivity.

The Simulation Model

The model that we have actually implemented is illustrated in Figure 2 and is the part of Figure 1 in boldface type. This simplified model removes the semantic and contextual levels, leaving only the orthographic level, the phonological level, and the interlevel of hidden units between these two. Furthermore, as an additional simplification, we have not implemented feedback from the phonological to the hidden units; this means, in effect, that phonological representations cannot in fact influence the construction of representations at the orthographic

Figure 2. Structure of the implemented model and number of units.

level. There is, however, feedback from the hidden units to the orthographic units. This feedback plays the role of the top-down word-to-letter connections in the interactive activation model of word perception, allowing the model to sustain, reinforce, and clean up patterns produced by external input to the orthographic level.

Several further assumptions were required in implementing this simplified model. These assumptions can be grouped into three types: Processing assumptions, specifying the way in which activations influence each other; learning assumptions, specifying how connection strength adjustment takes place as a result of experience; and representational assumptions, specifying how orthographic and phonological characteristics of words are to be represented.

Processing assumptions. At a fine-grained level, we believe it would be most accurate to characterize processing in terms of the gradual buildup of activation (McClelland, 1979; McClelland & Rumelhart, 1981), subject to a considerable amount of random noise. However, for simplicity, the simulation model actually computes activations deterministically in a single processing sweep. This simplification makes simulation of the learning process feasible because it speeds up simulation by a couple of orders of magnitude.

Details of the processing assumptions of the model are as follows. Each word-processing trial begins with the presentation of a letter string, which the simulation program then encodes into a pattern of activation over the orthographic units, according to the representational assumptions described later. Next, activations of the hidden units are computed on the basis of the pattern of activation at the orthographic level. For each hidden unit, a quantity called the *net input* is computed; this is simply the activation of each input unit, times the weight on the connection from that input unit to the hidden unit, plus a bias term that is unique to the unit. Thus, for hidden unit i, the net input is given by

$$\text{net}_i = \sum_j w_{ij} a_j + \text{bias}_i.$$

Here j ranges over the orthographic units, a_j is the activation of orthographic unit j, bias_i is the bias term for hidden unit i, and w_{ij} is the weight of the connection to unit i from unit j. The bias term may be thought of as an extra weight or connection to the unit from a special unit that always has activation of 1.0.

The activation of the unit is then determined from the net input using a nonlinear function called the *logistic function:*

$$a_i = \frac{1}{1 + e^{-\text{net}_i}}.$$

The activation function must be nonlinear for reasons described in Rumelhart, Hinton, and Williams (1986). It must be monotonically increasing and have a smooth first derivative for reasons having to do with the learning rule. The logistic function satisfies these constraints.

Once activations over the hidden units have been computed, they are used to compute activations for the phonological units and new activations for the orthographic units based on feedback from the hidden units. These activations are computed following exactly the same computations already described; first, the net input to each unit is calculated, based on the activations of all of the hidden units; then the activation of each of these units is computed, based on the net inputs.

Learning assumptions. When the model is initialized, the connection strengths and biases in the network are assigned random initial values between ±0.5. This means that each hidden unit computes an entirely arbitrary function of the input it receives from the orthographic units and sends a random pattern of excitatory and inhibitory signals to the phonological units and back to the orthographic units. This also means that the network has no initial knowledge of particular correspondences between spelling and sound, nor can its feedback to the orthographic units effectively sustain or reinforce inputs to these units. Thus, the ability to recreate the orthographic input and generate its phonological code arises as a result of learning from exposure to letter strings and the corresponding strings of phonemes.

Learning occurs in the model in the following way. An orthographic string is presented and processing takes place, as described, producing first a pattern of activation over the hidden units and then a feedback pattern on the orthographic units and a feedforward pattern on the phonological units. At this point, these two output patterns produced by the model are compared to the correct, target patterns that the model should have produced. The target for the orthographic feedback pattern is simply the orthographic input pattern; the target for the phonological output is the pattern representing the correct phonological code of the presented letter string. We assume that in reality the phonological pattern may be supplied as explicit external teaching input—as in the case in which the child sees a letter string and hears a teacher or other person say its correct pronunciation—or self-generated on the basis of the child's prior knowledge of the pronunciations of words and the contexts in which they occur.

For each orthographic and phonemic unit, the difference between the correct or target activation of the unit and its actual activation is computed as follows:

$$d_i = (t_i - a_i).$$

The learning procedure adjusts the strengths of all of the connections in the network in proportion to the extent to which this change will reduce a measure of the total error, E. Thus,

$$\Delta w_{ij} = -\varepsilon \frac{\partial E}{\partial w_{ij}}.$$

Here ε is a learning rate parameter, and E is the sum of the difference terms for each unit, each squared:

$$E = \sum_i d_i^2.$$

The term $\partial E/\partial w_{ij}$ is the partial derivative of the error measure with respect to a change in the weight to unit i from unit j.[1]

The algorithm that is used to compute the partial derivative for each weight is the *back-propagation* learning procedure of Rumelhart, Hinton, and Williams (1986). Readers are referred to Rumelhart, Hinton, and Williams for an explanation of how these partial derivatives are calculated. For our purposes the important thing to note is that the rule changes the strength of each weight in proportion to the size of the effect changing it will have on the error measure. Large changes are made to weights that have a large effect on E, and small changes are made to weights that have a small effect on E.

Representational assumptions. In reality, the orthographic and phonological representations used in reading are determined by learning processes, subject to initial constraints imposed by biology and prior experience. The learning of these representations is beyond the scope of the model; for simplicity, we have treated them as fixed in the simulations. Our choice of representations is not intended to be definitive; rather, it was motivated primarily by a desire to capture a few general properties that we would expect such representations to acquire through learning, while at the same time building in very little specifically about the correspondences between spelling and sound or about the particular kinds of letter and phoneme strings that are words in English.

In representing a word's orthographic or phonological content, it is not sufficient to activate a unit for each of the letters or phonemes in the word because this would yield identical representations for pairs such as BAT and TAB. It is necessary to use some scheme that specifies the context in which each letter occurs. We chose to use a variant of Wickelgren's (1969) *triples* scheme, following Rumelhart and McClelland (1986a), rather than the strict positional encoding scheme of McClelland and Rumelhart (1981). In this we have given the model a tendency to be sensitive to local context rather than absolute spatial position, because letters occurring in similar local contexts activate units in common. Thus, for example, the letter string MAKE is treated as the set of letter triples _MA, MAK, AKE, and KE_ (where _ is a symbol representing the beginning or ending of a word), whereas the phoneme string /mAk/ is treated as the set of phoneme triples _mA, mAk, Ak_.[2]

Note that we do not claim that this scheme in its present form is fully sufficient for representing all of the letter or phoneme sequences that form words (see Pinker & Prince, 1988). However, we are presently applying the model only to monosyllables, and the representation is sufficient for these (see general discussion). Extensions of the representation scheme can be envisioned in which more global properties such as approximate position with respect to particular vowel groups is also represented in conjunction with each triple. Such a scheme would largely collapse to the present one for monosyllables.

An important way in which our representations differ from Wickelgren's (1969) proposal in that we do not assume a one-to-one correspondence between triples and units; rather, each

triple is encoded as a distributed pattern of activation over a set of units, each of which participates in the representation of many triples. The representation used at the phonemic level is the same as that used by Rumelhart and McClelland (1986a). Each unit represents a triple of phonetic features, one feature of the first of the three phonemes in each triple, one feature of the second of the three, and one of the third.[3] For example, there is a unit that represents [vowel, fricative, stop]. This unit should be activated for any word containing such a sequence, such as the words POST and SOFT. Word boundaries are also represented in the featural representation, so that there is a unit, for example, that represents [vowel, liquid, word boundary]; this unit would come on in words like CAR and CALL. There are 460 such units, and each phoneme triple activated 16 of them; see Rumelhart and McClelland (1986a) for details.

The representation used at the orthographic level is similar to that used at the phonological level, except that in this instance 400 units were used, and each unit was set up according to a slightly different scheme. For each unit, there is a table containing a list of 10 possible first letters, 10 possible middle letters, and 10 possible end letters. These tables are generated randomly except for the constraint that the beginning or end of word symbol does not occur in the middle position. When the unit is on, it indicates that one of the 1,000 possible triples that could be made by selecting one member from the first list of 10, one from the second, and one from the third is present in the string being represented. Each triple activated about 20 units. Although each unit is highly ambiguous, over the full set of 400 such randomly constructed units, the probability that any two sequences of three letters would activate all and only the same units in common is effectively zero.[4] In sum, both the phonological and the orthographic representations can be described as coarse-coded, distributed representations of the sort discussed by Hinton et al. (1986). The representations allow any letter and phoneme sequences to be represented, subject to certain saturation and ambiguity limits that can arise when the strings get too long. Thus, there is a minimum of built-in knowledge

---

[1] In fact, the size of the adjustments made to the strengths of the connections in the model is given by a somewhat more complex expression, as follows:

$$\Delta w_{ij} = -\varepsilon \frac{\partial E}{\partial w_{ij}} + \alpha \cdot \Delta w'_{ij}.$$

Here w' refers to the previous increment to the weights, and α is a parameter between 0 and 1. α can be thought of as specifying how much momentum there is in the magnitude of the changes made to the weights.

[2] Here, and elsewhere in the article, we use the following notation for representing phonemes: A = a in GAVE; a = a in HAVE; O = o in POSE; U = o in LOSE; i = i in LINT; I = i in PINT; E = ee in SEED; $^\wedge$ = u in MUST; u = oo in BOOK; o = o in HOT; W = ow in HOW; * = aw = PAW.

[3] The set of phonological features used was somewhat simplified, so that certain phonemes pairs (e.g., the initial phonemes in CHIN and SHIN) were not in fact distinguished. See Rumelhart and McClelland (1986a) for details.

[4] Ghosts are capable of appearing in this representation when it becomes too "saturated"; that is, when too many of the units are on at one time. This is one reason why a richer representation would be required to represent multisyllabic words.

of orthographic or phonological structure. The use of a coding scheme sensitive to local context does promote the exploitation of local contextual similarity as a basis for generalization in the model; that is, what it learns to do for a grapheme in one local context (e.g., the M in MAKE) will tend to transfer to the same graphemes in similar local contexts (e.g., the M's in MADE and MATE and, to a lesser extent, M's in contexts such as MILE and SMALL).

Naming and Lexical Decision

The model produces patterns of activation across the orthographic and phonological units as its output. For naming, we assume that the pattern over the phonological units serves as the input to a system that constructs an articulatory–motor program, which in turn is executed by the motor system, resulting in an overt pronunciation response. In reality, we believe that these processes operate in a cascaded fashion, with the triggering of the response occurring when the articulatory–motor program has evolved to the point at which it is sufficiently differentiated from other possible motor programs. Thus, activation would begin to build up first at the orthographic units, propagating continuously from there to the hidden and phonological units and from there to the motor system in which a response would be triggered when the articulatory–motor representation became sufficiently differentiated.

The simulation model simplifies this picture. Activations of the phonological units are computed in a single step, and the construction and execution of articulatory motor programs are unimplemented. The activations that are computed in this way can be shown to correspond to the asymptotic activations that would be achieved in a cascaded activation process (Cohen, Dunbar, & McClelland, 1989). To relate the patterns of activation the model produces to experimental data on latency and accuracy of naming responses, we use what we call the *phonological error score*, which is the sum of the squared differences between the target activation value for each phonological unit and the actual activation computed by the network.

It is important not to treat the error score as a direct measure of the accuracy of an overt response made by the network. In fact, the error scores can never actually reach zero because the logistic function used in setting the activations of units prevents activations from ever reaching their maximum or minimum values. Rather, with continued practice, error scores simply get smaller and smaller, as activations of units approximate more and more closely the target values. This improvement continues well beyond the point at which the correct answer is the best match to the pattern produced by the network. To determine the best match, we simply use the error score as a measure of how closely the pattern computed by the net matches the correct pronunciation and each of several other possible pronunciations. In general—as we will present in detail later—we find that after training, the error score is lower for the correct pronunciation than for any other.

Even when the target code provides the best fit to the pattern of activation over the phonological units, there is still room for considerable variation in error scores. We assume that lower error scores are correlated with faster and more accurate responses under time pressure. The rationale for the accuracy assumption is simply that a low error score signifies that the pattern produced by the network is relatively clear and free from noise, and so provides a better signal for the articulatory–motor programming and execution processes to work with. The rationale for the speed assumption is as follows: In a cascaded system, patterns that are asymptotically relatively clear (low in error) will reach a criterion level of clarity relatively quickly. Simulations demonstrating this point are presented in Cohen et al. (1989).

Thus far, we have discussed the use of the phonological error score as a measure of the accuracy and speed of naming. We shall see that this measure is sensitive to familiarity; the more frequently the network has processed a particular word, the smaller the error score will be. The error score computed over the orthographic units is likewise related to familiarity. Because the input pattern is also the target pattern for the orthographic feedback, the orthographic error score is simply the sum of the squares of the differences between the feedback pattern computed by the network and the actual input to the orthographic units. For lexical decision, in which the subject's task is to judge whether the stimulus is a familiar word, we assume that a measure like the orthographic error score is actually used in making this judgment. Note that this differs from our use of the phonological error score in accounting for naming performance. The calculated phonological error score is simply a measure of the asymptotic clarity of the computed phonological representation, which we use to predict naming latencies. In contrast, a measure like the orthographic error score is assumed to be actually computed by subjects as part of the decision process. Because the orthographic input is in fact presented to the subject, it seems reasonable to assume that subjects can compare this input to the internally generated feedback from the hidden units and use the result of this comparison process as the basis for judgments of familiarity. This issue is considered again in the section on Lexical Decisions in the Model.

Our goal was to develop a working simulation model that exhibited many of the basic phenomena of word recognition and naming, based on a theory of what is learned and how it is represented. When it came to assessing the performance of the model, we discovered that there was a simple monotonic relationship between error scores and naming latencies. This result was quite surprising, given that the error scores depend on some of the more arbitrary aspects of the simulations, such as the number of orthographic encoding units, the number of entries in each unit's table, and the Wickelphone output scheme. In addition, the error scores reflect the effects of training on a corpus of less than 3,000 words, many fewer than a skilled reader would know. Finally, we calculated the error scores using the weights from 250 learning epochs; other weights could have been used. The net result is that, although the fit between error scores and latencies is very good, it is by no means perfect. In future research it will be necessary to determine whether a better fit could be achieved by addressing some of the limitations of the present implementation (see the General Discussion section).

Parameters

Once the input and output representations are specified, the model leaves us with very few free parameters. There are two

free parameters of the input representation, the number of letters in each unit's table and the number of such units. After picking plausible initial values for these, however, we did not manipulate them. There are two other parameters: the learning rate ε and the number of hidden units. For both of these parameters, the initial values we chose (0.05 and 200, respectively) have turned out to produce quite good quantitative accounts of the phenomena. It is interesting that manipulation of the learning rate parameter has rather little effect; acquisition is not so much slower as less noisy with a smaller learning rate. Manipulation of the number of hidden units, however, has interesting and illuminating effects, which are considered when we discuss individual differences in learning to read. For completeness, two other parametric details should be mentioned. First, as targets for learning, we used the values of 0.9 and 0.1; that is, the model was trained to set the activations of units that should be on to 0.9 and the activations of units that should be off to 0.1, rather than to the extreme values of 1.0 and 0.0. Second, the momentum parameter α, was set at 0.9. These values are commonly used in models of this type (see, e.g., Sejnowski & Rosenberg, 1986, and Footnote 1).

The Training Regime

There is one other factor that has profound effects on the model's performance, namely, the set of learning experiences with which it is trained. The training corpus we have used consists of all of the monosyllabic words in the Kucera and Francis (1967) word count that consist of three or more letters. From these we removed proper nouns, words we judged to be foreign, abbreviations, and morphologically complex words that were formed from the addition of a final -s or -ed inflection. Note that this is not a complete list of monosyllables; the word FONT, for example, is one of many that do not appear in Kucera and Francis. Nevertheless, the corpus provides a reasonable approximation of the set of monosyllables in the vocabulary of an average American reader. To this list we added a number of words that had been used in some of the experiments that we planned to simulate. Some of these words were inflected forms (e.g., DOTS); for these, the Kucera–Francis frequency of the base form was used. Others were simply entered into the word list with frequencies of 0. The resulting list contained 2,897 words. This total includes 13 homographs (words such as WIND and BASS that have two pronunciations) that were entered twice, once with each pronunciation. Thus, there were 2,884 unique orthographic patterns in the list.

The training regime was divided into a series of epochs. Within an epoch, each word had a chance of being presented that was monotonically related to its estimated frequency:

$$p = K \log(\text{frequency} + 2).$$

A value of K was chosen so that the most frequent word (THE) had a probability of about .93. Words occurring once per million had probabilities of about .09, and words not occurring in the Kucera–Francis count had probabilities of .057. Thus, the expected value of the number of presentations of a word over 250 epochs ranged from about 230 to about 14. Because the sampling process is in fact random, there was about a 5%

chance that one of the least probable words would be presented less than 7 times in 250 epochs.

The use of the logarithmic frequency transformation radically compresses the range of variation in the presentation frequencies of different words. For example, the word THE is presented only about 10 times as often as a word like RAKE, whereas in the Kucera and Francis (1967) corpus, THE occurs more than 69,000 times as frequently as RAKE. This compression was motivated in part by practical considerations. We did not think it feasible to run sufficient trials to achieve even the current level of exposure to the least frequent words without compressing the frequency range. Using compressed frequencies, we achieved this level of exposure with a total of 150,000 learning trials. Using uncompressed frequencies, something on the order of 5,000,000 learning trials would have been required; this would take several months given available computational resources.

There are several other reasons why some compression of the frequency range is preferable to the use of raw frequencies. First, the word frequencies found in a count such as Kucera and Francis (1967) are based on samples of written text taken from adult sources and do not reflect the relative frequencies of words experienced by beginning readers. In the early stages of learning to read, the words to which the child is exposed necessarily span a much narrower range of frequencies than in the adult norms. With additional experience, the relative frequencies of words begin to differentiate. The logarithmic transform, which compresses the range of frequencies, is thus more in keeping with the child's experience than with the adult's. We thought that it was important to approximate this aspect of the child's experience because the largest gains in reading skill occur early in training. This is true both for the model, as will be shown, and for children, whose knowledge of the spelling–sound correspondences of the language expands rapidly during the first year or two of instruction.

A second point is that the frequency transform compensates for the effects of another aspect of the implemented model, the restricted corpus of words used in training. The training corpus consists entirely of monosyllabic words and includes only a few morphologically complex words. Children learn the spelling–sound correspondences of the language on the basis of exposure to both mono- and multisyllabic words, including morphological relatives that were excluded from the simulations. For example, the model is trained on a word such as DUNK but does not gain additional feedback from related items such as DUNKED or DUNKING. The net effect is that the listed frequencies of the base words tend to underestimate their actual frequency of occurrence in the language. This factor will have little effect on the model's performance on higher frequency words; the morphological relatives tend to be much lower in frequency, and including these words would result in little additional learning. However, the morphological relatives of the lower frequency items tend to be as frequent or more frequent than the base words themselves; excluding these items eliminates an important source of feedback. Thus, the restrictions on the training set disproportionately penalize the lower frequency words, which the frequency transform tends to counteract.

The effects of the frequency compression must also be considered in light of the properties of the learning algorithm we used,

which is an error correcting learning procedure. This means that changes in connection strengths are made only to the degree that the network fails to match the target. It follows that the magnitudes of the changes tend to diminish with successive presentations of a word. The data to be presented indicate that the model reached nearly asymptotic performance on higher frequency words with less than 250 presentations; thus, additional presentations would have little effect. The net result is that the network itself effectively compresses the effects of frequency as it learns in any case. Where the compression in the frequency range does have an effect is on the relative speed with which high- and low-frequency words are mastered. Higher frequency words do not reach asymptote as quickly because they are presented less often.

In summary, it seems likely that our compression of the frequency range may distort to some extent the rate of mastery of words of different frequencies. However, several considerations suggest that the effects of this compression are less significant than one might initially suppose. The differences between high- and low-frequency words relevant to the child's experience are actually smaller than the norms suggest. Moreover, given the properties of the corpus we have used in these simulations, some compression of the frequency range seems appropriate. In the final section of this article, we also present data from an additional simulation indicating that the model's performance replicates when a broader range of frequencies is used.

We should stress that the model represents a claim about the types of knowledge that are acquired, but it is not a simulation of the child's experience in learning to read in the American educational system. In the model, all of the words are available for sampling throughout training, with frequency modeled by the probability of being selected on a given learning trial. In actual experience, however, frequency derives in part from age of exposure; words that are higher frequency for adults tend to be introduced earlier than are lower frequency items. In learning to read, then, words are introduced sequentially, and often in groups that emphasize salient aspects of the orthography. As shown later, however, the model nonetheless exhibits some of the basic developmental trends characteristic of the acquisition process.

Results

Pronunciation of Written Words

We consider first the model's account of the task of naming written words aloud. Words vary in terms of variables such as frequency of occurrence, orthographic redundancy and ortho-graphic–phonological regularity. Many studies have investigated the effects of these variables on naming performance (see Barron, 1986; Carr & Pollatsek, 1985; Patterson & Coltheart, 1987; Seidenberg, 1985b, for reviews). The basic research strategy has been to examine performance in naming words that differ systematically in terms of these structural variables. The central observation is that even among very skilled readers, there are differences among words in terms of ease of pronunciation. We now consider whether the model's performance on different types of words is comparable to that of people.

Phonological Output and Naming

Before characterizing the model's performance, it is necessary to consider further a theory of the naming task and how it relates to the output computed by the model. We assume that overt naming involves three cascaded processes: (a) the input's phonological code is computed, (b) the computed phonological code is compiled into a set of articulatory–motor commands, and (c) the articulatory motor code is executed, resulting in the overt response. Only the first of these processes is implemented in the model. In practice, however, the phonological output computed by the model is closely related to observed naming latencies.

A word is named by recoding the computed phonological output into a set of articulatory motor commands, which are then executed. Differences in naming latencies primarily derive from differences in the quality of the computed phonological output. Informally speaking, a word that the model "knows" well produces phonological output that more clearly specifies its articulatory–motor program than a word that is known less well. Thus, naming latencies are a function of phonological error scores, which index differences between the veridical phonological code and the model's approximation to it. Clearly, the computed phonological code and the compiled articulatory–motor program are closely related, which is why the error scores systematically relate to observed naming latencies. That the codes are distinct is suggested by evidence that subjects are able to use phonological information even when compilation of the articulatory–motor program is blocked by performance of a secondary articulatory task. For example, subjects can reliably judge phonological properties of stimuli when they are simultaneously mouthing a nonsense syllable (Besner & Davelaar, 1982). Other models have also distinguished between phonological and articulatory codes (e.g., LaBerge & Samuels, 1974).

Differences in naming latencies could also be associated with the execution of the compiled articulatory–motor programs. Consider, for example, a factor such as frequency. The distributions of phonemes in high- and low-frequency words differ; some phonemes and phoneme sequences occur more often in higher frequency words than low, and vice versa (Landauer & Streeter, 1973). Phonemes also differ in terms of ease of articulation (Locke, 1972); higher frequency words may contain more of the phonemes that are easier to pronounce, or it may be that the phonemes that are characteristic of high-frequency words are easier to pronounce because they are used more often. Thus, naming latencies for high- and low-frequency words could differ not because frequency influences the computation of phonological output or the translation of this output into an articulatory code, but because they contain phonemes that differ in terms of ease of articulation. We have ignored this aspect of the naming process for two reasons. First, we have not implemented procedures for producing articulatory output. More important, existing studies indicate that effects of variables such as frequency and orthographic–phonological regularity obtain even when articulatory factors are carefully controlled. For example, there are frequency effects even when articulatory factors are controlled by using homophones (e.g., high frequency: MAIN; low frequency: MANE; see McRae, Jared, & Seidenberg, in press; Theios & Muise, 1977). Among the

monosyllabic words under consideration, differences at the stage of producing articulatory–motor output contribute very little to observed naming latencies (see also Monsell, Doyle, & Haggard, 1989). In sum, naming latencies depend in part on factors related to the construction of an articulatory–motor program and its execution, processes the model does not simulate. It turns out, however, that we can give a fairly accurate account of a broad range of naming phenomena simply in terms of the computation from orthography to phonology.

In the sections that follow, we examine how the model performed on different types of words that were used in behavioral studies. Because the model was trained on a large set of words, we can examine the model's performance on the same items that were used in specific experiments. We evaluate the model's performance in the following way. Given a particular input string, the model produces a pattern of activation across the phonological units. We characterize this pattern by comparing it to different target patterns. For example, we can calculate an error score that reflects the difference between the obtained pattern and the one associated with the correct phonological code for the input string. We can also compare the output to other plausible phonological codes; for example, if the input were an exception word such as HAVE, we can compare the computed pattern of activation to the pattern for both the correct phonological code, /hav/, and the output for a plausible alternative, such as the regularized pronunciation /hAv/.

For the entire set of words after 250 learning epochs, the following results obtained. In general, the error scores calculated using the correct phonological codes as targets were much smaller than the error scores derived by using other targets. In order to be certain that the best fit to the computed output for a given word was the correct phonological code, it would be necessary to compare the output to all possible phonological patterns, which we have not done for obvious reasons. However, the following analysis provides a general picture of the model's performance. The phonological output computed for each word was compared to all of the target patterns that could be created by replacing a single phoneme with some other phoneme. For the word HOT, for example, the computed output was compared to the correct code, /hot/, and to all of the strings in the set formed by /Xot/, /hXt/, and /hoX/, where X was any phoneme. We then determined the number of cases for which the best fit (smallest error score) was provided by the correct code or one of the alternatives.

Among the 2,897 words in the corpus, there were 77 cases (2.7%) in which the best fit to the computed output was a pattern other than the correct one. The errors, which are listed in Tables 1 and 2, were of several types. The model produced 14 regularization errors, in which a word with an irregular pronunciation is given a "regular" pronunciation. These errors are also observed in children learning to read (Backman, Bruck, Hébert & Seidenberg, 1984) and in certain cases of dyslexia following brain injury (Marshall & Newcombe, 1973; Patterson et al., 1985). Thus, although the model was trained that the correct pronunciation of BROOCH is /brOC/, the best fit to the computed output was provided by the regularization /brUC/, similar to BROOM. For PLAID, the model produced /plAd/ instead of /plad/, and for SPOOK, it produced /spuk/ (as in BOOK) instead of /spUk/. All of the regularization errors were produced for words that occurred with very low frequencies during the training phase. In these cases, the model's output was determined on the basis of knowledge derived from exposure to other words, for which the regular spelling–sound correspondences predominate. These errors illustrate a basic characteristic of the model, the fact that the output for a word is affected by exposure to both the word itself and other words. This aspect of the model is discussed in greater detail later.

There were 25 other cases in which the model produced incorrect vowels that were not regularizations. For example, the best fit to BEAU was /bU/, and the best fit to ROMP was /ramp/. Vowels account for the bulk of the errors because they are the primary source of spelling–sound ambiguity in English. There were also 24 cases in which the model produced incorrect consonants. Some of these errors are systematic; for example, the model produced hard Gs instead of soft ones for the words GEL, GIN, and GIST (it performed correctly on other such words, including GENE and GEM, however). Finally, one other type of error occurred because some target pronunciations specified in the training list were miscoded by the experimenter. For example, the pronunciation of SKULL was incorrectly coded as /skull/; in our encoding scheme, the correct code is /skAl/. It is interesting that in 5 cases, the best fit to the computed output was the correct code rather than the one used in training; for JAYS, for example, the model was trained on the incorrect pronunciation /jAs/, but the best fit was provided by the correct code /jAz/. These self-corrections were based on knowledge derived from exposure to related words, such as DAYS.

This analysis of the errors should not be taken as comprehensive because it only tests the computed output against the set of codes containing the same number of phonemes as the target; hence, it does not reveal cases in which phonemes were deleted or added from the target pattern. Inspection of other cases, however, suggests that the model produced few errors of these types. Consider, for example, words containing silent letters, such as DEBT and CALM. We tested the computed phonological output for these words against both the correct pronunciations and the "regularizations" that would occur by pronouncing the silent letters. We found no cases in which the regularized pronunciation yielded a smaller error score. Thus, it appears that in a very high percentage of cases the best fit to the computed output was provided by the correct phonological code, and the number of errors was small.

Among cases in which the best fit was the correct code, the error scores varied, indicating that the model's response was not equally strong for all of the correct items. This, of course, parallels the finding that human subjects pronounce some words more quickly, or with greater accuracy under time pressure, than others. Our main concern is to relate the magnitudes of the error scores computed after 250 epochs of training to the naming latencies obtained in behavioral studies. The simulations reported later compare naming latencies for the words used in particular studies with the error scores for these items. In general, naming latencies are monotonically related to error scores; in most of the simulations, latencies are about 10 times the error score plus a constant of 500–600 ms. The constant varies from experiment to experiment, and we take it to reflect experiment-specific factors such as the quality of the stimulus

display, sensitivity in the voice key used, and other factors that influence the overall speed of the subjects.[5]

Frequency Effects

We begin by considering simple effects of word frequency on naming latency. In general, common, familiar words yield faster naming latencies than do uncommon, less familiar words (e.g., Forster & Chambers, 1973; Frederiksen & Kroll, 1976). The standard interpretation of these effects is that they reflect processes involved in lexical access (i.e., access to entries stored in the mental lexicon). Each vocabulary item is thought to have a frequency-coded entry in the mental lexicon; recognition involves accessing the appropriate entry. In Morton's (1969) model, the entries were called *logogens*, and frequency was encoded by their resting levels of activation (see McClelland & Rumelhart, 1981, for a similar proposal). Balota and Chumbley (1985) also observed small frequency effects that were not attributable to lexical access because they occurred even when subjects had more than 1 s to prepare their responses. These effects were thought to be due to processes involved in producing articulatory–motor output.

Table 1
Corpus of True Errors

Word	Output	Word	Output
Regularizations (n = 14)		Other vowel errors (continued)	
ACHE	AC	PLUME	plom
BROOCH	brUC	QUALMS	kwAlmz
CROW	krw	QUARTZ	kw^rts
DOSE	doz	QUEUE	kwU
DOUSE	dwz	ROMP	ramp
DROUGHT	dr*t	STARVE	starv
PLAID	plAd	SWARM	swlrm
SOOT	sUt	WONT	w^nt
SPA	spa		
SPOOK	spuk	Consonant errors (n = 24)	
SUEDE	swEd	ANGST	ondst
SWAMP	swamp	BREADTH	brebT
WASP	wasp	CORPSE	kOrts
WOMB	wOm	CYST	sist
		CZAR	vor
Other vowel errors (n = 25)		DREAMT	dremp
ALPS	^lps	EWE	wU
BEAU	bU	FEUD	flUd
BLITHE	bliT	GARB	gorg
BRONZE	branz	GEL	gel
CHEW	cw	GIN	gin
DRAUGHT	draet	GIST	gist
SCARCE	skers	HEARTH	hors
SCOUR	skAr	NERSE	mers
FRAPPE	frlp	NYMPH	mimf
FROST	fr^st	PHAGE	pAj
KNEAD	nAd	SPHINX	spinks
LEWD	lEd	SVELTE	swelt
MAUVE	mav	TAPS	tats
MOW	ml	THWART	Twert
NONCE	nans	TSAR	tar
OUCH	AC	WALTZ	w*lps
PLEAD	plAd	WARP	wOrb
PLUME	plom	ZIP	vip

Table 2
Corpus of Coding Errors (n = 14)

Word	Coded as	Output
CHAISE[a]	cez	CAz
DANG	dAng	dang
DAUNT[a]	dwnt	d*nt
FOLD	rold	DOld
SKULL	skull	skulk
JAYS[a]	jAs	jAz
MEW[a]	mYu	myU
PROWL	prowl	prwwl
SHOOT[a]	sUT	sUt
STRODE	strOs	strOz
SWATH	swoth	swoCh
VELDT	veldt	velvt
WOW	www	wwl
ZOUNDS[a]	zWnds	zWndz

[a] Self-corrections.

Our model differs from these kinds of accounts in a fundamental way: It contains no lexicon in which there are entries for individual words; hence, they cannot be "accessed," and there is no direct record of word frequencies. Instead, knowledge of words is encoded in the connections in the network. Frequency affects the computation of the phonological code because items that the model has encountered more frequently during training have a larger impact on the weights. Higher frequency words tend to produce phonological output that more closely approximates the veridical pattern of activation, yielding smaller error scores. As noted earlier, we have assumed that the more closely the computed phonological code corresponds to the veridical code, the easier it will be to compile the code into a sequence of articulatory–motor commands. Thus, frequency has important effects on the computation of the phonological code and therefore on the time it takes to produce an overt response. Although we have not implemented the process, frequency should also affect the computation that takes the phonological code into a set of articulatory–motor commands; McRae et al. (in press) have provided evidence concerning the scope of these effects.

Orthographic–Phonological Regularity

Consider next the contrast between regular words such as MUST, LIKE, and CANE, and exception words such as HAVE, SAID, and LOSE. Regular words contain spelling patterns that recur in a large number of words, always with the same pronunciation. MUST, for example, contains the ending -UST; all monosyllabic words that end in this pattern rhyme (JUST, DUST, etc.). The words sharing the critical spelling pattern are called the *neighbors* of the input string (Glushko, 1979). Neighbors have

---

[5] The simulations reported here involve comparisons between subjects' naming latencies and the model's performance on the same items. The naming latencies presented in the figures sometimes differ slightly from those reported in the original articles because some experiments included a small number of words that were not contained in the training set. Excluding these items did not alter the patterns of results in any of the experiments.

Table 3
Mean Naming Latencies and Percentage Errors

Type	Example	Latency	Errors
High frequency			
Regular	NINE	540	0.4
Exception	LOSE	541	0.9
Low frequency			
Regular	MODE	556	2.3
Exception	DEAF	583	5.1

Note. Data are from the Seidenberg (1985c) experiment.

been defined in terms of word endings, also called *rimes* (Trieman & Chafetz, 1987) or *word bodies* (Patterson & Coltheart, 1987), although as we shall see other aspects of word structure also matter (Taraban & McClelland, 1987). Exception words contain a common spelling pattern that is pronounced irregularly. For example, -AVE is usually pronounced as in GAVE and SAVE, but has an irregular pronunciation in the exception word HAVE. In terms of orthographic structure, regular and exception words are similar: Both contain spelling patterns that recur in many words. It is often said that regular words obey the pronunciation "rules" of English, whereas exception words do not. Thus, these types of words are similar in terms of orthography, and they can be equated in terms of other factors such as length and frequency. Differences between them in terms of processing difficulty must be attributed to the one dimension along which they differ, regularity of spelling–sound correspondences.

The studies examining the processing of such words have yielded the following results. As noted previously, there are frequency effects; higher frequency words are named more quickly than lower frequency words. In addition, regularity effects—faster latencies for regular words compared to exceptions—are larger in lower frequency items and are small or nonexistent in higher frequency words (Andrews, 1982; Seidenberg, 1985c; Seidenberg, Waters, Barnes, & Tanenhaus, 1984; Taraban & McClelland, 1987; Waters & Seidenberg, 1985). In short, there is a Frequency × Regularity interaction, as exemplified by the results from Seidenberg (1985c) presented in Table 3.

The number of higher frequency items for which irregular spelling–sound correspondences have little impact on overt naming is likely to be rather large because of the type/token facts about English (Seidenberg, 1985c). A relatively small number of word types account for a large number of the tokens that a reader encounters. In the Kucera and Francis (1967) corpus, for example, the 133 most frequent words in the corpus account for about one half of the total number of tokens. Hence, a small number of words recur with very high frequency, and for these words spelling–sound irregularity has little effect. Exception words tend to be overrepresented among these higher frequency items, largely due to the fact that the pronunciations of higher frequency words are more susceptible to diachronic change (Hooper, 1977; Wang, 1979). It is interesting to note that although written English is said to be highly irregular, the irregular items tend to cluster in the higher frequency range, in which this property has negligible effects on processing. Finally, the size of this higher frequency pool varies as a function of

reading skill. Seidenberg (1985c) partitioned the data in Table 3 according to overall subject naming speed, yielding fast-, medium-, and slow-reader groups (Table 4). Among these subjects, who were McGill University undergraduates, the fastest readers named *lower* frequency words more rapidly than the slowest readers named *higher* frequency words, and thus showed no regularity effect even for the lower frequency items. Thus, faster readers recognize a larger pool of items without interference from irregular spelling–sound correspondences. In effect, more words are treated as though they are high-frequency items; this may be an important source of individual differences in reading skill.

To examine the model's performance on these types of words, we used a somewhat larger stimulus set studied by Taraban and McClelland (1987, Experiment 1). Figure 3 presents the model's performance on this set of high- and low-frequency regular and exception words after different amounts of training. Each data point represents the mean phonological error score for the 24 items of each type used in the Taraban and McClelland experiment. The learning sequence is characterized by the following trends. Training reduces the error terms for all of the words following a negatively accelerated trajectory. Throughout training, there is a frequency effect: The model performs better on the words to which it is exposed more often. Note that although the test stimuli are dichotomized into high- and low-frequency groups, frequency is actually a continuous variable, and it has continuous effects in the model. Early in training, there are large regularity effects for both high- and low-frequency items; in both frequency classes, regular words produce smaller error terms than do exception words. Additional training reduces the exception effect for higher frequency words, to the point where it is eliminated by 250 epochs. However, the regularity effect for lower frequency words remains.

Taraban and McClelland's (1987) adult subjects performed as follows. First, lower frequency words were named more slowly than higher frequency words. Second, there was a Frequency × Regularity interaction; exception words produced significantly longer naming latencies than regular words only when they were low in frequency. For lower frequency words, the difference between regular and exception words was 32 ms, which was statistically significant; for higher frequency words,

Table 4
Mean Naming Latencies as a Function of Decoding Speed

Word type	Subject group		
	Fastest	Medium	Slowest
High frequency			
Regular	475	523	621
Exception	475	517	631
Difference	0	−6	+10
Low frequency			
Regular	500	530	641
Exception	502	562	685
Difference	+2	+32	+44

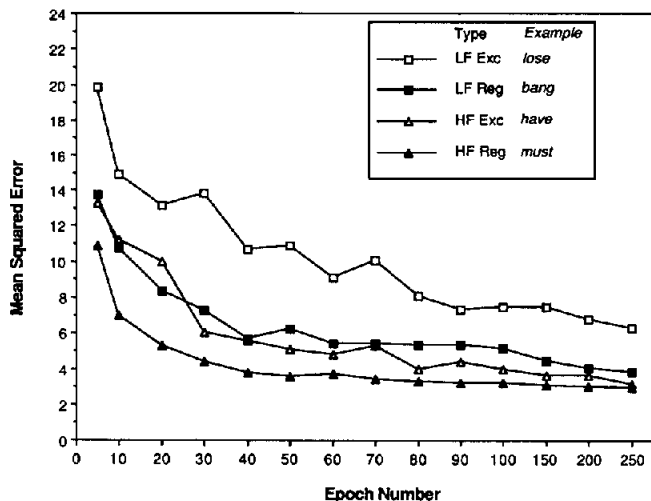*Note.* Numbers are in milliseconds.

Figure 3. Mean phonological error scores for the stimuli used by Taraban and McClelland (1987).

the difference was 13 ms and nonsignificant. The model produced similar results, as indicated in Figure 4.

Figure 5 shows two additional studies of this type, using slightly different stimulus sets. The Seidenberg (1985c, Experiment 2) data summarized in Table 3 are presented on the left; the results of Seidenberg, Waters, Barnes, and Tanenhaus (1984, Experiment 3) are on the right. The model's performance on the same stimulus words is also presented. In each case, both experiment and simulation yielded Frequency × Regularity interactions, with a good fit between the two.

In Figure 6 we summarize the results of 14 conditions from 8 experiments that examined differences between regular and exception words. The data represent the mean differences between exception words and regular words obtained in the experiments and in simulations using the same items. For conditions a–e, the differences between the naming latencies for regular and exception words were not statistically significant (these were higher frequency stimuli); the model also produced very small effects in these cases. In the remaining conditions, which yielded significant effects, the model also produces larger differences between the two word types. The correlation between experiment and simulation data is .915.

The simulation is revealing about the behavioral phenomena in two respects. First, it is clear that in the model, the Frequency × Regularity interaction occurs because the output for both types of higher frequency words approaches asymptote before the output for the lower frequency words. Hence, the difference between the higher frequency regular and exception words is eliminated, whereas the difference between the two types of lower frequency words remains. This result suggests that the interaction observed in the behavioral data results from a kind of "floor" effect due to the acquisition of a high level of skill in decoding common words. In the model, the differences between the two types of lower frequency words would also diminish if training were continued for more epochs. This aspect of the model provides an explanation for Seidenberg's (1985c) finding that there are individual differences among skilled read-

ers in terms of regularity effects. As Table 4 indicates, the fastest subjects in this study showed no regularity effect even for words that are lower in frequency according to standard norms. The model suggests that these subjects may have encountered lower frequency words more often than did the slower subjects, with the result that they effectively become high-frequency items.

Second, the model provides an important theoretical link between effects of frequency and regularity. Both effects are due to the fact that connections that are required for correct performance have been adjusted more frequently in the required direction for frequent and regular items than for infrequent or irregular items. This holds for frequent words simply because they are presented more often. It holds for regular words because they make use of the same connections as other, neighboring regular words. Hence, both frequency and regularity effects derive from the same source, the effects of repeated adjustment of connection weights in the same direction.

Performance on Other Stimulus Types

Several other types of words have been studied in naming experiments; research in this area has been marked by the devel-
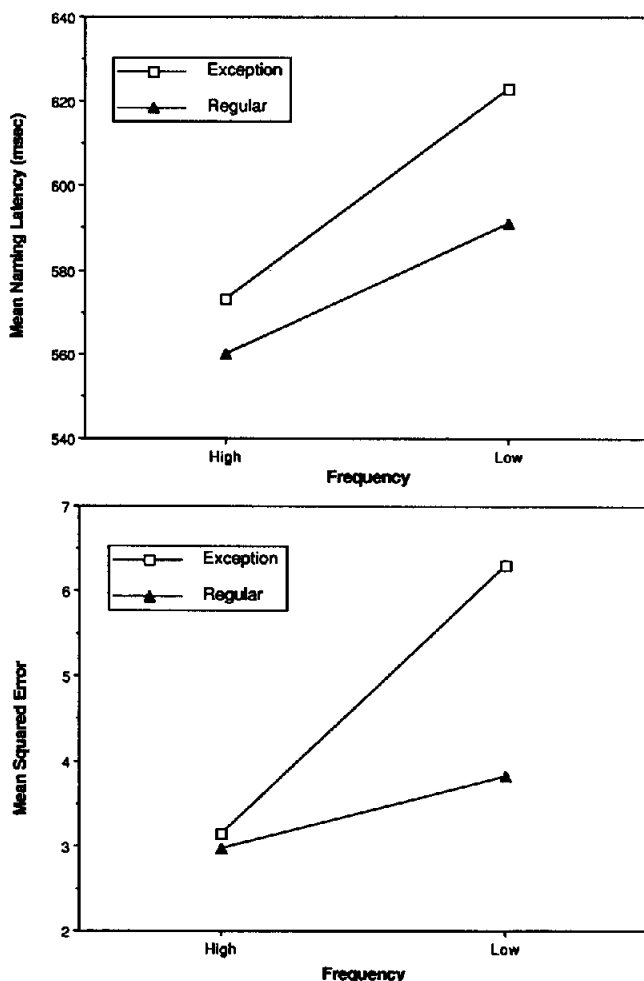

Figure 4. Results of the Taraban and McClelland (1987) study (upper graph) and the simulation data for 250 epochs (lower graph).
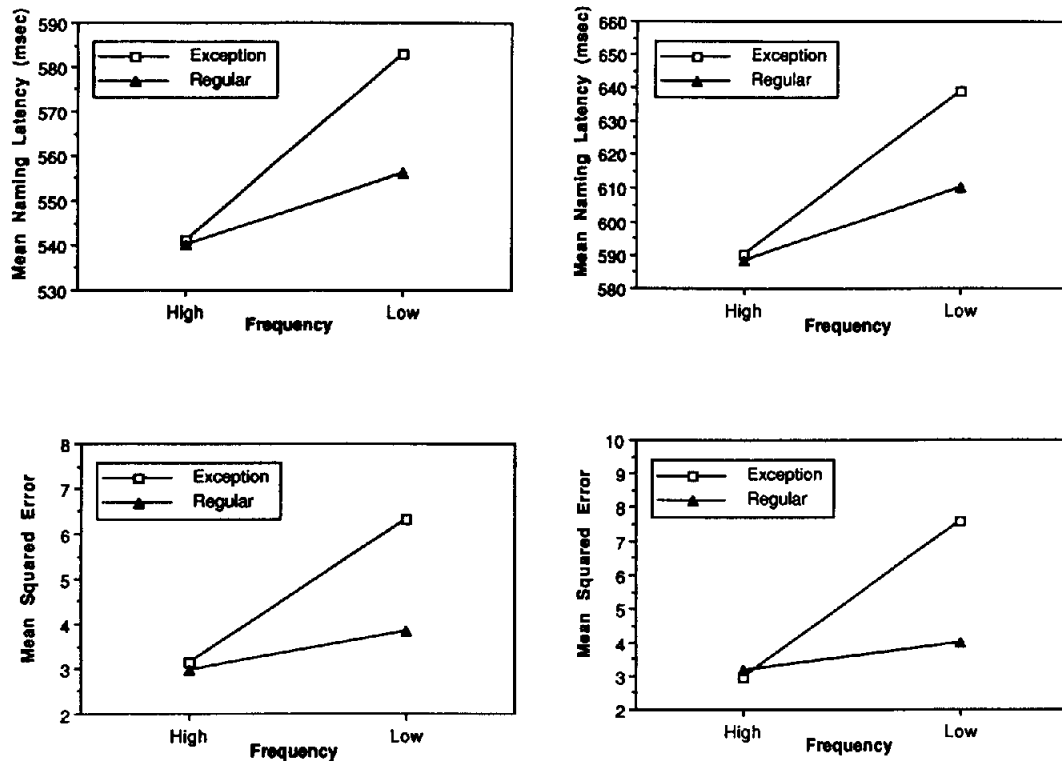
Figure 5. Results of the Seidenberg (1985c; left graphs) and Seidenberg, Waters, Barnes, & Tanenhaus (1984, Experiment 3; right graphs) studies: experiments (upper graphs) and simulations (lower graphs).

opment and revision of several taxonomies based on different properties of words or perceptual units thought to be theoretically relevant. In part, this research was motivated by the fact that several models, incorporating very different representational and processing assumptions, all predict longer naming latencies for exception words compared with regular words. In the dual-route model (Coltheart, 1978), longer latencies result because readers attempt to pronounce exception words by applying grapheme–phoneme correspondence rules, resulting in a temporary misanalysis. In Glushko's (1979) model, a word is pronounced by analogy to similarly spelled neighboring words. The fact that the neighbors of an exception word are all regular was thought to interfere with generating its pronunciation. According to Brown (1987), the factor that determines naming latencies is the number of times a spelling pattern (word body) occurs with a particular pronunciation. A regular word such as DUST contains a word body, -UST, that is pronounced /ust/ in many words. An exception word such as SWAMP contains a word body, -AMP, that is pronounced /omp/ in only one word, the exception itself. Hence, the frequency of a spelling–sound correspondence could be the source of the exception effect.

In the following sections, we consider the model's performance on several additional types of words and nonwords, showing that it closely simulates the behavioral data. We then consider the principles that govern the model's performance and compare them with ones in other models.

Regular inconsistent words. In an important article, Glushko (1979) studied a class of words called *regular inconsistent.*

These words, such as GAVE, PAID, and FOE, have two critical properties. Their pronunciations can be derived by rule; in fact, most of these words' neighbors rhyme (e.g., GAVE, PAVE, SAVE, BRAVE). However, each of these words has an exception word neighbor (e.g., HAVE, SAID, and SHOE, respectively). The view that readers pronounce words by applying spelling–sound rules predicts that regular inconsistent words should be named as quickly as regular words, other factors being equal; in both cases, the rules generate the correct pronunciations. Glushko (1979) proposed that words are pronounced by analogy to similarly spelled words, affording the possibility that pronunciation of a regular inconsistent word such as GAVE could be influenced by knowledge of an exception word such as HAVE. He reported experimental evidence that regular inconsistent words yield longer naming latencies than do regular words; he also found that nonwords derived from exception words (e.g., BINT from PINT) yielded longer latencies than nonwords derived from regular words (e.g., NUST from MUST). These findings have been taken as strong evidence against dual-route models (e.g., Henderson, 1982).

Subsequent studies of regular inconsistent words have yielded mixed results. Seidenberg et al. (1984b, Experiment 4) obtained the regular inconsistent effect only for lower frequency words, and several studies failed to yield statistically reliable effects at all (e.g., Seidenberg et al., 1984b, Experiment 1; Stanhope & Parkin, 1987; Taraban & McClelland, 1987). These mixed results suggest that the mere presence or absence of an exception word neighbor is not the only factor relevant to processing, an

Figure 6. Results of 14 conditions (a–n) examining exception effects: Experiment and simulation data. (Key: a = Seidenberg [1985c], Experiment 2, Set A, HF words; b = Seidenberg, Waters, Barnes, & Tanenhaus [1984], Experiment 3, HF words; c = Waters & Seidenberg [1985], Experiment 1, HF words; d = Taraban & McClelland [1987], Experiment 1, HF words; e = Seidenberg [1985c], Experiment 2, Set B, HF words; f = Glushko [1979], Experiment 3; g = Brown [1987], Experiment 1; h = Seidenberg [1985c], Experiment 2, Set B, LF words; i = Glushko [1979], Experiment 1; j = Seidenberg, Waters, Barnes, & Tanenhaus [1984], Experiment 3, LF words; k = Taraban & McClelland [1987], Experiment 1, LF words; l = Seidenberg [1985c], Experiment 1, Set A, LF words; m = Waters & Seidenberg [1985], Experiment 1, LF words; n = Seidenberg, Waters, Barnes, & Tanenhaus [1984], Experiment 1. HF = high frequency; LF = low frequency.)

issue to which we return later. We examined the model's processing of regular inconsistent words using stimuli from the Taraban and McClelland experiment described previously, which also included high- and low-frequency regular inconsistent words and matched regular word controls. This represents the largest set of regular inconsistent words used in any experiment. There were again 24 items of each type, all of which were included among the 2,897 words in our training set. Figure 7 shows the model's performance on these words after different amounts of training. Error scores again decreased with additional training, and higher frequency words again produced lower error scores than lower frequency words. However, after 250 epochs, there were only small differences between regular inconsistent words and regular words in both frequency ranges (high frequency: 0.0077; low frequency: 0.3128). These data are consistent with Taraban and McClelland's results; the differences between regular inconsistent words and regular controls in their experiment were 7 ms and 10 ms, respectively, for the high- and low-frequency items. Neither difference was statistically reliable. For comparison, note that the difference between lower frequency regular and exception words in their experiment was 32 ms, and 2.4804 in the simulation.

Seidenberg et al. (1984b) identified an aspect of Glushko's (1979) methodology that may have been responsible for the large regular inconsistent effect in his study. Glushko's experiment included matched exception/regular inconsistent pairs such as BEEN–SEEN, GIVE–DIVE, and NONE–CONE. Each spelling pattern in the stimulus list occurred at least twice with two different pronunciations; some spelling patterns were repeated

several times (e.g., the stimuli included NONE, CONE, GONE, DONE, SHONE, and BONE). Repetition of spelling patterns with different pronunciations may have introduced intralist priming effects that would tend to increase the magnitude of the regular inconsistent/regular difference. Seidenberg, Waters, Barnes, and Tanenhaus (1984, Experiment 2) showed that a large regular inconsistent effect occurs when stimuli are repeated in this way, but not when the stimuli are not repeated. The model provides additional support for this conclusion. We tested the model on the items from Glushko's Experiment 3, which had yielded a significant 17-ms difference between regular inconsistent and regular words. The model yielded a negligible difference of 0.1247 on the same items. The basis for this difference is clear: Unlike human subjects, the model's performance during testing is not influenced by previous trials. The model is tested on each stimulus without changing the weights in any way; hence, there are no intralist priming effects.

We consider the regular inconsistent words again later because they are theoretically important and because the studies examining these items did not control another important aspect of their structure. Here it is sufficient to note that the model gives a good account of the behavioral data obtained in studies using these words.

Strange words. Several studies (e.g., Parkin, 1982; Parkin & Underwood, 1983; Seidenberg, Waters, Barnes, & Tanenhaus, 1984; Waters & Seidenberg, 1985) have examined words that differ from the regulars, regular inconsistents, and exceptions in a basic way: They contain spelling patterns that occur in a very small number of words, often only one. Regular patterns such as -UST and inconsistent patterns such as -AVE are productive in the sense that they are realized in many words. Words such as GUIDE, AISLE, and FUGUE contain nonproductive spelling patterns that rarely occur in other words. For example, GUIDE is the only monosyllabic word ending in -UIDE. Henderson (1982) calls these words *lexical hermits;* in Glushko's (1979) terminology, they have few if any immediate neighbors. These words might be expected to be difficult to pronounce for three reasons: first, because they contain relatively unfamiliar spelling patterns and thus are low in terms of orthographic re-


Figure 7. Model's performance on regular inconsistent and regular words used in the Taraban and McClelland (1987) study.
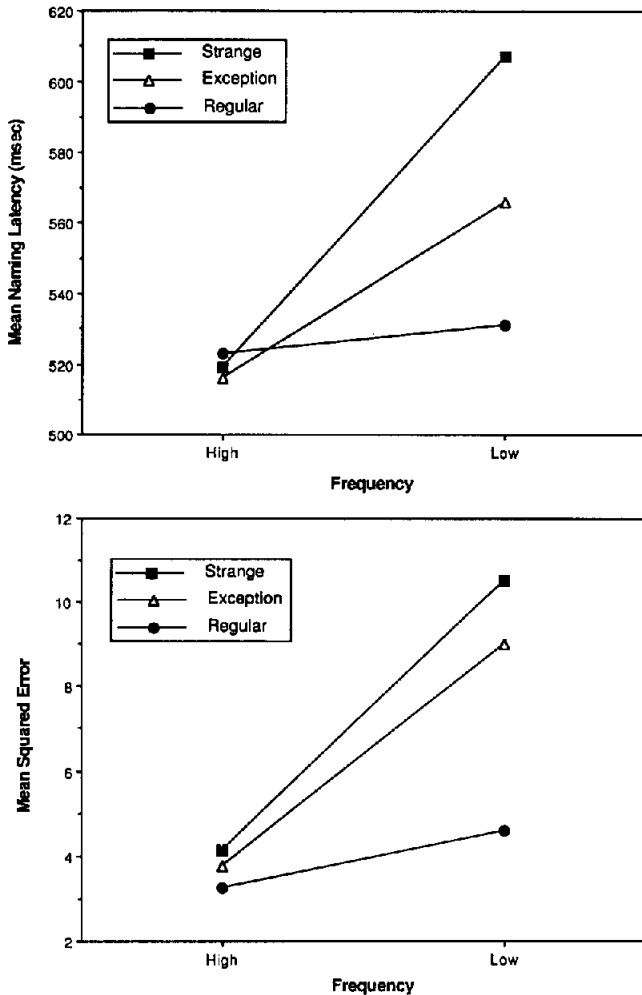
Figure 8. Results of the Waters and Seidenberg (1985) studies: Experiment (upper graph) and simulation (lower graph).

dundancy, a factor that would slow the identification of component letters; second, because the spelling-to-sound correspondences of these patterns are also relatively unfamiliar; and third, because these unusual spelling patterns are often associated with idiosyncratic pronunciations (as in CORPS).

Waters and Seidenberg (1985) compared the naming latencies for a set of these words (which they termed *strange*) with the latencies for regular and exception words. The words were again dichotomized into high- and low-frequency groups. Results of this study are presented in Figure 8. Among the higher frequency words, there were no reliable differences between word classes; for the lower frequency words, the ordering of latencies was strange > exception > regular. Strange words also produced a larger number of mispronunciation errors. The model's performance on these words is also presented in Figure 8, and shows the same interaction between frequency and word class. The results corroborate the conclusion that for higher frequency words, variations in word structure, such as the frequency of a spelling pattern or spelling–sound correspondence, have little impact on naming. Despite the various ways in which

regular, regular inconsistent, exception, and strange words differ, they yield similar naming latencies in this frequency range. Among the lower frequency words in the language, the strange items are the most difficult to name.

Unique words. We also tested the model on a set of words used by Brown (1987), who introduced another category of items, termed *unique*. These are words such as SOAP or CURVE that also contain word bodies that do not occur in other monosyllabic words. These words are somewhat less eccentric than the strange words mentioned earlier, as indicated by the fact that they produce lower orthographic error scores, which are a measure of orthographic redundancy (see discussion on p. 552). Brown also examined exception words such as LOSE and regular words such as MILL, which he termed *consistent*. The stimuli were used to examine the hypothesis that the factor critical to naming is the number of times a word body is associated with a given pronunciation. Both unique and exception words contain spelling patterns assigned a given pronunciation in only a single word (namely, the unique or exception item itself), whereas regular words contain word bodies associated with a given pronunciation in many words. Hence, Brown predicted that unique and exception words should yield similar naming latencies, and both should be slower than regular words. Data from Brown's naming experiment and the simulation are presented in Figure 9. Clearly, the fit between the two is very good.

Neighborhood size. Andrews (in press) reported a study that factorially varied word frequency and a measure of neighborhood size known as Coltheart's N (Coltheart, Davelaar, Jonasson, & Besner, 1977), which refers to the number of words that can be derived from a given word by changing one letter. There were 15 words in each of the four classes formed by crossing frequency (high, low) and neighborhood size (large, small). Results of the experiment and simulation are presented in Figure 10, with again a very good fit between the two. Both Andrews's data and the model suggest that as the frequency of a word increases, the effects of neighboring words diminish.

Nonword pronunciation. After training, the model has encoded facts about orthographic–phonological correspondences in the weights on connections. Although the model performs better on the training stimuli, it will compute phonological output for novel stimuli. In this respect, it simulates the perfor-
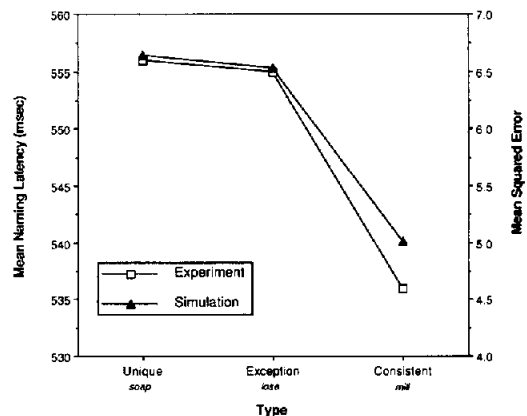

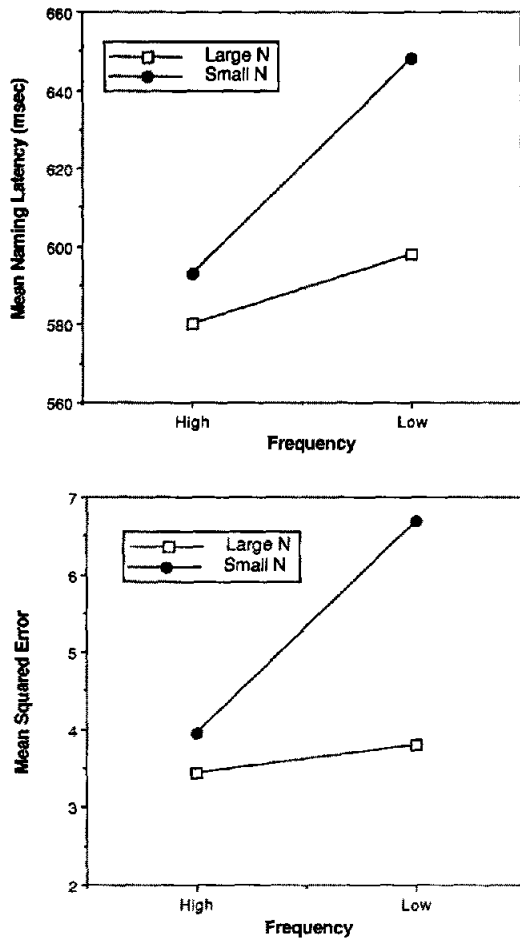Figure 9. Results of Brown (1987): Experiment and simulation data.

Figure 10. Results of Andrews (in press): Experiment and simulation data. (N refers to Coltheart's N, a measure of neighborhood size.)

Figure 12. Model's performance on Taraban and McClelland (1987) exception words. (Error scores for correct [exception] pronunciations and incorrect [regularized] pronunciations.)

mance of subjects asked to pronounce nonwords such as BIST or TAZE. Nonword performance provides important information concerning the naming process because, as we have seen, performance on many words reaches floor levels because of repeated exposure to the items themselves. Because nonwords
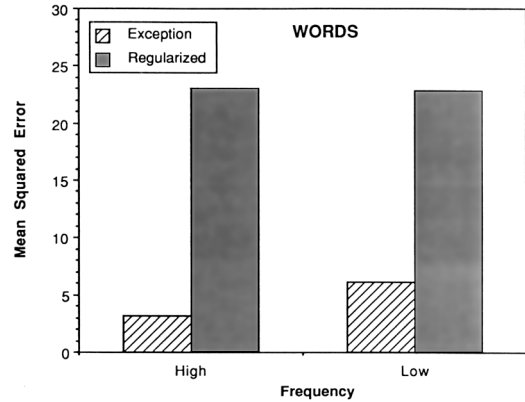

Figure 11. Results of the Glushko (1979) nonword experiment: Experiment and simulation data.

have not been encountered previously, pronunciation must be based on knowledge gained from similar words. A critical experiment was reported by Glushko (1979), who examined naming latencies for nonwords derived from regular words (e.g., NUST derived from MUST) and nonwords derived from exception words (e.g., MAVE derived from HAVE).[6] We tested the model on his set of nonwords; the results from experiment and simulation are presented in Figure 11. In both cases, performance is poorer on the exception nonwords. Note that the nonwords derived from exceptions are in effect "regular inconsistent." Whereas regular inconsistent words show little effect of a neighboring exception word, regular inconsistent nonwords do. The difference, of course, is that the model is actually trained on regular inconsistent words, but not the corresponding nonwords. Apparently, training on the item itself is sufficient to overcome the effect of training on the exception neighbor.

The model was also tested on a set of nonwords derived from the exception words used in the Taraban and McClelland (1987) study. These nonwords can be pronounced in two ways, either by analogy to the exception word (e.g., MAVE pronounced to rhyme with HAVE) or by analogy to a regular inconsistent word (e.g., MAVE rhymed with GAVE). Using the weights from 250 epochs, the model was tested to determine which pronunciation would be preferred. For each item, phonological error scores were calculated twice, using both exception and regular pronunciations as targets. We also calculated analogous scores for alternative pronunciations of the exception words themselves (e.g., HAVE pronounced correctly and pronounced to rhyme with GAVE). This is the regularization error discussed previously.

Figure 12 shows both types of error scores for the exception words in the Taraban and McClelland (1987) stimuli. For words, the correct "exception" pronunciations produce much

---

[6] Glushko's (1979) Experiment 2, which examined nonword naming, did not include repetitions of spelling patterns with different pronunciations; hence, it is not subject to the repetition priming hypothesis previously advanced in connection with his experiment on regular inconsistent words.
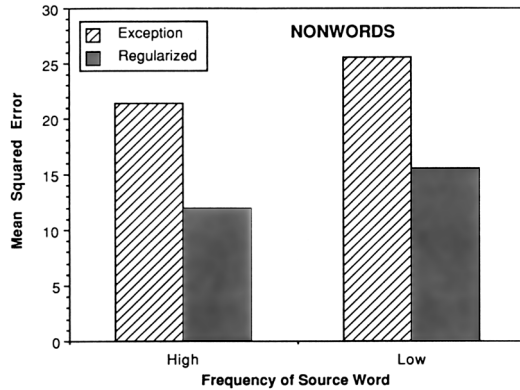
Figure 13. Model's performance on nonwords derived from Taraban and McClelland (1987) exception words. ("Exception" pronunciation rhymed with exception word [e.g., MAVE pronounced like HAVE]; "regularized" pronunciations rhymed with regular inconsistent word [e.g., MAVE pronounced like GAVE].)

smaller error scores than do the incorrect, "regularized" pronunciations. Thus, the model's output resembles the correct pronunciations rather than the regularized ones.

The opposite pattern obtains with the nonwords derived from these stimuli (Figure 13). Here the "regularized" pronunciations are preferred to the pronunciations derived from matched exception words. Note, however, that the difference between the two pronunciations is much smaller than in the corresponding word data, suggesting that the pronunciation of a nonword like MAVE is influenced by the fact that the model has been trained on exception words like HAVE.

Figure 14 shows the error scores for the *regular* pronunciations of nonwords derived from regular and exception words. The error scores are larger for nonwords such as MAVE (derived from an exception word) than for nonwords such as PAME (derived from a regular word). These results also indicate that the pronunciation of novel stimuli such as MAVE is affected by the fact that the model has been trained on both HAVE and regular words such as GAVE.

The model's performance on the nonwords is important for two reasons. First, it shows that performance generalizes to new items; the knowledge that was acquired on the basis of exposure to a pool of words can be used to generate plausible output for novel stimuli. Second, the nonword data provide additional information as to what the model has learned. Regular inconsistent words are little affected by training on exception word neighbors. However, the inconsistency in the pronunciation of -AVE is encoded by the weights, as evidenced by performance on regular inconsistent nonwords.

What the Model Has Learned

We have demonstrated that the model simulates a broad range of empirical phenomena concerning the pronunciation of words and nonwords. Why the model yields this performance can be understood in terms of the effects of training on the set of weights. The values of the weights reflect the aggregate effects of many individual learning trials using the items in the training

set. In effect, learning results in the recreation of significant aspects of the structure of written English within the network. Because the entire set of weights is used in computing the phonological codes for all words, and because all of the weights are updated on every learning trial, there is a sense in which the output for a given word is a function of training on all of the words in the set. Differences between words derive from facts about the writing system distilled during the learning phase. For words, the main influence on the phonological output is the number of times the model was exposed to the word itself. Number of times the model was exposed to closely related words (e.g., similarly spelled items) exerts secondary effects; there are also small effects due to exposure to other words. The magnitudes of these effects vary as a function of how similar these words are to a given item.

To see this more clearly, consider the following experiment. We test the model's performance on the word TINT; with the weights from 250 epochs, it produces an error score of 8.92. We train the model on another word, adjusting the weights according to the learning algorithm, and then retest TINT. By varying the properties of the training word, we can determine which aspects of the model's experience exert the greatest influence on the weights relative to the target. This procedure yields orthographic and phonological priming effects, which have been studied by Meyer, Schvaneveldt, and Ruddy (1974), Hillinger (1980), and Tanenhaus et al. (1980). For example, Meyer et al. observed that lexical decision latencies to a target word such as ROUGH were facilitated when preceded by the rhyme prime TOUGH but inhibited when preceded by the similarly spelled nonrhyme COUGH. For the purposes of the simulation, we examined the cumulative effects of a sequence of 10 prime (learn)—target (test) trials. The primes were a rhyme (MINT), a matched exception word (PINT), a word with the same consonants but a different vowel (TENT), and an unrelated control (RASP). The data are presented in Figure 15.

The results indicate, first, that priming with the orthographically similar rhyme MINT decreases the error for TINT; the overlap between the words is sufficient to improve performance.
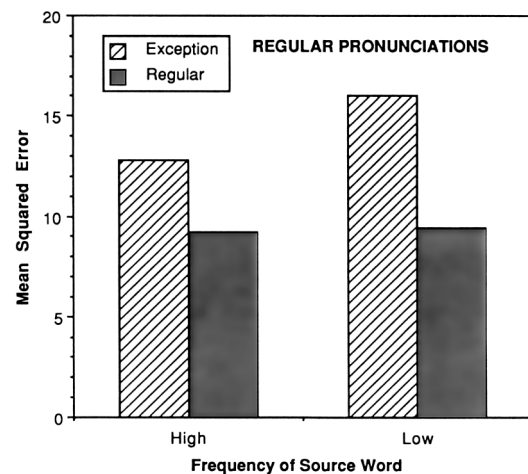

Figure 14. Error scores for regular pronunciations of nonwords derived from regular and exception words.
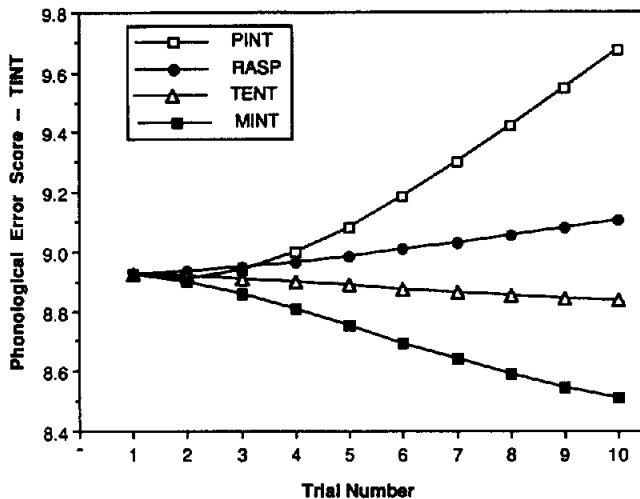
Figure 15. Effects of training on PINT, RASP, TENT, and MINT on phonological error score for TINT.

Other rhymes act in a similar manner. This outcome is consistent with Brown's (1987) proposal that the frequency with which a word body is associated with a given pronunciation influences performance; the number of times the pattern -INT = /int/ occurs in the training set affects performance on TINT. Note, however, that the other primes also have effects. Priming with the similarly spelled nonrhyme TENT also improves performance; the effect is smaller because vowels are the primary source of ambiguity in orthographic–phonological correspondences and, hence, the primary source of error. Training on MINT has a larger facilitating effect because it provides feedback concerning the primary source of ambiguity. The exception word PINT has interfering effects complementary to the facilitative effects of MINT. Finally, the unrelated prime RASP has very small negative effects.

Note that the priming effects illustrated in Figure 15 are not characteristic of all of the words in the training set after 250 epochs of training. TINT is somewhat unusual in that the model's performance is relatively poor, due in part to the fact that TINT is low in frequency and the fact that there are few -INT words in the corpus. There are smaller priming effects for target words that yield smaller error scores. Figure 15 accurately illustrates the influences of training on related words, but these effects are more salient earlier in the training sequence when error scores are larger.

The model clarifies why some effects of word type are obtained in behavioral studies and others are not. When experimenters compare performance on two types of words, they are attempting to observe the net effect of a particular aspect of word structure (e.g., regularity defined in terms of word bodies) against a background of noise provided by the effects of all other properties of the words. For this reason, experimenters routinely attempt to equate stimuli in terms of these other properties (e.g., frequency, length, initial phoneme). There is a net exception effect for lower frequency words because the regular correspondence is encountered many more times than the irregular one; repeated experience with words such as TINT, MINT,

and HINT has a negative impact on the weights from the point of view of PINT. Conversely, exposure to an exception such as PINT tends to have relatively small effects on a regular inconsistent word such as TINT because the exception word is encountered much less often than the set of rhyming regular inconsistent words. It is not that PINT has no effect on TINT; in the priming experiment, the effect was observed once it was magnified through repetition. The effect can also be observed earlier in the training sequence; eventually it recedes into the background provided by exposure to many other words. The model corroborates the common assumption that word bodies are relevant to naming; however, it suggests that other aspects of word structure also matter.

One other point should be noted. We also examined repetition priming, that is, the effects of 10 trials of training on TINT itself. This resulted in a much larger decrease in TINT's error score, from 8.92 to 2.50. As stated previously, the main factor that influences performance on a word is the number of times the model is exposed to the word itself; effects of neighboring words are relatively small. Thus, presenting an exception word such as PINT with much greater frequency would have less effect on TINT than a small number of exposures to TINT itself.

The model's behavior can be further clarified by examining yet another type of word, which contain what Seidenberg, Waters, Barnes, and Tanenhaus (1984) and Backman et al. (1984) called *ambiguous spelling patterns.* These spelling patterns, such as -OWN, -OVE, and -EAR, are associated with two or more pronunciations, each of which occurs in many words (e.g., BLOWN, FLOWN, KNOWN, GROWN; TOWN, FROWN, DROWN, GOWN). For inconsistent spelling patterns such as -INT or -AVE, the number of words with the regular pronunciation greatly exceeds the number of words with the exceptional pronunciation. For the ambiguous spelling patterns, however, the ratio is more nearly equal. Hence, during training, the model is exposed to many examples of each pronunciation. We constructed a set of 24 high-frequency and 24 low-frequency words containing these spelling patterns, matched with the stimuli in the Taraban and McClelland (1987) set in terms of frequency. Mean phonological error scores for these words (using the weights from 250 epochs) and the other stimuli in the Taraban and McClelland experiment, are presented in Figure 16. As before, there are negligible differences between the word types in the higher frequency range. Among the lower frequency words, the ambiguous items yield better performance than the exceptions, but worse performance than the regular inconsistents. Performance is better than on the exceptions because the model receives less training on the exceptional pronunciation than on either pronunciation of the ambiguous spelling pattern. Performance is worse than on the regular inconsistent words because the model is repeatedly exposed to both pronunciations. Thus, there are graded effects of regularity owing to the nature of the input during acquisition.[7]

---

[7] Ambiguous words have been used in only one study of skilled readers (Seidenberg, Waters, Barnes, & Tanenhaus, 1984, Experiment 1). The model simulates the results of this experiment quite closely. However, the ambiguous words were in the higher frequency range, in which they do not differ from regular words. In Backman, Bruck, Hébert, and Seidenberg's (1984) developmental study (described later), children's

Figure 16. Model's performance on Taraban and McClelland (1987) stimuli and on a set of ambiguous words (such as TOWN and LOVE).

Characteristics of the hidden units. Evidence as to how orthographic and phonological information are encoded by the network can be obtained by examining the patterns of activations over the hidden units produced by different words. Unlike the model's orthographic and phonological units, the hidden units do not have specific, predetermined roles. Rather, their representational and functional roles emerge as a result of experience in learning to perform the task that is imposed on the network by the training procedure. Recall that the activation of a hidden unit is a function of the weights on the connections coming into it. At first, each hidden unit has random incoming and outgoing connection strengths. Gradually these are adjusted through experience, so that units come to perform useful, generally partially overlapping parts of the task. Because of the task that these units need to perform—they must allow reconstruction of the orthography as well as construction of the phonology—the values of these weights are affected by feedback concerning both orthography and phonology.

Consider first the pattern of activation over the hidden units produced by the word LINT (Figure 17). LINT activates 23 units, 22 very strongly (net activation > .8) and one more weakly (net activation < .6). We can determine how many of these units are also activated by the orthographically similar rhyme MINT and by the unrelated word SAID. A total of 14 units are activated by both LINT and MINT, and 3 by LINT, MINT, and SAID; 1 unit was activated by both LINT and SAID. The remaining 5 units were "unique" to LINT, in the sense that they were not activated by either MINT or SAID. (Note that the "unique" units were activated by many other words outside this limited set.) Thus, a large number of units apparently reflect the orthographic and phonological similarity of LINT and MINT, and a smaller number are relevant to LINT itself. Fewer units are activated by both LINT and the unrelated word SAID.

performance on ambiguous words was better than on exceptions, but worse than on regular inconsistents. Thus, children show the pattern for lower frequency words shown in Figure 16. The stimuli in this experiment were words that are nominally high-frequency items for adults. As we argue later, younger readers' processing of higher frequency words is like skilled readers' processing of lower frequency words. Hence, the results are consistent with the data in Figure 16.
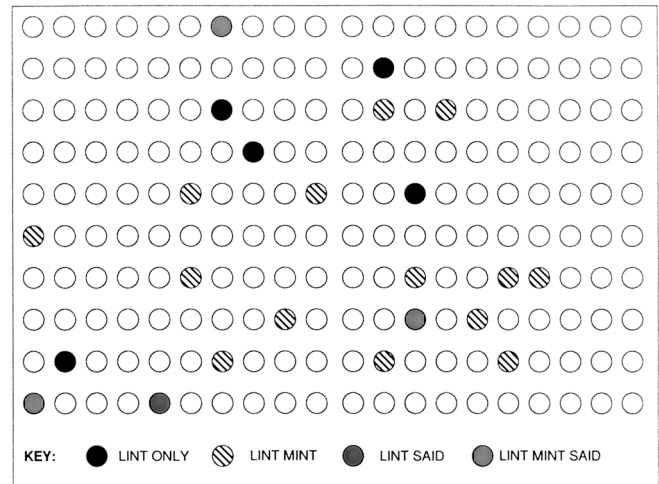
Figure 17. Hidden unit activations for LINT, MINT, and SAID.

This pattern contrasts with the one observed for the exception word PINT (Figure 18). PINT activates 22 units, 8 of which were activated by both PINT and MINT, 1 by PINT and the unrelated word SAID, and 3 by PINT, MINT, and SAID. There were 10 units activated by PINT only. Hence, compared with the pattern for LINT, there is a relatively larger number of units specific to PINT; moreover, the orthographically similar but nonrhyming stimuli LINT and PINT activate fewer units in common than do the orthographically similar, rhyming pair LINT and MINT. Finally, there is very little spurious overlap with an unrelated word such as SAID.

These snapshots of the hidden units indicate that they reflect generalizations concerning the regularities in the lexicon encoded by the weights on connections. Similarly spelled rhymes activate the largest number of common units (LINT/MINT = 14), similarly spelled nonrhymes a smaller number of common units (PINT/MINT = 8), and unrelated words a smaller number still (LINT/SAID and PINT/SAID both = 1). Six units are activated

Figure 18. Hidden unit activations for PINT, MINT, and SAID.

Table 5
Altering the Hidden Units for PINT: *Effects on Error Scores*

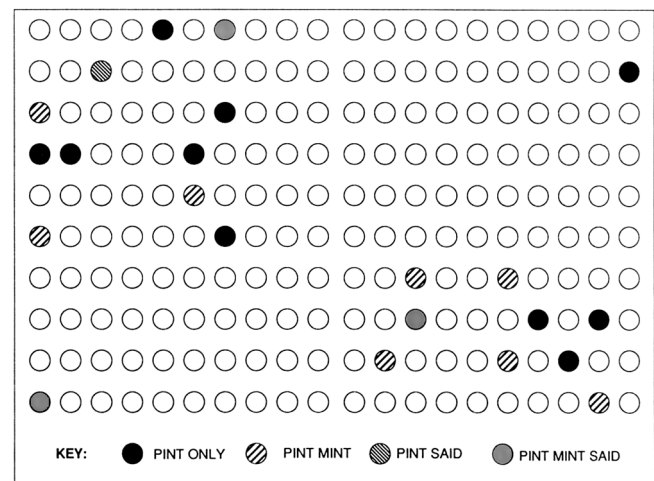|	Type of error score		
Case	Orthography	Correct pron	Reg pron
Normal baseline	6.47	6.64	34.6
Damage to individual units Unique to PINT			
Unit A	6.87	12.78	25.4
Unit B	9.61	7.19	25.4
Unit C	9.73	9.87	34.9
In LINT, MINT, & PINT Unit D	7.45	10.56	33.6
Turning on a unit in LINT & MINT	6.48	11.17	29.4

Note. Orthography = orthographic error score; Correct pron = phonological error score for correct pronunciation; Reg pron = phonological error score for regularized pronunciation.

by PINT, MINT, and LINT, and 3 by PINT, LINT, MINT, and SAID, reflecting some overlap among these items. Thus, inspection of the hidden units provides additional evidence that the model encodes orthographic and phonological relations among words.

It should also be noted that the units activated by a particular word contribute in different ways to the computed output. This point can be illustrated as follows. After 250 epochs of training, the word PINT produces the following results (Table 5): The orthographic error score is 6.47, the phonological error score computed for the correct pronunciation is 6.64, and the phonological error score computed for the incorrect, regularized pronunciation is 34.6. If we consider the patterns of activation for PINT, LINT, MINT, and SAID, there are 9 units unique to PINT. The contribution of an individual unit can be determined by temporarily excluding the unit (i.e., forcing its activation to remain fixed at zero) and then recalculating the output and error score. This procedure has different effects depending on which unit is zeroed. Shutting off one of the units unique to PINT (Unit A in Table 5) has little effect on the computed orthographic output, but dramatically increases the error score associated with the correct pronunciation and decreases the error score associated with the regularized pronunciation. Hence, this unit appears to be particularly relevant to the irregular pronunciation of PINT. Eliminating the output from Unit B has a somewhat different effect; it produces a large increase in the orthographic error score and smaller increases in the error scores for the correct and regularized pronunciations. Hence, this unit is primarily relevant to the orthography and may partially influence aspects of the pronunciation that are shared by the regular and exceptional pronunciations of PINT. Unit C produces a third pattern: substantial increases in the error scores for both the correct orthographic and phonological codes, with little effect on the score for the incorrect phonological code. Thus, each unit makes its own partial contribution to the model's performance on PINT.

We also examined the effects of zeroing a unit that is activated by LINT, MINT, and PINT. This produced a small increase in the orthographic error score; the effect on the phonological error

score for the correct pronunciation was intermediate between the effects of Units A and B. This appears to be a complex unit encoding information relevant to the correct spelling and to both pronunciations of -INT. Finally, consider the effects of activating a unit that is normally active in LINT and MINT but not normally active in PINT. This has virtually no effect on the orthographic output for PINT but yields an increase in the phonological error score for the correct pronunciation and a decrease in the error score for the incorrect, regularized pronunciation. Hence, the unit appears to be relevant to the regular pronunciation of -INT.

It can be seen, then, that the units contribute in complex ways to the computation of orthographic and phonological output. Some units must be on in order to produce correct output, and others must be off. Some units can be seen as contributing in relatively specific ways to the computed output (e.g., Unit A, which is critical to the pronunciation of -INT as in PINT, and Unit D, relevant to pronouncing -INT as in MINT). Other units can be seen as partially encoding several different types of information. This behavior is typical of models with hidden units. Often it is possible to identify the specific information encoded by individual units; however, many units contribute to the computed output in complex ways that do not reflect simple generalizations about the relations between two codes. To take another example, Hinton et al. (1986) described a small-scale model of the mapping from orthography to meaning. The hidden units in a model of this type will encode generalizations about correlations among semantic features. Some hidden units may be interpretable as encoding a generalization such as "large and yellow," whereas others will not because they encode complex, partial relations among several features.

Note also that generalizations concerning relations between orthography and phonology are encoded by several units rather than individual ones. For example, there is no single unit that encodes the pronunciation /int/ common to LINT, MINT, and other rhymes. Nor is there a single unit responsible for the irregular pronunciation of PINT. Although we identified a unit that is particularly salient to pronouncing -INT as /int/, other units also contribute to this pronunciation. Given this property of the model and the fact that units participate in many different words, spelling–sound correspondences cannot be seen as encoded by individual units.

To consider one more example, we examined the patterns of activation over the hidden units produced by the word MAID, the similarly spelled rhyme PAID, the similarly spelled non-rhyme SAID, the homophone MADE, and the unrelated word BASK. Sixteen units were activated by both MAID and PAID, 5 by MAID and SAID, and 4 by MAID and BASK, reflecting the differing degrees of orthographic and phonological similarity among these items. It is interesting that the homophonic pair MAID–MADE shared 13 units, somewhat fewer than the similarly spelled rhymes MAID and PAID, but more than would be expected if the words were unrelated. Thus, the degree of similarity between the words is systematically related to the activity of the hidden units.

Relationship to Other Models

With this picture of the model in hand, we can consider how it relates to previous proposals. In general, the model embodies

many of the principles that had been identified in previous work; however, it shows that they derive from a deeper generalization about the nature of the learning process.

Our model accounts for a number of phenomena that are problematical for the dual-route model, specifically, the interaction between frequency and regularity and the longer latencies for regular inconsistent nonwords compared with regulars. These effects are not predicted by the dual-route model and could only be accommodated by ad hoc extensions to it (Seidenberg, 1985a). The dual-route model also has other limitations that have been discussed extensively elsewhere (e.g., Humphreys & Evett, 1985; Seidenberg, 1985a). The model corroborates the common assumption that the ends of words—word bodies or rimes—are relevant to naming (Brown, 1987; Glushko, 1979; Meyer et al., 1974; Seidenberg, Waters, Barnes, & Tanenhaus, 1984; Treiman & Chafetz, 1987). This fact falls out from properties of the learning algorithm and training corpus. The ends of words turn out to be salient because of the properties of written English; the pronunciations of vowels are more influenced by the following letters than by the preceding ones. The learning algorithm picks up on these regularities, which have an impact on the weights. Importantly, the characteristics of the learning algorithm also dictate that the effective relations between words are not limited to word bodies. These units happen to be salient, but other regularities in the lexical corpus are also picked up.

The model incorporates Glushko's (1979) insight that the pronunciation of a word or nonword may be influenced by knowledge of the pronunciations of other, neighboring words. As the Andrews (in press) study showed, words with more neighbors tend to be named more quickly than words with fewer neighbors; in the model, this occurs because the neighbors of a word tend to modify the weights in the same direction as the word itself. These effects are smaller for higher frequency words, however, because of repeated exposure to the words themselves. The model also incorporates Glushko's assumption that inconsistencies in spelling–sound correspondences are relevant to performance; inconsistent neighbors push the weights away from the values that are optimal for pronouncing a given word. The representations and processes in our model differ in critical respects from his proposal, however. Glushko's model contains nodes for individual words, and pronunciations are synthesized on the basis of competition among partially activated entries. Our model contains no word-level nodes; the competition between words is realized in the effects of the connection weights, which are determined by exposure to many items. Our model captures the notion of lexical analogy that was central to Glushko's model in terms of the consequences of learning within a distributed system. A second difference between the accounts is that Glushko assumed that the ends of words—word bodies—have a special status in naming. This assumption has been widely accepted by reading researchers (see, e.g., Brown, 1987; Henderson, 1982; Parkin, 1982; Patterson & Coltheart, 1987; Seidenberg, Waters, Barnes, & Tanenhaus, 1984; Treiman & Chafetz, 1987). In our model, there is no single perceptual unit relevant to pronunciation; the model picks up on regularities in terms of word endings, but also regularities involving other parts of words.

The model is consistent with Brown's (1987) principle con-

cerning the number of times a word body is associated with a given pronunciation; again, it can be seen that the principle is simply one of the consequences of the learning process. However, Brown made an additional assumption that is not congruent with our model, namely, that inconsistencies in spelling–sound correspondences do not influence processing. In Brown's model, for example, the number of times -OSE is pronounced /uz/ and the number of times it is pronounced /oz/ are separate facts that do not interfere with one another. This assumption provided the basis for the prediction that exception words such as LOSE and unique words such as SOAP should yield similar naming latencies, despite the fact that LOSE has inconsistent neighbors. In our model, the effects of experience in naming LOSE and POSE (and all other words) are superimposed on the weights, rather than separated in the manner Brown suggested. Hence, our model predicts that consistency of a spelling–sound correspondence could affect naming, whereas Brown's does not. In effect, Brown's model suggests that repetition of a spelling pattern with a given pronunciation facilitates performance, with no interference due to exposure to an inconsistent pronunciation. In our model, performance is determined by the net effects of exposure to both pronunciations (and to other words); interference can result when training is inconsistent.

The experiment presented in Brown (1987) does not discriminate between the two theoretical alternatives because, as the data in Figure 9 indicate, our model simulates the results even though it does not conform to Brown's assumptions about inconsistency.[8] Critical cases are provided, however, by the regular inconsistent words and nonwords discussed earlier. Our model predicts inconsistency effects whose magnitude will depend on factors such as the frequencies of the regular inconsistent and exception words and their similarity to other items. According to Brown's model, regular inconsistent words should yield

---

[8] Note, however, that Brown's (1987) study does not provide clean evidence for his principle. The critical comparison in the experiment is between unique and exception words. These words are similar in terms of the factor Brown assumed to be relevant, the number of times their word bodies are associated with a given pronunciation (in both cases, the number is 1). They differ in terms of the factor thought to be irrelevant; only the exception words have inconsistent neighbors. Hence, the finding that the words yield similar naming latencies was taken as evidence that only the first of these factors is relevant. However, the words also differ in other respects relevant to processing (and to our model). Specifically, exception words contain higher frequency spelling patterns than do unique words. This is a necessary consequence of the fact that the exceptions have a large number of regular inconsistent neighbors. Hence, there is a confounding between the number of times a spelling pattern occurs in the orthography and consistency of pronunciation. In our model, both of these factors are relevant; they jointly account for why performance is similar on exception words such as LOSE and unique words such as SOAP. The model is trained on a large number of -OSE words, and the weights come to reflect the fact that these words typically rhyme with POSE. It then performs relatively poorly on the exception LOSE. Unique words such as SOAP fare differently. The fact that -OAP is pronounced /op/ is not very strongly encoded by the weights because this pattern is encountered so infrequently. This also means, however, that the model has not been given inconsistent feedback about the pronunciation of this pattern. The tradeoffs between these factors, which are realized in the learning process, are such that SOAP and LOSE are about equally difficult to name.

Table 6
*Characteristics of the Stimuli in the Seidenberg,
McRae, and Jared (1988) Study*

Characteristic	Type	
	Consistent	Inconsistent
n	40	40
KF freq	5.50	5.50
Friends	8.5	8.3
Total freq friends	496	467
Length in letters	4.5	4.5
Orth error score	8.63	8.23
Enemies	0	3.2
Total freq enemies	0	496
Phon error score	4.48	5.63

Note. KF freq = mean Kucera & Frances (1967) frequency; Friends = mean number of words in which word-body occurs with regular pronunciation; Total freq friends = average of the summed frequencies of the friends; Orth error score = mean orthographic error score from the model; Enemies = mean no. of words with same word body but different pronunciation; Total freq enemies = average of the summed frequencies of the enemies; Phon error score = mean phonological error score from the model.

longer latencies than regular words only if the two types of words are not equated in terms of the factor he assumed to be relevant, the number of times their word bodies occur with regular pronunciations. That is, if the word bodies in regular inconsistent words occur with regular pronunciations in fewer items than do the word bodies in regular words, the regular inconsistent items should yield longer naming latencies. If the items are equated in terms of word-body frequency (and other factors such as lexical frequency), no difference should obtain.

As we have noted, previous studies have not yielded reliable differences between regular inconsistent and regular words. However, these studies may not be definitive for two reasons. First, as we have seen, the mere presence of a single exception word neighbor may produce negligible effects on a regular inconsistent word because the effects are washed out by exposure to a large number of words containing the regular pronunciation, including the word itself. Our model does not predict appreciably longer latencies for *all* words defined as regular inconsistent compared to matched regular words; however, it does predict detectable consistency effects for some words, particularly lower frequency words that have more than a single exception word neighbor. For example, the pattern -ONE is highly inconsistent because it is associated with three pronunciations, one regular (as in BONE) and two exceptional (GONE; DONE/ NONE). This inconsistency might be expected to influence the processing of a lower frequency word such as HONE. Similarly, the pattern -OSE is associated with three pronunciations (as in POSE, LOSE, and DOSE). If, as our model suggests, there are effects due to inconsistencies in spelling–sound correspondences, they should be more apparent using stimuli that include such words. A second factor is that the stimuli in previous studies of regular inconsistent words were not equated in terms of the Brown (1987) factor, the frequencies with which their word bodies are associated with regular pronunciations. With these issues in mind, Seidenberg et al. (1988) conducted the following

experiment. The stimuli were 40 pairs of consistent (regular) and inconsistent words (see Appendix). The properties of these words are summarized in Table 6. They were equated in terms of the number of words in which the word bodies occur with regular pronunciations (termed *friends* in the table); this is Brown's factor. They were also equated in terms of the summed frequencies of these friends. Thus, both types of stimuli contain word bodies that are associated with regular pronunciations about equally often. Seidenberg et al. (1988) also matched the stimuli in terms of overall frequency, length, and initial phoneme. The two types of words were also equated in terms of orthographic error scores, so that any differences between them cannot be attributed to orthographic redundancy. The systematic difference between the words is that the inconsistent items have enemies—words that contain the same word body but are pronounced irregularly. As a result, the two types differ in terms of mean phonological error scores (inconsistent = 5.63; consistent = 4.48); this difference is statistically significant. Thus, our model predicts longer latencies for the inconsistent words, whereas Brown's model predicts no difference because the stimuli are equated in all respects relevant to his account. The study was run with 25 McGill University undergraduates as subjects, who named the words aloud as they appeared on a computer screen. The results, presented in Figure 19, showed a 13-ms inconsistency effect, which was significant in both subject and item analyses. The phonological error scores for these words also provide a good fit to the latency data.

The results of Glushko's (1979) nonword experiment (presented in Figure 11) also contradict Brown's (1987) model. The study showed that nonwords derived from inconsistent spelling patterns (e.g., MAVE from HAVE/GAVE) yield longer naming latencies than do nonwords derived from regulars (e.g., NUST from MUST). Note that the difference here is between the latencies to produce the *regular* pronunciations of these stimuli. According to Brown, this difference should only obtain if the word bodies in the inconsistent stimuli were associated with regular pronunciations in fewer words than the word bodies in the regu-
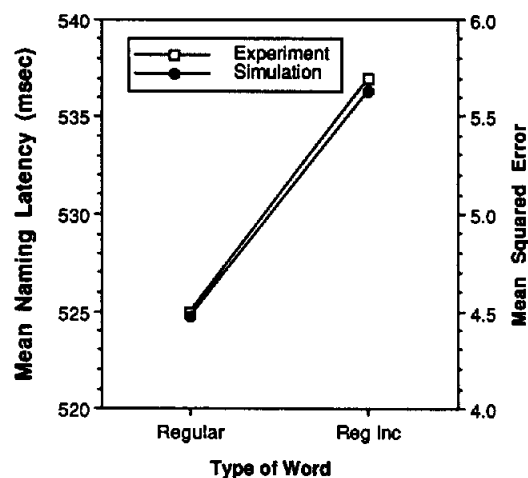

Figure 19. Results of the Seidenberg, McRae, and Jared (1988) experiment with regular inconsistent and regular words equated in terms of word-body frequency: Experiment and simulation data.

lar items. In Glushko's stimuli, however, the opposite pattern obtains: the inconsistent nonwords have an average of 9.5 regular neighbors, whereas the regular nonwords have an average of 6.2 regular neighbors. Because the regular inconsistent nonwords actually have *more* regular neighbors but yield *longer* latencies, the results are not consistent with Brown's model.

In summary, the model simulates the results of a broad range of empirical studies using many different sets of stimuli. The factor that Brown (1987) isolated—the number of times a word body is associated with a given pronunciation—has an impact on performance, one that must be considered in drawing comparisons between different types of items. However, this is not the only factor that influences performance; inconsistencies in spelling–sound correspondences also matter. Moreover, aspects of word structure other than word bodies also affect processing, such as overlap in terms of the beginnings of words (Taraban & McClelland, 1987). The pretheoretical distinctions between different types of stimuli (e.g., regular inconsistent and regular; unique and exception) are difficult to maintain because several different factors—overall frequency, word-body frequency, regularity, orthographic redundancy, and so on—are typically confounded in the language. These natural confoundings are neatly handled by the model in terms of the aggregate effects of training on the settings of the weights on connections.

Acquisition of Naming Skills

We have suggested that the model provides a good characterization of a broad range of phenomena related to the naming performance of skilled readers. As a learning model, it also speaks to the issue of how these skills are acquired; moreover, it provides an interesting perspective on the kinds of impairments characteristic of developmental and acquired dyslexias. Developmental dyslexia could be seen as a failure to acquire the knowledge that underlies word recognition and naming. Acquired dyslexias naturally correspond to impairments following damage to the normal system. Here, we focus on the acquisition of naming skills and their impairment in developmental dyslexia. Our studies of acquired forms of dyslexia are discussed in Patterson et al. (in press).

Studies of children's acquisition of word recognition skills (e.g., Backman et al., 1984; Barron & Baron, 1977; Jorm & Share, 1983; Seidenberg, Bruck, Fornarolo, & Backman, 1986) have addressed how children reach the steady state observed in adults; they have also addressed the bases of failure to acquire age-appropriate reading skills and of specific reading disability (dyslexia). Naming plays an important role in acquiring word recognition skills; children in the earliest stages of learning to read typically recognize words by "sounding out"; that is, they attempt to derive the pronunciation of a written word and match it to a known phonological form. A study by Backman et al. (1984) examined the acquisition of naming skill. Children named regular, exception, regular inconsistent, and ambiguous words and nonwords derived from these items. All of the stimuli were words that are high-frequency items in adult vocabularies. The subjects were children in Grades 2, 3, 4, and in high school, reading at or above age-appropriate levels ("good readers"), and children in Grades 3 and 4, reading below age-appropriate levels ("poor readers"). Response latencies showed the expected


Figure 20. Results of Backman, Bruck, Hébert, and Seidenberg (1984) naming study. (PR = poor readers in Grades 3 and 4; HS = high school students.)

developmental trends: Younger and poorer readers named words at longer latencies than older, better readers. The effects of word type were manifested in the number of mispronunciation errors.

The primary data are summarized in Figure 20. The developmental trends exhibited in these data are clear: Younger, less skilled readers have more difficulty with the words associated with multiple pronunciations (exception, regular inconsistent, ambiguous); they show larger regularity effects. The reader groups differed very little in performance on regular items. As children acquire reading skills, the differences between word classes shrink and disappear. The less skilled readers have weaker knowledge of spelling–sound correspondences; this lack of knowledge is a liability in the case of words with irregular, inconsistent, or ambiguous spelling–sound correspondences. Older children and adults are able to compute the pronunciations of high-frequency exemplars of all word classes about equally well; differences between word classes only persist for lower frequency items. The unskilled readers' performance in naming *higher* frequency words is therefore similar to that of skilled readers' naming of *lower* frequency words; in effect, the developmental data reveal the emergence of the modulating effects of experience on naming performance. At the same time that children are achieving the ability to name different types of words equally well, their knowledge of spelling–sound correspondences is expanding, as evidenced by the older readers' superior performance in reading nonwords (Backman et al., 1984).

Consider these facts in light of the simulation data presented earlier. The data for regular and exception words presented in Figure 3 show that early in training, the model produces poorer output for exception words compared to regular in both frequency ranges. Like children in the early stages of reading acquisition, the model performs more poorly even on higher frequency exception words. The effect of training is to decrease the error scores to a point at which the two types of higher frequency words reach floor values, yielding the Frequency × Reg-

ularity interaction also observed in adults. The data for the regular inconsistent words and regular controls (Figure 7) are also interesting, because early in training, there are small differences between regular inconsistent and regular words in both high- and low-frequency ranges. Backman et al.'s (1984) children also produced more errors on common regular inconsistent words than on regular. As in the model, performance was better for regular inconsistent words than for exceptions. It is clear why there are regular inconsistent effects early in acquisition but not late; early in training, both exception and regular inconsistent items appear about equally often. It is only after additional experience that the regular spelling–sound patterns gain the upper hand.

Thus, the model captures a key aspect of the child's acquisition of word naming skills. It is known that acquiring knowledge of spelling–sound correspondences is a key component of learning to read; disorders in phonological analysis skills are thought to be a primary source of reading disability, and children who are backward readers (Backman et al., 1984) or developmental dyslexics (Seidenberg et al., 1986) exhibit relatively poor performance in naming words and nonwords aloud (see Jorm & Share, 1983, and Stanovich, 1986, for reviews). One of the primary developmental trends observed in studies such as Backman et al. (1984) is that although children who are acquiring age-expected reading skills initially have more difficulty naming higher frequency exception words (and other items containing spelling patterns associated with multiple pronunciations) than regular words, this deficit is eliminated by about Grade 5 (10 years of age). During the first few years of instruction, children learn to name common exception words as efficiently as they do regular words. Even among skilled adult readers, however, lower frequency exception words continue to produce longer naming latencies and more errors than do lower frequency regular words. Note that the differences among younger good and poor readers in terms of the number of words read without interference from irregular spelling–sound correspondences are seen, at a higher level of performance, among skilled readers (Table 4). In both groups, the number of words in this pool is related to reading skill. In the model, these effects simply derive from the amount of training experience.

Both poor readers who are reading below age-expected levels and children who have been diagnosed as developmental dyslexics fail to show this improvement in naming higher frequency exception words. For example, the naming performance of the poor readers in Grades 3 and 4 in the Backman et al. (1984) study was like that of good readers in Grade 2. Both the younger and poorer readers made more errors on exception words and other items containing spelling patterns associated with multiple pronunciations.

Developmental Dyslexia

Developmental dyslexia is a term applied to children who are failing to acquire age-appropriate reading skills despite adequate intelligence and opportunity to learn (Vellutino, 1979). The nature of this disorder—whether it derives from a single or multiple cause(s), whether there are different subtypes, or whether the performance of children diagnosed as dyslexic differs from that of children who are merely poor readers—is

a matter of continuing debate. However, it is clear that many dyslexic children exhibit poor word decoding and naming skills, and there is some evidence that these impairments have a biological basis (Benton, 1975; Vellutino, 1979).

The model suggests a basis for the impaired performance of some dyslexic readers, who appear to be unable to master fully the spelling–sound correspondences of the language. Consider the results of an experiment in which we retrained the model with one half as many hidden units, 100 instead of 200. In all other respects, the training procedure was the same as before. At the start of training, all of the weights were given small random values. The model was again trained on the 2,897-word vocabulary. In the simulations reported here, we used a version of the training list in which the coding errors mentioned earlier were corrected. Training was also carried out for 500 epochs instead of 250. Figure 21 (upper graph) gives the mean phonological error scores for regular and exception words in the Taraban and McClelland (1987) stimulus set when the model was trained with 200 hidden units. This is a replication of the simulation reported in Figure 3 (note, however, the change of scale on the ordinate). In Figure 21 (lower graph) we summarize the data for the same words in the simulation using 100 hidden units. Two main results can be observed in comparing the two data sets. First, training with fewer hidden units yields poorer performance for all word types. High-frequency regular words, for example, asymptote at a mean squared error of about 2 in the 200-unit simulation but only 3.8 in the 100-unit simulation; other words yield similar results. Second, even after 500 epochs, exception words produce significantly poorer output than do regular words in both high- and low-frequency ranges in the 100-unit simulation; in the 200-unit simulation, exception words produce larger error scores only in the lower frequency range.

Eliminating one half of the hidden units, then, produced a general decrement in performance; more important, *higher* frequency words produced the patterns associated with *lower* frequency words in the 200-unit simulation (i.e., larger error scores for exception words compared to regular). Even with fewer hidden units, the model continued to encode generalizations about the correspondences between spelling and pronunciation; error scores were smaller for regular words than for other types. However, it performed more poorly on words whose pronunciations are not entirely regular. Thus, including fewer hidden units makes it more difficult to encode item-specific information concerning pronunciation.

These results capture a key feature of the data obtained in studies of poor readers and dyslexics. These children exhibit *larger* regularity effects than do good readers; they continue to perform poorly in naming even higher frequency exception words. At the same time, their performance shows that they have learned some generalizations about spelling–sound correspondences; for example, they are able to pronounce many nonwords correctly. One of the main hallmarks of learning to read English is acquiring knowledge of spelling–sound correspondences. Backward readers achieve some success in this regard, but cope poorly with the irregular cases. The model performs in a similar manner with too few hidden units; given the resources that are available, it is able to capture crude generalizations about regularity but at the expense of the exception
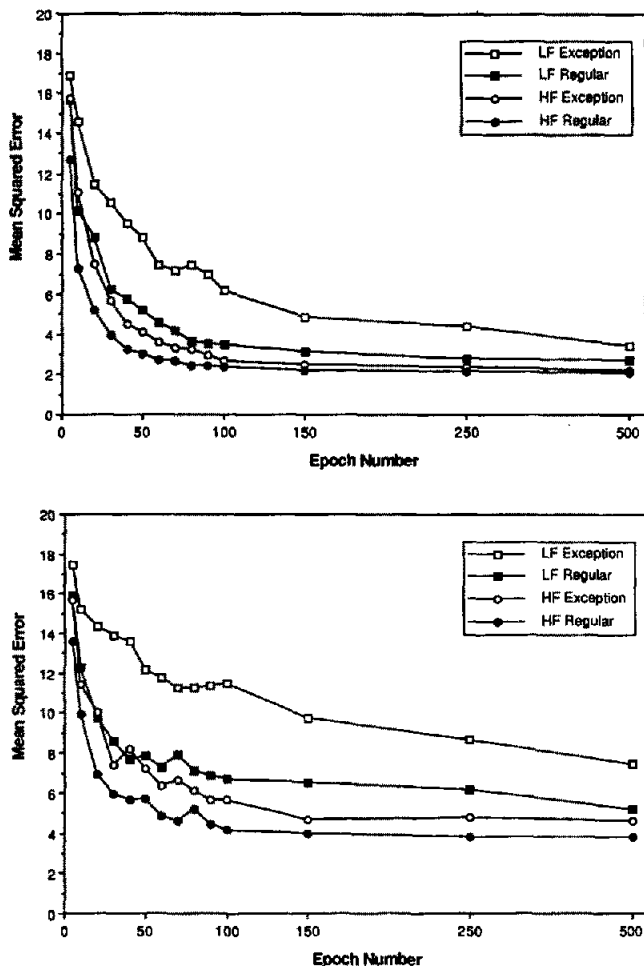
Figure 21. Model's performance on new simulations using the Taraban and McClelland (1987) regular and exception words: 200 (upper graph) and 100 (lower graph) hidden units.

to learn on a word-by-word basis, resulting in adequate performance on regular and exception words but very poor generalization to novel stimuli (see Barron, 1986, for discussion). These children apparently fail to encode generalizations concerning spelling–sound regularities. One possibility is that this type of performance results from the use of a somewhat different encoding of orthographic input or phonological output, or both. If the amount of overlap in the encoding of similar inputs or outputs is reduced, there will be less transfer of what is learned about one word to other words that are similar to it. Yet another possibility is that the pathway from orthography to phonology is so grossly deficient in such readers that they read primarily by accessing meaning from print, and then producing the pronunciation corresponding to the accessed meaning. Hence, only words that are within the child's vocabulary can be pronounced. This possibility is consistent with the full version of our model illustrated in Figure 1.[10]

Summary of the Naming Simulations

The model provides a basis for understanding the manner in which knowledge of orthographic–phonological correspondences is acquired, represented in memory, and used in naming. The generalization that governs the model's performance concerns the properties of the writing system that are picked up during learning. All of the various empirical phenomena observed in the behavioral studies we have reviewed (concerning neighborhood effects, lexical analogy, word-body frequencies, and the like) fall out of this single property of the model. The model goes beyond earlier proposals in suggesting that the best characterization of the knowledge relevant to pronunciation is given by the entire state of the network, rather than by generalizations concerning spelling–sound rules, perceptual units, or types of words.

The model differs from previous accounts in terms of the kinds of knowledge representations and processes used. In contrast to the dual-route model, there are no rules specifying the regular spelling–sound correspondences of the language, and

words. The main implication of the simulation, of course, is that failures to achieve age-expected reading skills may derive from limitations on the computational resources available for the task. There is another important implication, however. Apparently, the architecture of the model determines in an important way its ability to behave like humans. If there are too few units, the model can learn generalizations about the regularities in the writing system; however, it does not have the capacity to encode enough of the word-specific information relevant to exception words to perform as well as people perform. With a sufficient number of units, it is able to cope with both regular and irregular cases, although not equally well on all items. The important point is that human performance seems to reflect rather subtle constraints concerning computational resources. The idea that impaired performance might result from dedicating too few resources to a task is one that could be pursued in future research.[9]

We should stress that some dyslexic children exhibit other patterns of performance, suggesting that the normal system can be impaired in other ways. For example, some children appear

---

[9] Although it is tempting to equate the number of hidden units with the size of the population of neurons that might be dedicated to reading in the brain, one must be careful not to take this analogy too literally. First, the precision of the individual units used in our simulations could only be achieved by much larger numbers of actual neurons. Second, resource limits might arise in a number of ways, such as degree of noise or number of modifiable connections per neuron, rather than strictly in terms of numbers of neurons involved.

[10] We also considered the possibility that generalization would be reduced if the model were given too many hidden units. This has been observed in some experiments with back propagation (e.g., Hinton, 1986). This behavior would correspond to learning the pronunciations of words on an item-by-item basis, leading to poor performance on novel stimuli such as nonwords. We ran one simulation using 400 hidden units that yielded results very similar to the ones with 200 hidden units, except that learning was faster and lower error scores were achieved. Thus, in the present case at least, merely doubling the number of hidden units does not significantly reduce the generalization performance of the model. We are continuing to explore this and other possible computational bases for different patterns of dyslexic performance (see also Patterson, Seidenberg, & McClelland, in press).

there is no lexicon in which the pronunciations of all words are listed. All items—regular and irregular, word and nonword—are pronounced using the knowledge encoded in the same sets of connections. The main assumption of the dual-route model is that separate mechanisms are required in order to account for the capacity to name exception words and nonwords (Coltheart, 1986). Exception words cannot be pronounced by rule, only by consulting a stored lexical entry; hence, one route is called *lexical* or *addressed* phonology. Nonwords do not have lexical entries; hence, they can only be pronounced by rule. Hence the second route, termed the *nonlexical* or *subword* process. One of the main contributions of the model is that it demonstrates that pronunciation of exception words and nonwords can be accomplished by a single mechanism using weighted connections between units. The analysis of the hidden units also indicated that the model did not partition itself in a manner analogous to the routes in the dual-route model.

The model suggests that the distinction between words that conform to the spelling–sound rules of the language and those that do not (i.e., the contrast between regular and exception words), which motivated the dual-route model, is simply not rich enough to account for human performance. Connection weights reflect the cumulative effects of many learning trials, each of which imposes small changes on the weights. Correct predictions about performance follow from an understanding of what is learned on this basis, not merely whether or not a pronunciation obeys a putative rule. Thus, words whose pronunciations are equally well specified by the rules can differ in terms of naming performance; performance on words that violate the rules also differs depending on their similarity to other words. The distinction between rule-governed items and exceptions fails to capture these generalizations.

Our model also differs from proposals by Glushko (1979) and Brown (1987) in that there are no lexical nodes representing individual words and no feedback from neighbors. Where the model agrees with these accounts is in regard to the notion that regularity effects result from a conspiracy among known words. In our model, this conspiracy is realized in the setting of connection strengths. Words with similar spellings and pronunciations produce overlapping, mutually beneficial changes in the connection weights.

Following the work of Glushko (1979), a number of researchers have developed definitions of regularity or consistency based on assumptions as to which perceptual units or neighborhoods are relevant to pronunciation (e.g., Kay & Bishop, 1987; Parkin, 1982; Parkin & Underwood, 1983; Patterson & Coltheart, 1987). From the perspective of the model, these definitions miss relevant generalizations concerning the kinds of knowledge that underlie pronunciation, how this knowledge is represented in memory, and how it influences processing. There is no single perceptual unit relevant to pronunciation. The output that the model produces for a given letter string is determined by the properties of *all* of the words presented during training. From this perspective, the various definitions of regularity or neighborhood are simply imperfect generalizations about the nature of the input and its effects on what is learned.

Orthographic Output and Lexical Decision

We turn now to other aspects of the model that are of interest primarily because of their relevance to the lexical decision task,

which is probably the most widely used task in reading research. One of the main features of the model is that it uses distributed representations: The spellings and pronunciations of words are represented in terms of patterns of activation across output nodes. In this respect, the model differs radically from previous conceptions of lexical knowledge, which assumed that the spellings and pronunciations of words are stored as entries in one or more mental lexicons (e.g., Coltheart, 1978, 1987; Forster, 1976; Morton, 1969). We have shown that the model provides a good account of subjects' performance in naming words aloud. The question that arises is whether this type of knowledge representation can support performance on other tasks. Lexical decision presents an especially challenging case because standard accounts of the task assume that it is performed by accessing the kinds of lexical entries that our model lacks.

In the following sections we present an account of lexical decisions to isolated words and nonwords, and show that the model simulates the results of many experiments. Our main point is that distributed representations provide a basis for making lexical decisions; moreover, the model provides an enlightening account of some complex lexical decision phenomena. Interestingly, the model simulates many of the main lexical decision phenomena despite the absence of any representation of meaning at all; thus, our account of the task runs contrary to the standard view that decisions are necessarily made by determining whether or not the target stimulus has a meaning. We do not doubt that meaning is sometimes relevant, and we note that our account of lexical decision is necessarily limited because we have not implemented a semantic system or provided a way for contextual information to influence processing, as it is of course known to do (e.g., Fischler & Bloom, 1979; Schwanenflugel & Shoben, 1985; Seidenberg, Waters, Sanders, & Langer, 1984). Both of these components are relevant to lexical decision performance under conditions that are beyond the scope of the present model.

Although considerations of contextual and semantic factors have often entered into lexical decision experiments, the task has also been widely used as a way to investigate the structural properties of words relevant to "lexical access." The subject is presented with a string of letters and must decide whether it forms a word. Use of the task was predicated on the observation that words and pronounceable nonwords differ in an essential respect: Words have conventional meanings, and nonwords do not. It was initially assumed that this distinction between the stimuli provided the basis for making the word/nonword decision: word decisions are made by identifying the target as a particular word and accessing its meaning; if this process fails, the target is a nonword. Hence, the task could be used to study the properties of words (e.g., frequency, orthographic redundancy, orthographic–phonological regularity) that influence access to lexical representations and then meaning (Henderson, 1982; McCusker, Hillinger, & Bias, 1981). However, words and nonwords also differ in other respects, providing other bases for making the decision; for example, words are more familiar orthographic and phonological patterns than nonwords. The task requires subjects to discriminate between the two types of stimuli. As in a signal detection task, the subject must establish decision criteria that allow fast responses with acceptable error

rates. These criteria could in principle involve any of the several dimensions along which words and nonwords differ. Perhaps the primary conclusion from extensive use of the task is that response criteria vary as a function of the properties of the stimuli in an experiment. As subjects' response criteria vary, so do the effects of variables such as frequency, orthographic–phonological regularity, and contextual congruence (e.g., Forster, 1981b; Neely, 1977; Seidenberg et al., 1984; Stanovich & West, 1981).

The general framework given in Figure 1 suggests that the presentation of a word results in the computation of several types of information or codes in parallel, resulting in what Donnenwerth-Nolan, Tanenhaus, and Seidenberg (1981) called *multiple code activation*. We have emphasized the computation of the phonological code and shown that the model provides a good account of the empirical naming data. We envision an analogous process, which has not been implemented, by which readers compute the meaning of a word, corresponding to a pattern of activation across a set of semantic nodes; see Kawamoto (1988) and Hinton et al. (1986) for initial steps toward modeling this process. Finally, the implemented model also includes the computation of orthographic output, resulting from feedback from the hidden units to the orthographic units. This code represents the retention or recycling of the orthographic input in a short-term sensory store; the computed code provides the basis for performing tasks such as tachistoscopic recognition and thus accounting for the phenomena that motivated the McClelland and Rumelhart (1981) word recognition model.

Presentation of a stimulus string will activate orthographic, phonological, and semantic information in parallel, each of which could provide information relevant to the decision process, depending on the conditions in an experiment. Consider, for example, a case in which the stimuli consist of familiar words and nonwords that are random letter strings. Subjects could respond correctly simply on the basis of orthographic information; the words contain letter patterns that are "legal" according to English orthography, whereas the nonwords will contain letter patterns that do not occur in any words (e.g., PSKT). Properties of the words related to phonology (e.g., orthographic–phonological regularity) or meaning (e.g., concreteness/abstractness) would have little effect on performance if decisions were made on the basis of this orthographic strategy. If the stimuli included familiar words and orthographically legal nonwords (such as NUST), this simple orthographic strategy might be disabled. However, the stimuli still provide a nonsemantic basis for responding; the subject could decide if the target is a word by determining whether it has a familiar pronunciation. When the decision is based on phonological information, we might expect factors such as orthographic–phonological regularity to affect performance. In principle, this strategy might in turn be disabled if the nonword stimuli were so-called *pseudohomophones* such as BRANE, which sound like words (Dennis, Besner, & Davelaar, 1985). Because these stimuli look and sound like words, subjects might be required to use semantic information in making their decisions. Later, we consider evidence that subjects do in fact modify their decision strategies in such ways (see also Bradshaw & Nettleton, 1974; James, 1975).

Similar considerations apply when target words and nonwords appear in word or sentence contexts. Although subjects could in principle base their decisions on the properties of the word and nonword targets, they find it very difficult to inhibit comparing targets to the contexts in which they occur. Here, decision latencies are influenced by the perceived congruence of target and context. Neely (1977), for example, showed that contextual information influences both word and nonword decisions; moreover, decision latencies depend on factors such as the types of contextual information provided and the proportions of trials of different types (Seidenberg, Waters, Sanders, & Langer, 1984; Tweedy, Lapinski & Schvaneveldt, 1977). Again, subjects respond intelligently to the information provided by the stimuli in the experiment and modify their response strategies to improve performance.

The logic of the lexical decision task, then, does not necessarily require the subject to access the meanings of word targets; rather, it requires the subject to find a basis for reliably discriminating between words and nonwords. The model suggests that there are at least three types of information that could enter into the decision process for isolated stimuli. When targets appear in meaningful contexts, there is a fourth source of information. Which information is used depends on the properties of the stimuli, which afford different response strategies. A theory of lexical decision performance must provide a principled account of how strategies vary as a function of the stimulus conditions. We illustrate this aspect of the model by considering some data that have been the source of considerable puzzlement.

Variable Effects of Orthographic–Phonological Regularity

There have been many lexical decision studies, analogous to the naming studies described earlier, using regular and exception words (these include Andrews, 1982; Bauer & Stanovich, 1980; Coltheart, Besner, Jonasson & Davelaar, 1979, Parkin, 1982; Parkin & Underwood, 1983; Seidenberg, Waters, Barnes, & Tanenhaus, 1984; Waters & Seidenberg, 1985). As in the naming studies, orthographic–phonological regularity has negligible effects on lexical decisions for higher frequency words. Whereas the naming studies have yielded robust exception effects for lower frequency words, the results of the lexical decision experiments have been inconsistent. In studies such as those of Coltheart et al. (1979) and Seidenberg, Waters, Barnes, and Tanenhaus (1984, Experiment 3), no effects of orthographic–phonological regularity were observed, whereas in others (such as Parkin, 1982, and Bauer & Stanovich, 1980), they were.

These inconsistent effects have been interpreted as indicating that words can be recognized by either direct (visually based) or mediated (phonologically based) processes (Barron, 1986; Carr & Pollatsek, 1985; Seidenberg, 1985b). In cases in which there were no effects of phonological regularity, it was inferred that recognition is direct; in cases in which there were such effects, recognition was thought to be phonologically mediated. Use of these alternative strategies was thought to be under the reader's control (Coltheart, 1978). This account left a key question unresolved, however: It did not explain the factors that determined why a particular strategy seemed to be used in a particular experiment. Note that the inconsistent results that led to this view involved the same types of stimuli (regular and ex-

Figure 22. Results of the Waters and Seidenberg (1985) lexical decision study.

ception words) used in different experiments. Hence, it cannot be the case that direct access is used for one type of word (e.g., exceptions) and mediated access for the other (e.g., regular), as suggested by some versions of the dual-route model.

Waters and Seidenberg (1985) discovered a generalization that accounts for these seemingly inconsistent outcomes. They noted that the lexical decision results depended on the types of words and nonwords included in a stimulus set. When the stimuli in an experiment contain only regular and exception words and pronounceable nonwords, no exception effect obtains (Coltheart et al., 1979; Waters & Seidenberg, 1985). Under these conditions, the effect of irregular spelling–sound correspondences for lower frequency words obtained with the naming task is eliminated. The situation changes when the stimuli contain a third type of item, the so-called *strange words* first studied by Seidenberg, Waters, Barnes, and Tanenhaus (1984). These are items, such as ONCE, AISLE, and BEIGE, that contain unusual spelling patterns. In a naming study, Waters and Seidenberg (1985) obtained the results presented in Figure 8. Among the higher frequency words there were again very small differences among word types; among the lower frequency items, the strange items produced the longest naming latencies, followed by exception and then regular. The model yields similar results. In a second experiment, subjects made lexical decisions to these stimuli, yielding the results in Figure 22, similar to those obtained in naming. Waters and Seidenberg then repeated these experiments deleting the strange words from the stimulus set, which eliminated the difference between regular and exception words in lexical decision but not naming (Figure 23).

Thus, phonological effects in lexical decision (the differences between regular and exception words) depend on the composition of the stimuli in the experiment; the presence or absence of strange words accounts for the seemingly inconsistent results of previous lexical decision studies. It is important to recognize that the results on the naming task are not affected by this factor; there are robust exception effects for lower frequency words, whether or not strange words are included.

Waters and Seidenberg (1985) proposed the following account of these results. When the stimuli consist of regular and exception words and pronounceable nonwords, subjects base their decisions on the results of orthographic analyses. Hence,

no effects of phonological regularity obtain. Including the strange stimuli increases the difficulty of the word/nonword discrimination. Subjects are asked to respond "word" when they see an item with an unfamiliar spelling pattern such as AISLE and to respond "nonword" when they encounter stimuli that contain common spelling patterns but are, nonetheless, not words (e.g., NUST). Making this discrimination on the basis of orthographic information is difficult; thus, subjects change their response strategy, turning to phonological information as the basis for their decisions. In effect, the subject now responds "word" if the stimulus has a familiar pronunciation and "nonword" if it does not. Thus, subjects could make correct decisions for words that are in their spoken vocabularies even when they are unsure of their spellings. Under these conditions, the task is much like naming: It requires computing the phonological code. Thus, results are similar to those in naming, with a regularity effect for lower frequency words.

Analogous results involving semantic information were reported by James (1975). When the stimuli consisted of words and very wordlike nonwords, decision latencies were faster for concrete words than for abstract ones, suggesting that subjects used semantic information in making their decisions. When the nonwords were changed to orthographically illegal letter strings, the difference between the concrete and abstract words was eliminated, suggesting that decisions were based on orthographic information alone. It can also be seen how this account generalizes to the case of targets presented in sentence contexts. If the word/nonword discrimination is difficult, subjects judge the perceived congruence of sentence context and target; they respond "word" if the target forms a meaningful continuation of the sentence, and "nonword" if the target does not (Stanovich & West, 1982). Because language comprehension normally involves integrating words and contexts, subjects find it very difficult to inhibit this process in making lexical decisions.

In sum, lexical decision allows considerably more flexibility in response strategy than does naming. In the former task, the orthographic, phonological, and semantic codes may all provide a basis for responding depending on list composition, instructions, and other experiment-specific factors. Naming is more constrained because the subject must produce the correct pro-
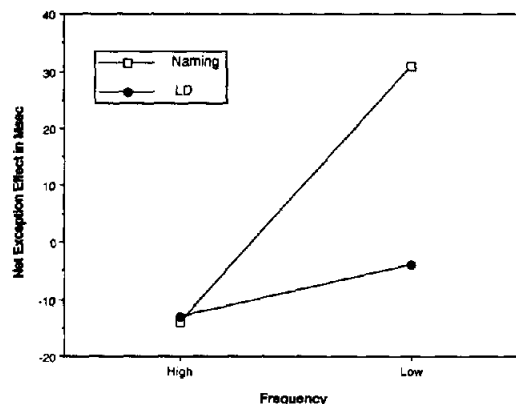

Figure 23. Magnitude of the exception effect when strange words are excluded.

nunciation, which requires computation of the phonological code.

Lexical Decisions in the Model

As noted previously, we assume that lexical decision makes use of the orthographic output that is computed in parallel with phonological and semantic output; orthographic output provides the basis for the familiarity judgment described by Balota and Chumbley (1984) in their account of lexical decision. We will explicitly examine the simplest case, in which only this orthographic input is used, but as we have noted, experimental variables will determine whether this strategy is sufficient. The subject computes a measure of orthographic familiarity by comparing the input string to the computed orthographic output. In our model this corresponds to comparing the pattern of activation produced across the orthographic units by the input to the pattern produced through feedback from the hidden units. The subject compares the obtained orthographic error score to a criterion value adopted on the basis of experience with prior word and nonword error scores, the relative frequency of words and nonwords, and instructional factors, as standardly assumed in signal detection experiments. If the error score is less than the criterion, the subject makes the word response; if greater than the criterion, the subject makes the nonword response. Words and nonwords falling on the wrong side of the criterion are assumed to be responded to incorrectly. Items with scores farther from the criterion are assumed to be responded to more rapidly than those with scores close to criterion. If information about the orthographic error scores accrues gradually over time, as we assume it does in reality, more extreme values would exceed criterion more rapidly than less extreme values (cf. Ratcliff, 1978).

This lexical decision strategy will lead to an unacceptably high error rate under some conditions, specifically when the words and nonwords are orthographically similar. Under these conditions, we assume that subjects also assess the familiarity of the stimuli in terms of the computed phonological output. Feedback from other parts of the system would provide the basis for judging the familiarity of this code. Indeed, the phonological representation computed by our existing orthography → phonology pathway can be seen as an input pattern over the phonological units. If this pattern were passed through a set of hidden units reciprocally connected to the phonological units and trained through experience with the sounds of words, the difference between the incoming phonological stimulus and this feedback could serve as the basis for a familiarity judgment.

The simulations reported in the next section are concerned with cases in which orthographic and phonological information provide a basis for making lexical decisions. This account is completely consistent with the possibility that there may be other cases in which subjects must consult information provided by the computation from orthography to semantics. Our main point is that, contrary to standard views of lexical decision, access to individuated lexical representations associated with particular words is not required by the task. Instead, information about familiarity of the pattern produced by the stimulus at one or more levels of representation provides a sufficient basis for lexical decision performance. In some cases, familiar-

ity of semantic patterns may need to be assessed, but in others, orthographic or phonological information may be sufficient. Our simulations show that this can indeed be the case, inasmuch as they indicate that we can capture the results of a number of lexical decision experiments with the existing version of the model, in which the computation of semantic representations is not implemented.

Simulation results. We tested this account by using the model to compute orthographic error scores for the Waters and Seidenberg (1985) word and nonword stimuli using, as before, the weights from 250 learning epochs. The word stimuli had been included in the 2,897-word training set. Figure 24 (top) presents the data for the condition in which the stimuli consist of high- and low-frequency regular and exception words; Figure 24 (middle) presents the data for the pronounceable nonwords. Although the distributions of error scores overlap a bit, inspection suggests that a decision criterion can be established that yields an error rate similar to that observed in the actual experiment. Because the decision can be based on orthographic output, no effect of phonological regularity is predicted. Figure 24 (bottom) presents the same data as in the top figure but with the addition of the high- and low-frequency strange items. Now there is considerable overlap between the word and nonword distributions. This is primarily because the mean orthographic error score for the lower frequency strange words is 13.1770, whereas the mean for the nonwords is 15.450, with a standard deviation of 5.610. This overlap makes it impossible to establish a decision criterion that yields an acceptably low error rate. Under these circumstances, we argue, subjects begin to look to phonological output (and possibly semantic as well). Decision latencies should now exhibit the pattern associated with the naming task, longer latencies for lower frequency exception words compared to regular. This was the result obtained in the Waters and Seidenberg (1985) experiment.

In effect, the orthographic error scores provide a measure of orthographic familiarity. The validity of this measure is supported by the observation that it accounts for other data as well. For example, the lower frequency, orthographically irregular strange words yield larger orthographic error scores than do regular or exception words. Hence, when the nonword stimuli are sufficiently unwordlike to permit an orthographic response strategy, the model predicts that strange items will still yield longer lexical decision latencies than the other types (as Waters & Seidenberg, 1985, found). This measure is also interesting because it derives from everything that the model has encoded about the frequency and distribution of letter patterns in the lexicon. Error scores are a function of the input stimulus and the weights on connections that derive from the entire training experience. Other measures of orthographic familiarity have been used in word recognition experiments (e.g., positional letter frequencies, bigram frequencies, Coltheart's N measure), with mixed results. These inconsistent results, we suggest, may be due to the fact that orthographic familiarity as it is reflected in the performance of the adult reader is better captured by the overlaid effects of the full range of experiences with the structure of words, as in our model, than by these other measures, which reflect only part of the information that is acquired through experience. It is a characteristic of this measure, and therefore an implication of our model, that the orthographic

Figure 24. Simulation of the Waters and Seidenberg (1985) studies. (Orthographic error scores for regular and exception words [top]; pronounceable nonwords [middle]; regular, exception, and strange words [bottom].)

familiarity of a letter string reflects frequency of exposure to the string itself, as well as exposures to other, orthographically overlapping letter strings.

Homographs. Additional evidence consistent with this account is provided by performance on homographs, words such as LEAD or WIND that contain common spelling patterns but are associated with two pronunciations. Thirteen homographs were included in the training set; the model was trained on both pronunciations of each word. The Kucera and Francis (1967) norms provide estimates of the overall frequencies of these words; we arbitrarily assigned a frequency equal to one half of the listed frequency to each pronunciation. Thus, the model was equally likely to receive feedback concerning both pronunciations. These words represent the limiting case in terms of orthographic–phonological inconsistency because the model is given inconsistent feedback about the entire words, not merely parts such as word bodies. Given this inconsistent feedback, it is not surprising that the model performed relatively poorly on these items, producing high phonological error scores. Even though the model was exposed to both pronunciations equally often, after 250 epochs of training it typically "preferred" one pronunciation. For example, for the word WIND, the error score for the pronunciation /wInd/ was much smaller than the score for the pronunciation /wind/, probably because the training corpus contained several /Ind/ words and no other /ind/ words. Similarly, the model preferred LEAD = /led/ and BASS = /bas/, again on the basis of regularities elsewhere in the corpus. It is interesting that human subjects asked to name isolated homographs aloud also produce very long latencies (Seidenberg, Waters, Barnes, & Tanenhaus, 1984). Presumably, the correct pronunciations of these words are normally determined by establishing which meaning is appropriate to a given context and computing the pronunciation from meaning; subjects perform poorly when contextual information is not provided, forcing them to rely on the computation from orthography to phonology, which is ambiguous. Our account of lexical decision suggests that if the stimuli consist of words containing common spelling patterns and orthographically distinct nonwords, no effects of factors related to phonology should be observed because the decision can be based on orthographic output; hence, homographs should behave like other words with common spelling patterns. This outcome has been observed empirically: Whereas homographs yield longer naming latencies than nonhomographs, they do not yield longer lexical decision latencies (Seidenberg, Waters, Barnes, & Tanenhaus, 1984). The model predicts that if this experiment were repeated with nonwords whose orthographic error scores overlapped with those of the word stimuli, the orthographic response strategy would be disabled, forcing subjects to consult phonological information as well. Under these conditions, homographs should yield longer lexical decision latencies than nonhomographs, as in naming. This prediction has not been tested, however.

In sum, the model provides a simple account of observed differences between lexical decision and naming performance. The naming task requires the subject to compute a word's phonological code; thus, it is affected by factors such as orthographic–phonological regularity. Under many conditions, the lexical decision task can be performed on the basis of orthographic information, and latencies are affected by orthographic

properties of words, but not by orthographic-phonological regularity. If the stimuli in a lexical decision experiment include very wordlike nonwords, or very unwordlike words, subjects' decisions take into account the computed phonological codes. Under these conditions, lexical decision results are like those that obtain in naming, because both responses are based on the same information.

Orthographic and Phonological Priming

The preceding account generalizes to a somewhat different phenomenon studied by Meyer et al. (1974). Rather than orthographic-phonological regularity, they examined orthographic and phonological priming effects. The stimuli consisted of words and nonwords presented in pairs. Subjects responded "yes" if both stimuli were words, and "no" if the pair contained a nonword. The word pairs included orthographically similar rhymes (e.g., BRIBE-TRIBE), orthographically similar nonrhymes (e.g., FREAK-BREAK), and unrelated control items (e.g., BRIBE-TIGHT, FREAK-TOUCH). Rhyme pairs yielded faster latencies and nonrhyme pairs slower latencies than did controls. This mixed pattern of facilitation and inhibition indicates that phonological relations between the words influenced subjects' decisions. Meyer et al. interpreted the results as indicating that processing of the prime biased the encoding of the target. Having computed the phonological code for BRIBE biased the subject to assign the same code to TRIBE; this strategy yielded interference when the stimuli were nonrhymes such as FREAK-BREAK. However, Hillinger (1980) obtained facilitation on trials containing rhymes with different spellings (e.g., CAKE-BREAK), suggesting that phonological relations between words affect subjects' decision strategies rather than target encoding.

According to our account, phonological information will bias lexical decisions only when the use of orthographically based decision criteria is disabled because of the similarity of words and nonwords along this dimension. It is easy to see why Meyer et al.'s (1974) stimuli would have this effect; the stimuli included word/nonword pairs such as MOIST-SOIST and DRUNK-FRUNK, which differ by only one letter. It follows from our account that phonological information would not be used if the word and nonword stimuli were more discriminible in terms of orthography. Shulman, Hornak, and Sanders (1978) reported this result. They replicated the Meyer et al. study using the same word stimuli but varying the properties of the nonwords, which were either pronounceable pseudowords (like Meyer et al.'s) or random letter strings. With pseudoword stimuli, the results replicated Meyer et al.'s, with facilitation for orthographically similar rhymes and inhibition for orthographically similar nonrhymes, indicating the use of phonology. With random letter strings as nonwords, there was facilitation for both rhymes and nonrhymes, indicating the use of orthographic but not phonological information.

Frequency Blocking Effect

Glanzer and Ehrenreich (1979) and Gordon (1983) reported a seemingly anomalous lexical decision phenomenon called the *frequency blocking effect*, which can also be understood within the account of lexical decision performance given earlier. The

Table 7
Results of the Gordon (1983) Frequency Blocking Experiment

List type	Word frequency class		
	Low	Medium	High
Mixed-frequency	710	566	520
% error	8.9	0.3	0.1
Pure-frequency	710	547	480
% error	8.1	0.5	0.2
Difference (in ms)	0	19	40

Note. Main entries are lexical decision latencies in milliseconds.

phenomenon is the finding that, in this task, the magnitude of the effect of frequency depends on the composition of the stimuli in an experiment. Gordon (1983), for example, reported an experiment in which the stimuli were high-, medium-, and low-frequency words, presented in either mixed or blocked conditions. In the mixed condition, stimuli from all three frequency bands were randomly intermixed; in the blocked conditions, the same stimuli were presented, but blocked according to frequency. Gordon's results are given in Table 7. In both conditions, there were frequency effects, with the order of lexical decision latencies being high < medium < low. Whereas latencies for the lower frequency words were identical in the mixed and blocked conditions, they were faster in the blocked condition than in the mixed condition for both medium- and high-frequency words. This change in the magnitude of the frequency effect is the frequency blocking phenomenon. Gordon presented a signal detection model, much like the one given earlier, in which subjects vary their decision criteria in response to the properties of the stimulus set.

We simulated Gordon's experiment by computing the orthographic error scores for high-, medium-, and low-frequency words like the ones used in his experiment. There were 24 items of each type, matched in length. We also tested 69 pronounceable nonwords similar to the ones he used. The distributions of orthographic error scores are presented in Figure 25. Assume that decisions are based on a weighted combination of orthographic and other types of information (e.g., phonological and/or semantic). As the overlap between words and nonwords in terms of orthography decreases, subjects should weigh orthographic information more heavily. As the overlap increases, subjects should weigh the other types of information more heavily. When the stimuli are intermixed, the distributions for words and nonwords show considerable overlap, predicting that the lexical decision should be difficult. Under these conditions, subjects might be expected to weigh the other types of information more heavily in making their responses. The overlap is due primarily to the lower frequency words, some of which produce error scores like the pronounceable nonwords. Hence, presenting only low-frequency words and pronounceable nonwords would not facilitate performance, as Gordon (1983) observed. The situation improves when medium- and high-frequency words are blocked. Because the distribution for the high-frequency words overlaps little with the nonwords, blocking would allow the subject to establish decision criteria based on ortho-

graphic information alone; other types of information would not need to be consulted. Orthographic information is closer to the input stimulus than either phonological or semantic information; therefore, decisions based on this code should be more rapid. Because the distributions for the medium-frequency words and nonwords overlap a bit more, blocking would yield a smaller benefit. These predictions are entirely consistent with Gordon's results.[11] One other point should be noted. The frequency blocking phenomenon derives from the fact that lexical decision performance depends on the discriminibility of word and nonword stimuli. Because naming depends on the computation of phonological output, rather than the discriminibility of words and nonwords, it follows that frequency blocking should have little effect on naming performance. Forster (1981a) reported this result, providing strong support for this analysis of the differences between the tasks.

Pseudohomophone Effects

The final simulations concern the processing of pseudohomophones—nonwords such as BRANE or PRUVE that sound like words. Pseudohomophone effects refer to differences in naming or lexical decision latencies for these items compared to nonpseudohomophones such as BRONE or PRAVE. Performance on these stimuli has been thought to provide evidence concerning the role of phonology in the access of meaning. In the original study employing these stimuli (Rubenstein, Lewis, & Rubenstein, 1971), subjects performed the lexical decision task. On nonword trials, latencies were longer for pseudohomophones such as BRANE than for nonpseudohomophones. Rubenstein et al. assumed that the task was performed by determining whether the target stimulus has a meaning. Thus, the longer latencies for stimuli such as BRANE suggested that the stimulus was phonologically recoded and that this phonological code was used to access the meaning associated with BRAIN. This information interfered with the decision that BRANE is a nonword. Latencies for pseudohomophones derived from high- and low-frequency words (e.g., BRANE, high frequency; BRUME, low frequency) did not differ. Subsequent studies of pseudohomophone effects have yielded inconsistent results (see, e.g., Coltheart et al., 1977; Dennis et al., 1985; Van Orden, 1987). McCann and Besner (1987; McCann, Besner, & Davelaar, 1988) recently reported three findings concerning these stimuli. First, when the task was to name the stimuli aloud, pseudohomophones yielded faster latencies than did nonpseudohomophones. Second, when the task was lexical decision, the pattern was reversed: Pseudohomophones yielded longer latencies than nonpseudohomophones. Third, neither the lexical decision nor naming latencies for pseudohomophones were correlated with the frequencies of the base words from which they were derived. That is, the latency to name or make a lexical decision to an item such as BRANE was unrelated to the frequency of BRAIN. Besner and McCann interpreted these results in terms of a model concerning the role of frequency in lexical access.

These results are relevant to the model we have proposed for the following reason. Pseudohomophone effects are thought to reflect the influence of the lexical entry for the base word on the pseudohomophone. That is, BRANE differs from BRONE because only BRANE is influenced by a neighboring homophone. BRAIN

facilitates the naming of BRANE but interferes with making a lexical decision to it. Pseudohomophone effects would appear to be a problem for our model because it lacks word-level representations; there does not seem to be a way for the spelling or pronunciation of BRAIN to directly influence BRANE because there is no lexical entry for BRAIN.

It is interesting to note, however, that the model actually performs differently on McCann and Besner's (1987) pseudohomophone and nonpseudohomophone stimuli. When the stimuli (which were nearly identical in the two studies) were tested on the model, the pseudohomophones yielded smaller orthographic and phonological error scores. Hence, the model predicts that they should be easier to name and yield longer lexical decision latencies, just as McCann and Besner found.

The model simulates these effects because it is sensitive to a general difference between the two types of stimuli: Pseudohomophones tend to be more wordlike than the nonpseudohomophones. That is, the pseudohomophones tend to contain spelling patterns and spelling–sound correspondences that occur more often in words; hence, they are better approximations to actual words. This tendency derives from two factors. First, some pseudohomophones benefit from the model's exposure to orthographically similar base words. Training on a word such as BRAIN or CAUGHT tends to modify the weights in a direction that facilitates processing on pseudohomophones such as BRANE or COUGHT. The magnitude of this effect will depend on the similarity of pseudohomophone and base word; much smaller effects will occur for dissimilar pairs such as CAUGHT and CAWT. Second, pseudohomophones tend to be more wordlike because of constraints that govern the construction of the stimuli. The constraint that pseudohomophones sound like words may require using more of the spelling patterns and spelling–sound correspondences that actually occur in words; conversely, the constraint that nonpseudohomophones not sound like words may require using structures that do not occur very often. Because the error scores reflect the aggregate effects of exposure to a large vocabulary of words, they tend to pick up on these systematic differences between the stimuli.

In short, the model produces pseudohomophone effects because these stimuli tend to be closer approximations of words than are the nonpseudohomophone controls. Still, it is possible that there could be pseudohomophone effects above and beyond those accounted for by general orthographic and phonological properties of the stimuli. If the processing of a target such as BRANE were influenced by the entry for a word such as BRAIN, the model would fail to pick up this effect. Hence, there might be differences between the stimuli even when they are equated in terms of the error scores generated by the model. On the other hand, the model predicts no differences between the two types of stimuli if they are equated in terms of error scores. We tested these predictions by using the orthographic and phonological error scores generated by the model to create two sets of stimuli. In the unbalanced set, the stimuli were like the ones in

[11] Even with the most optimistic setting of the decision criteria, the simulation predicts somewhat more errors in the medium- and high-frequency conditions than Gordon (1983) actually observed. However, it should be noted that this simulation did not use his stimuli because they were not published with the study.
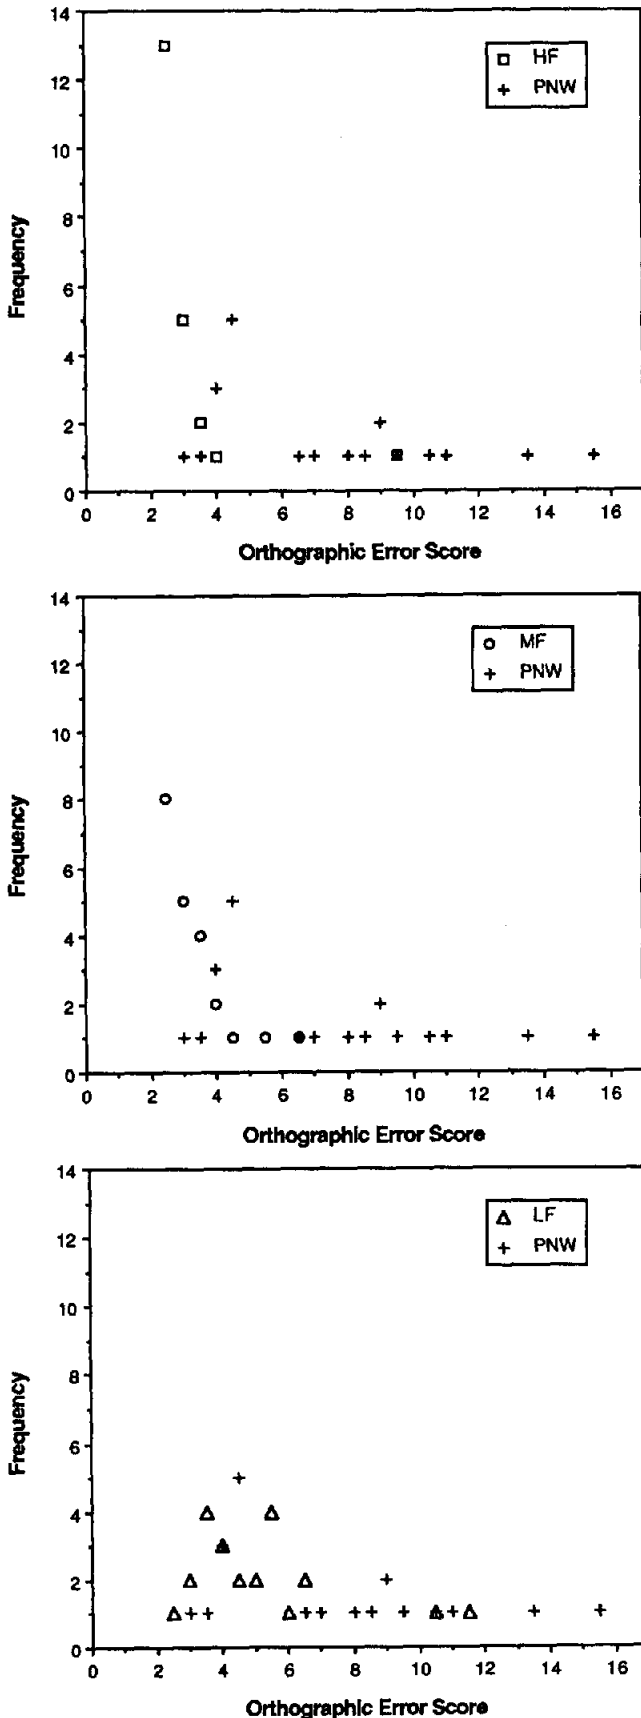
the McCann et al. (1988) studies, in that the pseudohomophones produced significantly smaller orthographic and phonological error scores than the nonpseudohomophones. In the balanced set, the two types of nonwords were equated in terms of both error scores. In the lexical decision version, 24 subjects were presented with all of the stimuli randomly intermixed with a set of monosyllabic words. In the naming version, a second group of 24 subjects were presented with each nonword and required to name it aloud. Results for the unbalanced stimuli (Figure 26, top) replicate the McCann and Beser (1987) and McCann et al. (1988) findings: Pseudohomophones were easier to name than nonpseudohomophones, but yielded longer lexical decision latencies. This pattern did not replicate with the balanced stimuli, however (Figure 26, bottom). There was a main effect of task, with faster latencies on naming than on lexical decision, but no interaction with type of nonword.

In sum, the model replicates the pseudohomophone effects in the McCann et al. (1988) studies, even though it does not contain explicit lexical entries to influence pseudohomophone processing. These effects are realized in the model's error scores, which reflect the extent to which pseudohomophones and nonwords resemble words in the lexicon. Our experiment suggests that the general tendency for pseudohomophones to be closer approximations to words can be eliminated by other facts that affect the error scores. The error scores are effective because they provide summary measures that capture influences that arise not only from experience with a particular word, but also with other words that overlap with it in a wide variety of ways.

Summary of the Lexical Decision Simulations

The model gives a good account of simple word/nonword discrimination, including some more subtle phenomena related to changes in decision criteria, as well as differences between naming and lexical decision. Several points emerge from this analysis. First, we have shown that the model can account for lexical decision performance despite the absence of word-level representations. This represents a substantial change from previous accounts that assumed that lexical decisions involved accessing such representations. The simulations also show that the types of knowledge representations that we found useful in accounting for naming performance can support the lexical decision process.

A second point is that the types of information used in making lexical decisions vary systematically in response to properties of the stimulus set. Under the conditions that are characteristic of many lexical decision experiments, subjects can base their decisions on orthographic information alone. When this strategy is disabled, they can use phonological information. In principle there should be other conditions in which semantic information must be consulted. The model provides an inde-

Figure 25. Simulation of the frequency blocking effect. (Distributions of orthographic error scores for higher frequency words and orthographically legal, pronounceable nonwords [top]; medium frequency words and nonwords [middle]; and lower frequency words and nonwords [bottom].)

Figure 26. Replications of the McCann, Besner, & Davelaar (1988) pseudohomophone effects: Experiment and simulation data. (Upper graph: Stimuli that are not equated in terms of error scores. Lower graph: Stimuli equated in terms of error scores.)

pendent basis for determining when orthography will or will not provide a sufficient basis for the decision, allowing us to correctly predict when lexical decision results will or will not mimic those obtained with naming.

The model also suggests that under the conditions that often obtain in single-word studies, lexical decisions can be based on nonsemantic types of information. This observation is important because it calls into question the assumption that lexical decision performance necessarily provides evidence concerning processes related to the access of meaning. Our model accounts for performance in many single-word studies even though it contains no representation of meaning at all. Of course, subjects can ultimately determine that words have meanings and that nonwords do not; in our model, this information would be provided by the computation from orthography to semantics. However, the lexical decision process does not actually proceed on this basis under many conditions.

Finally, it is clear that the conditions we have examined (concerning the presence or absence of strange words and the use of pronounceable nonwords vs. random letter strings) do not

exhaust the range of possible circumstances afforded by the lexical decision paradigm. Our main point is that the results of any given experiment must be interpreted in regard to the response strategies permitted by the stimulus conditions. The results of each experiment represent a point in a space of possibilities determined by the properties of the stimuli, instructions to the subjects, and other experiment-specific factors. The complexity of the task increases greatly when targets appear in sentence contexts. A more complete theory than ours would provide an account of the types of information and decision processes involved in the judgments of contextual congruity typical of performance in sentence context experiments (see Forster, 1981b; Stanovich & West, 1981, for discussion).

General Discussion

The model of lexical processing that we have described can be summarized in terms of a number of main features. Lexical processing entails the computation of several types of output in parallel. We have described the computation of the orthographic and phonological codes in some detail and shown that the model provides a quantitative account of various behavioral phenomena. The model accounts for differences among words in terms of processing difficulty, differences in reading skill, and facts about the course of acquisition. Lexical decision and naming are characterized in terms of how the computed codes are used in making these types of responses. A task such as naming focuses on the use of one type of code, phonology; a task such as lexical decision may involve all of the codes. The same types of knowledge representations and processes are involved in the computation of all three codes (although the implemented model is restricted to orthography and phonology). Knowledge is represented by the weights on connections between units. These weights are primarily determined by the nature of the English orthography that acts as input, in conjunction with feedback during the learning phase. Our claim is that representing knowledge of the orthography in this way is felicitous given the quasiregular nature of the system; the characteristics of English orthography are more congruent with this type of knowledge representation than with the kinds of pronunciation rules proposed previously. The computation of the orthographic code is affected by the facts about the distribution of letter patterns in the lexicon; computation of the phonological code is affected by facts about correlations between orthography and phonology.

The main theoretical implications of the model can be characterized in terms of a number of recurring issues in reading research.

Role of Phonology in Word Recognition

A large amount of research has been directed at questions concerning the use of phonological information in visual word recognition. Three issues have been studied, although they have not always been distinguished. One concerns access to phonology: Does the processing of a word necessarily result in access to phonological information? The second concerns the nature of the computation involved in accessing phonology: What kinds of knowledge are involved and is there a single process or

more than one? The third issue concerns the relation between phonological access and meaning: Is the phonological code computed as part of the process by which the meaning of a word is identified?

Concerning the first issue, the primary question is whether access of phonological information is an automatic consequence of processing or the result of a recoding strategy under the control of the perceiver. Clearly, the task of understanding a text does not necessarily require access of phonological information, and the task can be accomplished by individuals who lack any knowledge of orthographic–phonological correspondences at all (e.g., nonspeaking deaf persons). It might nonetheless be useful to access phonological information, as suggested by early information-processing models of memory such as Atkinson and Shiffrin (1968), which proposed that subjects recode visual stimuli into phonological representations for the purpose of retaining information in short-term memory, consistent with the results of studies such as Conrad (1964). Thus, phonological recoding was thought to be a strategy relevant to maintaining information in short-term memory, rather than a necessary consequence of stimulus encoding. Some reading researchers retained this idea and attempted to identify the factors that determined when phonological recoding was used. For example, it was proposed that phonological information might be used for certain types of words (e.g., regular rather than exception; Coltheart, 1978), by certain types of readers (e.g., poor readers: Jorm & Share, 1983; good readers: Barron, 1981), or for certain tasks (e.g., naming rather than lexical decision; Coltheart et al., 1979).

Our model differs from these proposals in that it incorporates the idea that visual word recognition results in the activation of phonological information in parallel with other representations (Donnenwerth-Nolan et al., 1981; Seidenberg & Tanenhaus, 1979). In acquiring word recognition skills, children learn to associate the orthographic codes for words with both their meanings and pronunciations. Once this skill is acquired, processing of a written stimulus results in activation of multiple types of information, even though only one may be required for performing a given reading task. Tversky and Kahneman (1983) have observed other phenomena of this type. Their studies show that individuals find it difficult to ignore information that is correlated with information that is relevant to problem solving but is not itself relevant to the solution. According to this view, activation of phonological information is a result of stimulus encoding processes rather than recoding strategies. What varies is whether this information is used in performing tasks such as lexical decision, as illustrated by the experiments we simulated earlier. The activation of phonological representations in parallel with meaning may account for the "voice in the head" experienced by many individuals in silent reading.

Additional support for this view is provided by studies such as Tanenhaus et al. (1980), in which a modified Stroop paradigm created a situation in which access of phonological information had a *negative* effect on performance. This result is inconsistent with the idea that access of phonological information is due to a subject strategy intended to facilitate performance. Rather, subjects accessed this information even when it was optimal to avoid doing so. The ubiquitous effects of phonological information on various reading tasks observed by Baron (1979), Klei-

man (1975), and others simply reflect the fact that phonological information, like meaning, is rapidly activated in reading; they further show that this information is used in performing tasks such as making a lexical decision or judging the meaningfulness of an utterance.[12]

In regard to the nature of the computation involved in accessing phonology, our model refutes what Seidenberg (1988) has termed the *central dogma* linking different versions of the dual-route model of naming, namely, that separate processes are required for naming exception words on the one hand and novel items on the other. Our model demonstrates that a single computation that takes spelling patterns into phonological codes is sufficient to account for naming of these types of items and others. Moreover, it provides an explicit account of quantitative differences between stimulus types in terms of naming difficulty.

Note, however, that within the architecture illustrated in Figure 1 there is a second, indirect way to generate the pronunciations of words: by computing the meaning of a word from orthography and computing its pronunciation from meaning, as in speech production. In this respect, our account is similar to the dual-route model, which also holds that there are two ways to pronounce letter strings. It is important to recognize the differences between the models, however; they are not notational variants (Seidenberg, 1988, in press-b). The evidence that there is a second naming mechanism is compelling; as we have noted, the indirect method is relevant to generating the contextually appropriate pronunciations of homographs such as WIND. Moreover, the indirect method is implicated in certain types of dyslexia that occur following brain injury. For example, so-called *phonological dyslexics* are able to name familiar words but are impaired in naming nonwords (Shallice & Warrington, 1980). This would follow if the patient's capacity to compute pronunciations from orthography were impaired but the indirect route from orthography to meaning to phonology were not. Perhaps the primary difference between the two models concerns the role of the indirect route in normal reading. According to the dual-route model, words with irregular pronunciations can *only* be pronounced by the indirect method. This follows from the assumption that readers' knowledge of spelling–sound correspondences is represented in terms of rules that, by definition, are only capable of generating the pronunciations of regular words and nonwords. In our model, knowledge of spelling–sound correspondences is represented in terms of the weights on connections between units involved in the computation from orthography to phonology. As we have demonstrated, this type of knowledge representation is sufficient to

---

[12] The Tanenhaus, Flanigan, and Seidenberg (1980) results, and related phenomena such as the visual tongue-twister effect (McCutchen & Perfetti, 1982), have suggested that subjects cannot shut off phonological processing completely, even when it would be beneficial to do so. However, it may be that this computation can be regulated to some extent. Cohen, Dunbar, and McClelland (1989) have recently proposed a model of attention that has this implication. For example, the instruction to attend to colors of Stroop stimuli may facilitate the encoding of this information. Thus, although phonological information is activated under a broad range of conditions, the manner in which it is computed may vary.

account for facts about the pronunciation of regular and irregular words and nonwords. Moreover, the type of computation we have described is necessary in order to account for consistency effects of the type illustrated in Figure 19. The dual-route model is silent about cases in which the pronunciation of a putatively rule-governed word is influenced by knowledge of words not covered by the rule. In sum, there are similarities between the dual-route model and the account presented here, but the models use different types of knowledge representations and processes and make different predictions about inconsistent words. Ours is a dual-route model, but it is not an implementation of any previous model.

The picture is similar when we turn to the third issue, the role of phonological information in accessing the meanings of words, probably the single most widely studied question in reading research. A large number of studies have been directed at distinguishing between direct and phonologically mediated routes to meaning (see Carr & Pollatsek, 1985; Henderson, 1982; McCusker et al., 1981, for reviews). The direct access hypothesis is that readers recognize a letter pattern as a particular word, providing access to a representation of its meaning stored in semantic memory. The phonological mediation hypothesis holds that readers first compute the phonological code for a word and then use this code to search semantic memory. Despite extensive research, empirical studies have not yielded a clear resolution of the issue (contrast, e.g., Baron, 1973; Van Orden, Johnston & Hale, 1988). The model presented in Figure 1 provides a framework for integrating many of the conflicting results in the literature. As Figure 1 indicates, the model entails computations from orthography to meaning and from orthography to phonology. The default assumption, then, is that meanings are activated on the basis of a direct computation from orthography. The computation from orthography to phonology occurs in parallel, however, with the result that the phonological code becomes available and, as suggested earlier, it can influence performance on many tasks, even when it is not logically required. This aspect of the model underscores an ambiguity in much of the research on phonological mediation: Many studies have provided evidence that subjects use phonological information in reading, but as the model suggests, this fact does not itself necessarily indicate that access of meaning was phonologically mediated. In general, it has proven difficult to empirically discriminate between activation of phonological information and phonologically mediated access of meaning.

These two assumptions—that there is a direct computation from orthography to meaning and a separate, equally direct computation from orthography to phonology—are consistent with a large body of empirical findings in this area. However, the framework presented in Figure 1 also affords the possibility that phonological information could influence the activation of meaning by means of feedback from the computed phonological code, the third side of the triangle in Figure 1. Just as there is an indirect route from orthography to meaning to phonology, there is an indirect route from orthography to phonology to meaning. Other factors being equal, the feedback from phonology to meaning should develop relatively slowly because it requires a prior computation from orthography to phonology. Thus, feedback from phonology to meaning should depend on amount of time available for this process to occur (Seidenberg,

1985b, 1985c). In general, this feedback will have an effect when the primary computation from orthography to meaning is itself relatively slow. There are a number of conditions under which this might occur. For example, readers are sometimes more familiar with the pronunciation of a word than its spelling. In such cases, the computation from orthography to meaning might fail to yield a clear pattern, but the reader could attempt to determine the word's meaning from phonology. This process may be characteristic of children in the earliest stages of learning to read, who identify the meanings of words by sounding them out, matching the phonological codes that are generated to words in their spoken vocabularies. Similarly, the computation from phonology to meaning might be used when it provides information relevant to performing a particular task. For example, if subjects are required to make a difficult lexical decision or categorization judgment, the information provided by feedback from phonology to meaning may provide an additional basis for responding (e.g., Van Orden et al., 1988). In general, feedback from phonology to meaning should be associated with words that have unfamiliar spelling patterns, readers who are relatively poor at computing meanings from orthography, conditions under which accessing the information facilitates performance, or difficult tasks that yield relatively long response times (Seidenberg, 1985b, in press-a).

One other case should be mentioned. Several researchers have examined the hypothesis that the extent to which phonology is used in accessing the meanings of words depends on the properties of the orthography. More phonological mediation is thought to be observed in the reading of "shallow" orthographies with relatively simple and direct spelling–sound correspondences (e.g., Serbo-Croatian; Katz & Feldman, 1981; Turvey, Feldman, & Lukatela, 1984). "Deep" orthographies are thought to discourage the use of phonology in accessing meaning. The model presented in Figure 1 provides a framework for considering both the universal and language-specific aspects of processing. We assume that this general architecture underlies visual word recognition in all languages. Differences among orthographies are realized in terms of the characteristics of the orthographic and phonological encodings and the weights on connections between units. The weights on the connections between orthographic and phonological units will reflect the degree of consistency or regularity of spelling–sound correspondences in a given orthography. Other factors being equal, then, the computation from orthography to phonology should be more rapid in the "shallow" orthographies, allowing more opportunity for feedback from phonology to meaning. Note, however, that many other factors need to be considered before concluding that there is more phonological mediation in a given orthography. There are other differences between orthographies that could also influence the difficulty of the computation from orthography to phonology; for example, languages differ in terms of the average length and number of syllables per word. Moreover, within the framework given in Figure 1, there is also a direct route from orthography to meaning, which may be used even in "shallow" orthographies (see Seidenberg, in press-a, for discussion).

In sum, many of the controversies in the study of visual word recognition have been concerned with the questions concerning the number of processes involved in identifying the meanings

or pronunciations of words. The framework presented in Figure 1 clarifies how these questions are related. Both of the codes can be derived on the basis of primary, direct computations from orthographic input. In both cases, however, there is an indirect method of generating the relevant code. The existence of both direct and indirect routes is a consequence of the architecture presented in Figure 1, which reflects interconnections among the readers' knowledge of the written, spoken, and semantic codes for words.

The Lexicon and Lexical Access

Our model differs from previous accounts in regard to the manner in which lexical knowledge is represented and processed. A standard view, common to models such as Coltheart (1978), Forster (1976), Morton (1969), and others, is that lexical memory consists of entries corresponding to the different codes of words. For example, Forster (1976) suggested that lexical memory consists of a set of files or bins, including a master file containing entries for all of the vocabulary items, and slave files containing entries for different codes (e.g., a file containing word pronunciations). The models described by Coltheart (1987) and Monsell (1987) contain multiple lexicons, including separate orthographic lexicons used in reading and writing, and separate phonological lexicons used in listening and speaking. Research within this framework has focused on questions concerning what has been termed *lexical access:* how the entries for different codes are accessed in reading, the order in which they are accessed, and how access of one code affects access of other codes.

The present model departs from these precursors in a fundamental way: Lexical memory does not consist of entries for individual words; there are no logogens. Knowledge of words is embedded in a set of weights on connections between processing units encoding orthographic, phonological, and semantic properties of words, and the correlations between these properties. The spellings, pronunciations, and meanings of words are not listed in separate stores; hence, lexical processing does not involve accessing these stored codes. Rather, lexical information is computed on the basis of the input string in conjunction with the knowledge stored in the network structure, resulting in the activation of distributed representations. Thus, the notion of lexical access does not play a central role in our model because it is not congruent with the model's representational and processing assumptions.

The view that lexical processing involves the *activation* of different types of information rather than *access* to stored lexical codes represents more than a change in terminology. Access to a lexical code is often taken to be an all-or-none phenomenon, whereas our alternative framework replaces this concept with a partial or graded activation of representations. In an activation model with distributed representations, a code is represented as a pattern of activation across a set of units. The activations of the units can differ in strength. Moreover, the representations in our model are not "lexical" in two senses: The units of representation do not correspond to words, and they support the processing of nonwords as well as words. These conceptions raise different questions and generate different empirical predictions. For example, within the access framework, it is rele-

vant to ask how many of the meanings of an ambiguous word are accessed; Swinney (1979; Onifer & Swinney, 1981) has proposed that lexical access results in all of the meanings of an ambiguous word becoming available with equal strengths. In contrast, a network with distributed representations, such as ours, affords the possibility of partial activation of one or more meanings (see Hinton et al., 1986; Hinton & Sejnowski, 1986; Kawamoto, 1988; McClelland & Kawamoto, 1986; McClelland & Rumelhart, 1985). The latter view is more congruent with evidence concerning the effects of contextual information on the activation of meaning (Barsalou, 1982; Burgess, Tanenhaus, & Seidenberg, 1989; Schwanenflugel & Shoben, 1985; Tabossi, 1988).

Similarly, within the lexical access framework, research has focused on whether factors such as frequency influence lexical access or postaccess processes involved in making lexical decisions or in naming words aloud (Balota & Chumbley, 1984, 1985; McCann & Besner, 1987). In our model, there is no lexical access stage common to all word recognition tasks; there are simply orthographic, phonological, and semantic computations. Within this framework, the primary question concerns how the readers' knowledge of the correlations among these codes is represented, how they are computed, and how the computed codes are used in performing different tasks. Frequency—the reader's experience in reading, hearing, and pronouncing words—affects these computations, but there are no separate effects due to lexical access.

In sum, the notion of *lexical access* carries with it a concern with certain types of theoretical questions. The primary questions concern the number of lexicons, how they are organized and linked, and whether it is orthographic or phonological information that provides access to meaning. The primary processing mechanism is search through one or more ordered lists. In our model, the codes are distributed, they are computed on the basis of three orthogonal processes, and the primary processing mechanism is spread of activation. The primary theoretical questions concern the properties of these computations, which are determined by the properties of the writing system that are picked up by the learning algorithm on the basis of experience.

If, in keeping with much of previous usage, we take the term *lexical access* to refer to access of information concerning the meanings of a word, then an implication of our model is that neither naming nor lexical decision latencies necessarily reflect this process. The model simulates many aspects of single-word naming and lexical decision performance even though meaning is not represented at all. Naming simply involves a direct mapping from spelling to pronunciation. Lexical decision often involves simply a judgment based on nonsemantic properties of the word and nonword stimuli. Hence, the results of experiments using these tasks may have no direct bearing on the question, How do readers access the meanings of words from print? The model calls into question the common assumption that these tasks necessarily provide evidence as to how readers identify the meanings of words.

Acquisition of Reading Skill

The model suggests that learning to read words involves learning to compute orthographic, phonological, and semantic

codes from visual stimuli. Acquiring this skill is a function of three factors: the nature of the stimulus, the nature of the learning rule, and the architecture of the system.

Nature of the stimulus. The model suggests that learning to read involves creating a network structure that encodes facts about the orthography. The model works as well as it does because it is trained on a significant fragment of written English, which contains a complex latent structure. Measures of orthographic redundancy (such as positional letter frequencies and bigram frequencies), lists of spelling–sound rules (such as Venezky, 1970), and definitions of regularity or phonological neighborhoods (e.g., Parkin, 1982) are partial characterizations of what is actually a very complex correlational structure concerning relations between letters and between letters and phonemes. Like the child learning to read, the model is exposed to this complex input in the training phase.

The learning rule. This elaborate structure would be of no importance were it not for the fact that there is at least one learning algorithm (there may be more) capable of extracting it. The effect of the learning rule is that the weights on connections come to encode regularities present in the input. This is a good thing to be doing if the input does in fact exhibit a rich set of regularities. It is an especially good thing to be doing if the regularities are statistical (as in written English) rather than categorical (as in rules, as they are normally construed). Thus, there is a good match between what the learning algorithm does and what is to be recovered from the input.

The architecture of the system. We have demonstrated that the model's capacity to simulate human behavior critically depends on one aspect of the architecture, the number of hidden units. This aspect of the model illustrates what may be a general characteristic of connectionist models. In order to capture facts about human behavior, the models apparently have to obey a kind of "Three Bears" principle concerning computational resources. The experiments with the number of hidden units suggest that if there are too few, the model will learn some of the basic regularities but will not be able to cope well enough with exceptions. Although we have not established this point in regard to the present model, it is known that in some cases, networks with too many hidden units "memorize" the training examples, but fail to extract implicit regularities, and thus lack the ability to respond to novel inputs (Hinton, 1986). Apparently, the number of hidden units has to be "just right," to capture both the regularities and the exceptions as people do. A detailed understanding of these characteristics of network models will require considerable mathematical analysis of network capabilities. In the meantime, the empirical discovery that something as general as the number of hidden units contributes in specifiable ways to the solution of a problem is interesting insofar as it suggests how biological constraints—the human architecture—influence what is learnable.

In sum, it will probably turn out that having the right amount of computational machinery (and the right organization of that machinery) is necessary to be able to encode the regularities that are found in the input and extracted by the learning algorithm. There may be other general architectural constraints as well.

The characterization of our model in terms of environment, learning rule, and architecture provides a useful framework for thinking about other connectionist models and about behavior in general because it incorporates some of the most important approaches to understanding behavior that have emerged in modern psychology. With Gibson, it shares the emphasis on understanding the structure of the input. With learning theory, it shares the notion of general laws of learning. With Chomsky, it shares an emphasis on how biological constraints contribute to what is learnable. Which of these elements contributes most to the solution of a given problem will probably vary. In the case of learning to read and pronounce written English, the biological constraints are probably fairly minimal: The system has to devote the right amount and kind of resources to the problem. The solution is largely driven by the highly structured input and the power of the learning rule. In language acquisition, where the input to the system is thought to be impoverished relative to what is learned, biology may impose stronger constraints on the solution space. Thus, depending upon the nature of the problem, one or another component may contribute more or less to its solution; nonetheless, all three need to be considered.

Generality of the Simulation Results

It is important to consider the generality of the conclusions we have reached on the basis of the model's performance, an issue that arises in connection with every simulation model. Our concerns focus on two aspects of the simulations. First, the model's scope is limited; it deals with only some aspects of visual word recognition. Second, there are questions as to how specific aspects of the implementation contribute to the model's performance. Both of these factors could limit the generality of the results. For example, the model might perform as well as it does only because it deals with only selected phenomena; similarly, it might perform very differently if certain features of the implementation were changed.

Scope limitations. The model's scope is restricted in three primary respects: (a) It is only concerned with monosyllabic words, (b) we have not implemented a process that yields an articulatory–motor response on the basis of the computed phonological code, and (c) we have not addressed issues related to meaning. Our primary concerns are whether these limitations compromise the conclusions that we have drawn and whether the model would need to be changed in important ways in order to deal with them.

The restriction to monosyllabic words could be important for two reasons. First, it might be that the model performs as well as it does only because the learning problem has been constrained in this way. It is possible, for example, that the learning algorithm would function much differently if the model were exposed to a wider variety of words. If the set of monosyllabic words is more homogeneous than the set of words in English, this might contribute in important ways to the behavior of the model. This is an empirical question that awaits further experimentation with this model and others like it. We should note, however, that we obtained essentially similar results for simulations using lists of 1,200 and 2,897 monosyllabic words; although the larger list was more heterogeneous, this fact had little effect on its behavior. Moreover, Lacouture (1989) has developed a model similar to ours based on a training corpus of 2,100 words, including both mono- and multisyllabic items. This

model exhibits similar behavior on monosyllabic words even though the training corpus is quite different. Hence, it does not appear that our results are specific to the particular corpus that we used or to the use of only monosyllabic words.

A second issue is that complex words exhibit additional types of structure, such as syllables and morphemes, which could be relevant to processing. Moreover, the pronunciation of multisyllabic words raises difficult issues concerning the assignment of syllabic stress. There have been a large number of studies examining the role of structures such as syllables and morphemes in visual word recognition (see Seidenberg, 1989, for discussion). These studies have led to models in which the processing of complex words involves parsing into sublexical syllabic or morphemic components. For example, Spoehr and Smith (1973) obtained evidence that syllables play a role in tachistoscopic recognition and proposed a model in which word recognition involves the recovery of syllabic structures. Other studies have been taken as providing evidence that words are decomposed into component morphemes as part of the recognition process (e.g., Murrell & Morton, 1974; Taft, 1985). Treiman and Chafetz (1987) have provided evidence indicating the salience of subsyllabic onset and rime units. This research would seem to require representations of syllables, morphemes, and onset/rime that are accessed as part of the recognition or pronunciation of letter strings. This would represent a substantial elaboration of our minimal model for monosyllabic words.

We consider these to be unresolved questions. Clearly, the model in its current form is silent about the complex processes involved in assignment of syllabic stress. The basic question is whether these phenomena can be accommodated by extensions to the present model or whether they require a model with very different types of representations and processes. For example, stress assignment is determined in part by grammatical category, a type of knowledge the current model lacks. However, it is easy to imagine extensions to the model in which grammatical category is directly encoded and learned according to similar principles. Similarly, in some theories, stress is represented by a feature associated with the representations of vowels (Chomsky & Halle, 1968), which could be accommodated by adding a feature to the scheme used here to encode phonemes. More recent theories, however, suggest that stress assignment involves access to an explicit syllabic level of representation (see Selkirk, 1980, for discussion), which might entail a major modification of the present account.

These issues can only be addressed by further research. However, there is good reason to think that a model very much like ours could account for at least the effects of sublexical structures such as syllables, morphemes, and onset/rime that have been observed with tasks such as lexical decision and naming without additional representational or processing assumptions. Specifically, the model may provide an account of the effects of complex word structure that is an alternative to parsing rules. Studies of the role of syllables and morphemes in visual word recognition have yielded inconsistent results, with some yielding evidence for decomposition into these components, whereas others have not (see Henderson, 1982; Seidenberg, 1989, for reviews). These inconsistent results may indicate that what is relevant to processing is not syllables or morphemes, but properties of words that are *correlated* with these structures. As we ob-

served at the beginning of this article, syllables and morphemes are inconsistently realized in English orthography. Just as the properties of written English make it difficult to formulate a set of rules governing orthographic–phonological correspondences, they also make it difficult to formulate parsing rules that will yield the correct decomposition into component parts. Moreover, there has been little agreement among linguists concerning the definition of the syllable (see Hoard, 1971; Kahn, 1976; Seidenberg, 1987; Selkirk, 1980). The inconsistency of spelling–sound correspondences in English led us to abandon the notion of mapping rules in favor of weighted connections between units; the analogous inconsistencies in terms of syllables and morphemes might require abandoning parsing rules for the same reason. At the same time, the orthography does provide cues to syllabic and morphological structures. Morphemes, for example, are sublexical components that recur in a large number of words. As such they tend to be very high frequency spelling patterns. Consider, for example, a prefix such as PRE-, which recurs at the beginning of a large number of words. Empirical studies have suggested that the prefix and stem of a word act as perceptual groups (Taft, 1985). Does this grouping occur because the reader decomposes the word into morphemic components or because prefixes tend to be extremely high-frequency spelling patterns? Similar considerations hold in the case of syllables. The syllabic structures of words will tend to be realized in the orthography by inhomogeneities in the distributions of letters because syllables are properties of the spoken language and the orthography is alphabetic. Hence, "syllabic" effects could occur in word recognition not because readers recover syllabic structures per se, but only because they are affected by orthographic properties that are correlated with syllables. In sum, the hypothesis is that effects of units such as syllables and morphemes in visual word recognition are secondary to facts about how these units are realized in the writing system. Thus, effects of these structures would be an emergent property of a model, like ours, that only encodes facts about orthographic redundancy and orthographic–phonological regularity. We are currently examining this hypothesis (see Seidenberg, 1987, 1989, for discussion). There is already some suggestive evidence in this regard. Treiman and Chafetz (1987) have shown that subjects are sensitive to the division of syllables into onset and rime. In the word SPLASH, for example, the onset is SPL- and the rime is -ASH. We have already shown that rime units tend to be salient to pronunciation because of the structure of English orthography, as in the simulations of effects of different words on performance on TINT. Training with PINT or MINT has large effects on processing TINT, but training with TENT or TINS has much smaller effects. This is simply a consequence of the fact that vowel pronunciations—the most sensitive and least predictable aspect of the word—are sensitive to the letters that follow them, and the model picks up on this fact.

The scope of the model is also limited in that we have not implemented a process that takes computed phonological output into a set of articulatory–motor commands. We cannot be certain, then, that this process can be implemented in a manner consistent with facts about speech production. We think it highly unlikely that the model will prove to be inconsistent with facts about speech production, given the simple monotonic re-

lationship between phonological error scores and pronunciation latencies, but it does represent an unresolved issue. Minimally what is required is a mechanism that would take the imperfect specification of the phonological code provided by the model into an explicit representation of the pronunciation. The sequential networks described by Jordan (1986) are quite suggestive in this regard; these networks take patterns of activation representing entire words as input and learn to produce the corresponding phonemes one at a time in sequence. Ultimately, we would hope that a model of this type would encompass many of the phenomena described by Dell (1986) in a mechanism that incorporated learning procedures.[13]

Finally, the model does not address issues related to meaning. Insofar as the primary goal of word recognition is to identify the contextually appropriate meaning of a word, this represents a serious limitation. What we have demonstrated is that a large number of lexical decision and naming phenomena thought to bear on issues concerning access of meaning can be simulated by a model in which meaning is not represented at all. However, questions concerning the representation and access of meaning remain to be addressed; we have not, for example, even touched on the role of semantic priming or contextual constraint in word processing. As we have noted, promising work by Kawamoto (1988), Hinton and Sejnowski (1986), McClelland and Rumelhart (1985), and others have used principles very similar to the ones we have used to address the computation of meaning. Further exploration of these issues is an important topic for future research.

Details of the implementation. We have argued that aspects of our model are critical to understanding how words are recognized and pronounced. The critical aspects include the use of distributed representations, the existence of a layer of hidden units, the adjustment of weights on connections through learning, and the idea that pronunciation involves a direct mapping from orthography to phonology. There are details of the present implementation that are less theoretically relevant, however, and it is prudent to consider how they might contribute to its behavior. The main questions in this regard concern the representations of orthographic and phonological knowledge. The method of encoding phonemes was also used by Rumelhart and McClelland (1986a) in their model of the acquisition of past tense morphology. Pinker and Prince (1988) have noted several limitations of this encoding scheme.

We are aware of these limitations and have not claimed that the model embodies an adequate characterization of English phonology. The important question is, Does the model exhibit the behavior that it does (in terms of regularity effects and the like) *because* of specifics of the phonological (or orthographic) encoding schemes that we have chosen to use? This question can be addressed empirically by developing models that perform the same task as ours (learning about the structure of English orthography) but do not use the same representational schemes. Two additional models (Lacouture, 1989; Sejnowski & Rosenberg, 1986) provide evidence on this score. Sejnowski and Rosenberg's model uses letters and phonemes as representational units, rather than the triples used in our model. Although context sensitivity is not built into their representations, it is introduced in another way: Each letter is presented to the network for processing centered in a seven-letter window, so

that there are three letters of context on either side of the central letter. The task of the network is to produce the correct output for the central letter, given this context. In other respects, their model is similar to ours; it learns the correspondences between graphemes and phonemes using a network with a layer of hidden units and the back-propagation learning algorithm to adjust the weights on connections. Because the two models yield similar behavior in many respects, it appears that the use of the triples notation is not necessary in order to obtain many aspects of our own model's performance.

Lacouture's (1989) model, in contrast, uses a position-specific representational scheme similar to the one proposed by McClelland and Rumelhart (1981), rather than a locally context sensitive scheme like the one used here. That is, there was a complete set of 28 graphemic primitives (featural components of letters) for each of the letter positions in a word, counting from left to right. In spite of several obvious drawbacks of this sort of scheme, Lacouture's model also behaves similarly to ours, yielding, for example, the frequency by regularity interaction and other phenomena. Once again it appears that models with widely differing representation schemes yield qualitatively similar results. What is common to all of these models is the use of representations in which similar words with similar spelling produce overlapping input patterns, and words with similar pronunciations produce overlapping output patterns.

Of course, the specific details of the representations do affect the degree of overlap of input and output representations, and ultimately it will turn out that there are some choices of representation that will be superior to others, particularly if multisyllabic items are included. However, we do not think that the choice of representation is an a priori process independent of learning. Although there may be constraints that come originally from evolution or prereading experience, or both, we believe these predispositions are subject to considerable reorganization with experience. Our choice of representation was intended to approximate the one that people learn to use, rather than to serve as an exact characterization.

One other aspect of the implementation of the model deserves to be re-examined in light of our results: the fact that we compressed the range of word frequencies rather drastically in training our network. Two questions arise concerning this compression: Was it justifiable and was it responsible for any important aspects of the results?

We have already argued that some compression was justifiable, in that the untransformed Kucera-Francis (1967) word frequencies provide a biased picture of the experience we might expect a child to have with the words in our corpus. This is particularly true when we consider the fact that the spelling patterns and spelling-sound correspondences represented in low-frequency words tend to show up in words derived from the base forms of these words, as well as in the base forms themselves.

---

[13] Lacouture's (1989) model is suggestive in this respect. It computes phonological output in a manner very similar to ours; however, the computed phonological representation is then input to an auto-associative network (Anderson, Silverstein, Ritz, & Jones, 1977), which essentially completes the phonological code based on this partial input. This pattern-completion process might be seen as analogous to assembling an articulatory code.

Nevertheless, we cannot definitively assert that the actual degree of compression that we used is completely justified. This issue is important because Bever (in press) has suggested that the model closely simulates human performance only because of the frequency transform, which he considers to be unrealistic. Bever's conjecture is that the model would fail to learn the correct pronunciations of many words if a broader range of frequencies were used. As we have noted, exception words tend to be overrepresented among the higher frequency items in the lexicon. Bever's intuition is that if words such as HAVE or SAID were presented more often, the model would not be able to learn the regular pronunciations of regular inconsistent words such as RAVE or PAID.

Although this conjecture certainly deserves careful consideration, there is no reason to suppose that it is correct. Because of the error-correcting character of the learning rule that we use in training the network, performance on high-frequency items reaches asymptote relatively early; after this point, they exert relatively little influence on performance because the network has sufficient resources (in the form of units and connections) to master less frequent items in its environment. Under these circumstances, repeated presentation of high-frequency items keeps accuracy with these items high and, at the same time, allows gradual acquisition of the capacity to deal with other items in the corpus. We can see this pattern clearly in the simulations reported in this article. As Figure 3 shows, performance on words of relatively high frequency approaches asymptote by about 70 epochs, leaving room for continued improvement on lower frequency words. To be sure, a change in the frequency compression function that we used would tend to increase the importance of the word frequency factor, relative to the orthographic regularity, but it should not change the fact that both frequency and regularity influence performance nor the fact that regularity is a more important factor among less frequent words.

Still, in light of these considerations, it seemed prudent to explore whether similar results would obtain if a less drastic compression of the frequency range were used. Hence, we repeated the simulation using the same corpus of words and training procedure with one change: Words were sampled during the training phase as a function of the square root of their Kucera-Francis (1967) frequencies. Results of this simulation for the words in the Taraban and McClelland (1987) set are presented in Figure 27. The simulation was run for many more epochs because only about 60 items were presented in each one. The results replicate the Frequency × Regularity interaction seen in Figure 3. Looking at the regular inconsistent words, the correct pronunciations of these words again yielded much smaller error scores than the "exceptional" pronunciations, contrary to Bever's (in press) conjecture. Increasing the relative frequency of the higher frequency words did have one effect: It eliminated the regularity effect for high-frequency words early in training. In effect, the simulation says that if children were drilled repeatedly on a small number of high-frequency words, they would quickly learn to perform about equally well on both regular and irregular items.

In sum, the model is clearly limited in some respects, and details of its performance depend on some of the specific assumptions incorporated in the model. However, we see no rea-


Figure 27. Results of the replication simulation using a square root transform of frequency.

son to think that the theoretical conclusions we have offered are contingent on these aspects of the model.

Conclusions

We have presented a model of visual word recognition that synthesizes a broad range of empirical phenomena and provides an account of the types of knowledge relevant to this task, the manner in which this knowledge is represented in memory, and the course of acquisition. Our basic claim is that the model can account for these phenomena because of the close fit between the nature of the task (learning the structure of English orthography) and the capabilities of models of this type. English orthography is not strictly regular, and so it is not well captured by mechanisms involving systems of rules. Attempts to patch up this problem by proposing two routes (rules and lexical lookup) have been offered by others, but they have not been entirely successful. Our model, and others like it, offers an alternative that dispenses with this two-route view in favor of a single system that also seems to do a better job of accounting for the behavioral data. It remains for future research to establish whether the present approach can be successfully extended to longer words and to other aspects of word reading, and to integrate the word reading process, here artificially isolated, back into the process of understanding texts.

References

Adams, M. (1981). What good is orthographic redundancy? In H. Singer & O. J. L. Tzeng (Eds.), *Perception of print* (pp. 197–221). Hillsdale, NJ: Erlbaum.

Anderson, J. A., Silverstein, J. W., Ritz, S. A., & Jones, R. S. (1977). Distinctive features, categorical perception, and probability learning: Some applications of neural model. *Psychological Review, 89,* 369–406.

Andrews, S. (1982). Phonological recoding: Is the regularity effect consistent? *Memory & Cognition, 10,* 565–575.

Andrews, S. (1988). *Effects of frequency and neighborhood size on naming and lexical decision*. Manuscript submitted for publication.

Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. Spence & J. Spence (Eds.), *The psychology of learning and motivation* (Vol. 2, pp. 90–195). New York: Academic Press.

Backman, J., Bruck, M., Hébert, M., & Seidenberg, M. (1984). Acquisition and use of spelling-sound information in reading. *Journal of Experimental Child Psychology, 38,* 114–133.

Baddeley, A. D. (1979). Working memory and reading. In P. A. Kolers, M. E. Wrolstad, & H. Bouma (Eds.), *The processing of visible language* (pp. 355–370). New York: Plenum.

Balota, D., & Chumbley, J. (1984). Are lexical decisions a good measure of lexical access? The role of frequency in the neglected decision stage. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 340–357.

Balota, D., & Chumbley, J. (1985). The locus of the word frequency effect in the pronunciation task: Lexical access and/or production? *Journal of Memory and Language, 24,* 89–106.

Baron, J. (1973). Phonemic stage not necessary for reading. *Quarterly Journal of Experimental Psychology, 25,* 241–246.

Baron, J., & Strawson, C. (1976). Use of orthographic and word-specific knowledge in reading words aloud. *Journal of Experimental Psychology: Human Perception and Performance, 4,* 207–214.

Barron, R. W. (1981). Reading skill and reading strategies. In A. M. Lesgold & C. A. Perfetti (Eds.), *Interactive processes in reading* (pp. 299–327). Hillsdale, NJ: Erlbaum.

Barron, R. W. (1986). Word recognition in early reading: A review of the direct and indirect access hypotheses. *Cognition, 24,* 93–119.

Barron, R. W., & Baron, J. (1977). How children get meaning from printed words. *Child Development, 48,* 587–594.

Barsalou, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition, 10,* 82–93.

Bauer, D. W., & Stanovich, K. E. (1980). Lexical access and the spelling-to-sound regularity effect. *Memory & Cognition, 8,* 424–432.

Benton, A. L. (1975). Development dyslexia: Neurological aspects. In W. J. Friedlander (Ed.), *Advances in neurology* (Vol. 7, pp. 1–47). New York: Raven Press.

Besner, D., & Davelaar, E. (1982). Basic processes in reading: Two phonological codes. *Canadian Journal of Psychology, 36,* 701–711.

Besner, D., Waller, T. G., & MacKinnon, G. E. (1985). *Reading research: Advances in theory and practice* (Vol. 5). New York: Academic Press.

Bever, T. G. (in press). The demons and the beast—Modular and nodular kinds of knowledge. In C. Georgopouls, & R. Ishihara (Eds.), *Interdisciplinary approaches to language: Essays in honor of S.-Y. Kuroda*. Dordrecht, Holland: Kluwer.

Bradshaw, J. L., & Nettleton, N. C. (1974). Articulatory interference and the mown-down heterophone effect. *Journal of Experimental Psychology, 102,* 88–94.

Brown, G. D. A. (1987). Resolving inconsistency: A computational model of word naming. *Journal of Memory and Language, 26,* 1–23.

Burgess, C., Tanenhaus, M. K., & Seidenberg, M. S. (1989). Context and lexical access: Implications of nonword interference for lexical ambiguity resolution. *Journal of Experimental Psychology: Learning, Memory and Cognition, 15,* 620–632.

Bybee, J. L. (1985). *Morphology: A study of the relation between meaning and form.* Philadelphia, PA: John Benjamins.

Carr, T. H., & Pollatsek, A. (1985). Recognizing printed words: A look at current models. In D. Besner, T. G. Waller, & G. E. MacKinnon (Eds.), *Reading research: Advances in theory and practice* (Vol. 5, pp. 2–82). New York: Academic Press.

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

Cohen, J., Dunbar, K., & McClelland, J. L. (1989). *On the control of automatic processes: A parallel distributed processing model of the Stroop task.* Manuscript submitted for publication.

Coltheart, M. (1978). Lexical access in simple reading tasks. In G. Underwood (Ed.), *Strategies of information processing* (pp. 112–174). New York: Academic Press.

Coltheart, M. (1986). Functional architecture of the language-processing system. In M. Coltheart, G. Sartori, & R. Job (Eds.), *Cognitive neuropsychology of language* (pp. 1–25). London: Erlbaum.

Coltheart, M. (Ed.). (1987). *Attention & performance XII: Reading.* London: Erlbaum.

Coltheart, M., Besner, D., Jonasson, J., & Davelaar, E. (1979). Phonological recoding in the lexical decision task. *Quarterly Journal of Experimental Psychology, 31,* 489–508.

Coltheart, M., Davelaar, E., Jonasson, J., & Besner, D. (1977). Access to the internal lexicon. In S. Dornic (Ed.), *Attention & performance IV* (pp. 535–555). Hillsdale, NJ: Erlbaum.

Conrad, R. (1964). Acoustic confusions in immediate memory. *British Journal of Psychology, 55,* 75–84.

Daneman, M., & Carpenter, P. (1980). Individual differences in working memory and reading. *Journal of Verbal Learning and Verbal Behavior, 19,* 450–466.

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

Dennis, I., Besner, D., & Davelaar, E. (1985). Phonology in visual word recognition: Their is more two this than meats the I. In D. Besner, T. G. Waller, & G. E. MacKinnon (Eds.), *Reading research: advances in theory and practice* (Vol. 5, pp. 170–197). New York: Academic Press.

Donnenwerth-Nolan, S., Tanenhaus, M. K., & Seidenberg, M. S. (1981). Multiple code activation in word recognition: Evidence from rhyme monitoring. *Journal of Experimental Psychology: Human Learning and Memory, 7,* 170–180.

Fischler, I., & Bloom, P. (1979). Automatic and attentional processes in the effects of sentence contexts on word recognition. *Journal of Verbal Learning and Verbal Behavior, 18,* 1–20.

Forster, K. I. (1976). Accessing the mental lexicon. In E. C. T. Walker & R. J. Wales (Eds.), *New approaches to language mechanisms* (pp. 257–287). Amsterdam: North Holland.

Forster, K. I. (1981a). Frequency blocking and lexical access: One mental lexicon or two? *Journal of Verbal Learning and Verbal Behavior, 20,* 190–203.

Forster, K. I. (1981b). Priming and the effects of sentence and lexical contexts on naming time: Evidence for autonomous lexical processing. *Quarterly Journal of Experimental Psychology, 33A,* 465–495.

Forster, K. I., & Chambers, S. (1973). Lexical access and naming time. *Journal of Verbal Learning and Verbal Behavior, 12,* 627–635.

Fowler, C. A. (1977). *Timing control in speech production.* Bloomington: Indiana University Linguistics Club.

Frederiksen, J. R., & Kroll, J. F. (1976). Spelling and sound: Approaches to the internal lexicon. *Journal of Experimental Psychology: Human Perception and Performance, 2,* 361–379.

Glanzer, M., & Ehrenreich, S. L. (1979). Structure and search of the internal lexicon. *Journal of Verbal Learning and Verbal Behavior, 18,* 381–398.

Glushko, R. J. (1979). The organization and activation of orthographic knowledge in reading aloud. *Journal of Experimental Psychology: Human Perception and Performance, 5,* 674–691.

Gordon, B. (1983). Lexical access and lexical decision: Mechanisms of frequency sensitivity. *Journal of Verbal Learning and Verbal Behavior, 22,* 24–44.

Gough, P. B. (1972). One second of reading. In J. P. Kavanagh & I. G. Mattingly (Eds.), *Language by eye and by ear* (pp. 331–358). Cambridge, MA: MIT Press.

Henderson, L. (1982). *Orthography and word recognition in reading.* London: Academic Press.

Hillinger, M. L. (1980). Priming effects with phonemically similar words: The encoding-bias hypothesis reconsidered. *Memory & Cognition, 8,* 115–123.

Hinton, G. E. (1984). Parallel computations for controlling an arm. *Journal of Motor Behavior, 16,* 171–194.

Hinton, G. E. (1986). Learning distributed representations of concepts. *Proceedings of the Eighth Annual Conference of the Cognitive Science Society* (pp. 1–12). Hillsdale, NJ: Erlbaum.

Hinton, G. E., McClelland, J. L., & Rumelhart, D. E. (1986). Distributed representations. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1, pp. 77–109). Cambridge, MA: MIT Press.

Hinton, G. E., & Sejnowski, T. J. (1986). Learning and relearning in Boltzman machines. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1, pp. 282–317). Cambridge, MA: MIT Press.

Hoard, J. (1971). Aspiration, tenseness, and syllabification in English. *Language, 47,* 133–140.

Hooper, J. (1977). Word frequency in lexical diffusion and the source of morphological change. In W. Christie (Ed.), *Current progress in historical linguistics* (pp. 95–105). Amsterdam: Elsevier.

Humphreys, G. W., & Evett, L. J. (1985). Are there independent lexical and nonlexical routes in word processing? An evaluation of the dual-route model of reading. *Behavioral and Brain Sciences, 8,* 689–740.

James, C. T. (1975). The role of semantic information in lexical decisions. *Journal of Experimental Psychology: Human Perception and Performance, 104,* 130–136.

Jordan, M. I. (1986). Attractor dynamics and parallelism in a connectionist sequential machine. *Proceedings of the Eighth Annual Meeting of the Cognitive Science Society* (pp. 10–17). Hillsdale, NJ: Erlbaum.

Jorm, A. F., & Share, D. L. (1983). Phonological recoding and reading acquisition. *Applied Psycholinguistics, 4,* 103–147.

Kahn, D. (1976). *Syllable-based generalizations in English phonology.* Unpublished doctoral dissertation, University of Massachusetts.

Katz, L., & Feldman, L. B. (1981). Linguistic coding in word recognition: Comparisons between a deep and a shallow orthography. In A. M. Lesgold & C. A. Perfetti (Eds.), *Interactive processes in reading* (pp. 85–106). Hillsdale, NJ: Erlbaum.

Kawamoto, A. (1988). Distributed representations of ambiguous words and their resolution in a connectionist network. In S. Small, G. Cottrell, & M. Tanenhaus (Eds.), *Lexical ambiguity resolution* (pp. 195–228). San Mateo, CA: Morgan Kaufmann.

Kay, J., & Bishop, D. (1987). Anatomical differences between nose, palm, and foot, or the body in question. In M. Coltheart (Ed.), *Attention & performance XII: Reading* (pp. 449–469). London: Erlbaum.

Kleiman, G. M. (1975). Speech recoding in reading. *Journal of Verbal Learning and Verbal Behavior, 14,* 323–339.

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

LaBerge, D., & Samuels, S. J. (1974). Toward a theory of automatic information processing in reading. *Cognitive Psychology, 6,* 293–323.

Lacouture, Y. (1989). From mean square error to reaction time: A connectionist model of word recognition. In D. Touretsky, G. Hinton, & T. Sejnowski (Eds.), *Proceedings of the 1988 connectionist models summer school.* San Mateo, CA: Morgan Kauffman.

Landauer, T., & Streeter, L. (1973). Structural differences between common and rare words: Failure of equivalence assumptions for theories of word recognition. *Journal of Verbal Learning and Verbal Behavior, 12,* 119–131.

Locke, J. L. (1972). Ease of articulation. *Journal of Speech and Hearing Research, 15,* 194–200.

Marshall, J. C., & Newcombe, F. (1973). Patterns of paralexia: A psy-

cholinguistic approach. *Journal of Psycholinguistic Research, 2,* 175–199.

Marslen-Wilson, W. D. (1975). Sentence perception as an interactive parallel process. *Science, 189,* 226–228.

McCann, R. S., & Besner, D. (1987). Reading pseudohomophones: Implications for models of pronunciation assembly and the locus of word-frequency effects in naming. *Journal of Experimental Psychology: Human Perception and Performance, 13,* 14–24.

McCann, R. S., Besner, D., & Davelaar, E. (1988). Word recognition and identification: Do word-frequency effects reflect lexical access? *Journal of Experimental Psychology: Human Perception and Performance, 14,* 693–706.

McClelland, J. L. (1979). On the time relations of mental processes: An examination of systems of processing in cascade. *Psychological Review, 86,* 287–330.

McClelland, J. L. (1987). The case for interactionism in language processing. In M. Coltheart (Ed.), *Attention & performance XII: Reading* (pp. 3–36). London: Erlbaum.

McClelland, J. L., & Kawamoto, A. (1986). Mechanisms of sentence processing: Assigning roles to constituents. In J. L. McClelland & D. E. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 2, pp. 272–325). Cambridge, MA: MIT Press.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review, 88,* 375–407.

McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General, 114,* 159–188.

McClelland, J. L., & Rumelhart, D. E. (Eds.). (1986b). *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 2). Cambridge, MA: MIT Press.

McCusker, L. X., Hillinger, M. L., & Bias, R. G. (1981). Phonological recoding and reading. *Psychological Bulletin, 89,* 217–245.

McCutchen, D., & Perfetti, C. A. (1982). The visual tongue-twister effect: Phonological activation in silent reading. *Journal of Verbal Learning and Verbal Behavior, 21,* 672–687.

McRae, K., Jared, D., & Seidenberg, M. S. (in press). Frequency, lexical access, and naming. *Journal of Memory and Language.*

Meyer, D. E., Schvaneveldt, R. W., & Ruddy, M. G. (1974). Functions of graphemic and phonemic codes in visual word recognition. *Memory & Cognition, 2,* 309–321.

Minsky, M., & Papert, S. (1969). *Perceptrons.* Cambridge, MA: MIT Press.

Monsell, S. (1987). On the relation between lexical input and output pathways for speech. In D. A. Allport, D. G. MacKay, W. Prinz, & E. Scheerer (Eds.), *Language perception and production: Relationships among listening, speaking, reading, and writing* (pp. 273–311). London: Academic Press.

Monsell, S., Doyle, M. C., & Haggard, P. N. (1989). Effects of frequency on visual word recognition tasks: Where are they? *Journal of Experimental Psychology: General, 118,* 43–71.

Morton, J. (1969). Interaction of information in word recognition. *Psychological Review, 76,* 165–178.

Murrell, G. A., & Morton, J. (1974). Word recognition and morphemic structure. *Journal of Experimental Psychology, 102,* 963–968.

Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General, 106,* 226–254.

Onifer, W., & Swinney, D. (1981). Accessing lexical ambiguities during sentence comprehension: Effects of frequency, meaning, and contextual bias. *Memory & Cognition, 9,* 225–236.

Parkin, A. J. (1982). Phonological recoding in lexical decision: Effects

of spelling-to-sound regularity depend on how regularity is defined. *Memory & Cognition, 10,* 43-53.

Parkin, A. J., & Underwood, G. (1983). Orthographic vs. phonological irregularity in lexical decision. *Memory & Cognition, 11,* 351-355.

Patterson, K. E., & Coltheart, V. (1987). Phonological processes in reading: A tutorial review. In M. Coltheart (Ed.), *Attention & performance XII: Reading* (pp. 421-447). London: Erlbaum.

Patterson, K. E., Coltheart, M., & Marshall, J. (1985). *Surface dyslexia.* London: Erlbaum.

Patterson, K. E., Seidenberg, M. S., & McClelland, J. L. (in press). Connections and disconnections: Acquired dyslexia in a computational model of reading processes. In R. G. M. Morris (Ed.), *Parallel distributed processing: Implications for psychology and neuroscience.* London: Oxford University Press.

Perfetti, C. A. (1985). *Reading ability.* New York: Oxford University Press.

Pinker, S., & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition, 28,* 73-194.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review, 85,* 59-108.

Rayner, K., & Pollatsek, A. (1987). Eye movements in reading: A tutorial review. In M. Coltheart (Ed.), *Attention & performance XII: Reading* (pp. 327-362). London: Erlbaum.

Rozin, P., & Gleitman, L. (1977). The structure and acquisition of reading II: The reading process and the acquisition of the alphabetic principle. In A. Reber & D. Scarborough (Eds.), *Toward a psychology of reading* (pp. 55-141). Hillsdale, NJ: Erlbaum.

Rubenstein, H., Lewis, S. S., & Rubenstein, M. A. (1971). Evidence for phonemic recoding in visual word recognition. *Journal of Verbal Learning and Verbal Behavior, 10,* 647-657.

Rumelhart, D. E. (1977). Toward an interactive model of reading. In S. Dornic (Ed.), *Attention and performance VI* (pp. 573-603). Hillsdale, NJ: Erlbaum.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1, pp. 318-362). Cambridge, MA: MIT Press.

Rumelhart, D. E., & McClelland, J. L. (1986a). On learning the past tenses of English verbs. In J. L. McClelland & D. E. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 2, pp. 216-271). Cambridge, MA: MIT Press.

Rumelhart, D. E., & McClelland, J. L., Eds. (1986b). *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1). Cambridge, MA: MIT Press.

Schwanenflugel, P. J., & Shoben, E. J. (1985). The influence of sentence constraint on the scope of facilitation for upcoming words. *Journal of Memory and Language, 24,* 232-252.

Seidenberg, M. S. (1985a). Constraining models of word recognition. *Cognition, 20,* 169-190.

Seidenberg, M. S. (1985b). The time course of information activation and utilization in visual word recognition. In D. Besner, T. Waller, & G. E. MacKinnon (Eds.), *Reading research: Advances in theory and practice* (Vol. 5, pp. 200-252). New York: Academic Press.

Seidenberg, M. S. (1985c). The time course of phonological code activation in two writing systems. *Cognition, 19,* 1-30.

Seidenberg, M. S. (1987). Sublexical structures in visual word recognition: Access units or orthographic redundancy? In M. Coltheart (Ed.), *Attention and performance XII: Reading* (pp. 245-263). Hillsdale, NJ: Erlbaum.

Seidenberg, M. S. (1988). Cognitive neuropsychology and language: The state of the art. *Cognitive Neuropsychology, 5,* 403-426.

Seidenberg, M. S. (1989). Reading complex words. In G. N. Carlson &

M. K. Tanenhaus (Eds.), *Linguistic structure in language processing* (pp. 53-105). Amsterdam: Reidel.

Seidenberg, M. S. (in press-a). Lexical access: Another theoretical soup-stone? In D. Balota, G. Flores D'arcais, & K. Rayner (Eds.), *Comprehension processes in reading.* Hillsdale, NJ: Erlbaum.

Seidenberg, M. S. (in press-b). Visual word recognition and pronunciation: A computational model and its implications. In W. D. Marslen-Wilson (Ed.), *Lexical representation and process.* Cambridge, MA: MIT Press.

Seidenberg, M. S., Bruck, M., Fornarolo, G., & Backman, J. (1986). Word recognition skills of poor and disabled readers: Do they necessarily differ? *Applied Psycholinguistics, 6,* 161-180.

Seidenberg, M. S., McRae, K., & Jared, D. (1988, November). *Frequency and consistency of spelling-sound correspondences in naming.* Paper presented at the 29th annual meeting of the Psychonomic Society, Chicago.

Seidenberg, M. S., & Tanenhaus, M. K. (1979). Orthographic effects on rhyme monitoring. *Journal of Experimental Psychology: Human Learning and Memory, 5,* 546-554.

Seidenberg, M. S., Tanenhaus, M. K., Leiman, J. L., & Bienkowski, M. (1982). Automatic access of the meanings of ambiguous words in context: Some limitations of knowledge-based processing. *Cognitive Psychology, 14,* 489-537.

Seidenberg, M. S., Waters, G. S., Barnes, M. A., & Tanenhaus, M. K. (1984). When does irregular spelling or pronunciation influence word recognition? *Journal of Verbal Learning and Verbal Behavior, 23,* 383-404.

Seidenberg, M. S., Waters, G. S., Sanders, M., & Langer, P. (1984). Pre- and postlexical loci of contextual effects on word recognition. *Memory and Cognition, 12,* 315-328.

Sejnowski, T. J., & Rosenberg, C. R. (1986). *NETtalk: A parallel network that learns to read aloud* (EE and CS Techn. Rep. No. JHU/EECS-86/01). Baltimore, MD: Johns Hopkins University.

Selkirk, L. (1980). *On prosodic structure and its relation to syntactic structure.* Bloomington: Indiana University Linguistics Club.

Shallice, T., & Warrington, E. (1980). Single and multiple component central dyslexic syndromes. In M. Coltheart, K. Patterson, & J. C. Marshall (Eds.), *Deep dyslexia* (pp. 119-145). London: Routledge & Kegan Paul.

Shulman, H., Hornak, R., & Sanders, E. (1978). The effects of graphemic, phonemic, and semantic relationships on access to lexical structures. *Memory & Cognition, 6,* 115-123.

Spoehr, K. T., & Smith, E. E. (1973). The role of syllables in perceptual processing. *Cognitive Psychology, 5,* 71-89.

Stanhope, N., & Parkin, A. J. (1987). Further exploration of the consistency effect in word and nonword pronunciation. *Memory & Cognition, 15,* 169-179.

Stanovich, K. E. (1986). Matthew effects in reading: Some consequences of individual differences in the acquisition of literacy. *Reading Research Quarterly, 21,* 360-407.

Stanovich, K. E., & West, R. F. (1981). The effect of sentence context on ongoing word recognition: Tests of a two-process theory. *Journal of Experimental Psychology: Human Perception and Performance, 7,* 658-672.

Swinney, D. A. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior, 18,* 645-659.

Tabossi, P. (1988). Effects of context on the immediate interpretation of unambiguous nouns. *Journal of Experimental Psychology: Learning, Memory and Cognition, 14,* 153-162.

Taft, M. (1985). The decoding of words in lexical access: A review of the morphographic approach. In D. Besner, E. MacKinnon, & T. G. Waller (Eds.), *Reading research: Advances in theory and practice* (Vol. 5, pp. 83-123). New York: Academic Press.

Tanenhaus, M. K., & Carlson, G. (1989). Lexical structure and language comprehension. In W. D. Marslen-Wilson (Ed.), *Lexical representation and process*. Cambridge: MIT Press.

Tanenhaus, M. K., Dell, G., & Carlson, G. (1987). Context effects in lexical processing: A connectionist approach to modularity. In J. Garfield (Ed.), *Modularity in knowledge representation and natural language processing* (pp. 83-108). Cambridge, MA: MIT Press.

Tanenhaus, M. K., Flanigan, H., & Seidenberg, M. S. (1980). Orthographic and phonological code activation in auditory and visual word recognition. *Memory & Cognition, 8,* 513-520.

Taraban, R., & McClelland, J. L. (1987). Conspiracy effects in word recognition. *Journal of Memory and Language, 26,* 608-631.

Theios, J., & Muise, J. G. (1977). The word identification process in reading. In N. J. Castillan, D. Pisoni, & G. R. Potts (Eds.), *Cognitive theory* (Vol. 2, pp. 289-327). Hillsdale, NJ: Erlbaum.

Treiman, R., & Chafetz, J. (1987). Are there onset- and rime-like units in printed words? In M. Coltheart (Ed.), *Attention & performance XII: Reading* (pp. 281-298). London: Erlbaum.

Turvey, M. T., Feldman, L. B., & Lukatela, G. (1984). The Serbo-Croatian orthography constrains the reader to a phonologically analytic strategy. In L. Henderson (Ed.), *Orthographies and reading* (pp. 81-89). London: Erlbaum.

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review, 90,* 293-315.

Tweedy, J. R., Lapinski, R. H., & Schvaneveldt, R. W. (1977). Semantic context effects on word recognition: Influence of varying the proportion of items presented in an appropriate context. *Memory & Cognition, 5,* 84-89.

Van Orden, G. C. (1987). A ROWS is a ROSE: Spelling, sound, and reading. *Memory & Cognition, 15,* 181-198.

Van Orden, G. C., Johnston, J. C., & Hale, B. L. (1988). Word identification in reading proceeds from spelling to sound to meaning. *Journal of Experimental Psychology: Learning, Memory and Cognition, 14,* 371-385.

Vellutino, F. (1979). *Dyslexia.* Cambridge, MA: MIT Press.

Venezky, R. (1970). *The structure of English orthography.* The Hague, The Netherlands: Mouton.

Wang, W. S.-Y. (1979). Language change: A lexical perspective. *Annual Review of Anthropology, 8,* 353-371.

Waters, G. S., & Seidenberg, M. S. (1985). Spelling-sound effects in reading: Time course and decision criteria. *Memory & Cognition, 13,* 557-572.

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

Appendix

Stimuli in the Seidenberg, McRae, and Jared (1988) Experiment

Inconsistent words		Consistent words	Inconsistent words		Consistent words
Word	Enemy	Word	Word	Enemy	Word
bead	head	barge	loot	foot	lame
booth	smooth	bean	lull	bull	lilt
braid	said	beep	mall	shall	lure
brood	good	bin	noose	choose	mince
broth	both	bliss	pear	fear	nerve
brow	flow	brute	pleat	sweat	peach
cave	have	bunch	plied	skied	peel
chase	phase	cane	poll	doll	pier
crouch	touch	cape	pose	lose	poise
crush	bush	cheer	rut	put	probe
dome	some	coy	sneak	break	rust
drown	flown	den	sour	tour	scrub
floe	shoe	dime	stew	sew	steal
frost	post	doom	stint	pint	stole
gloss	gross	fade	stool	wool	strait
growl	bowl	flask	tease	cease	stunt
haste	caste	gloat	toad	broad	taint
hive	give	groan	tough	cough	teen
leaf	deaf	haunt	valve	halve	vain
lone	gone	hike	wove	love	weld

Note. Some inconsistent words have more than one enemy.