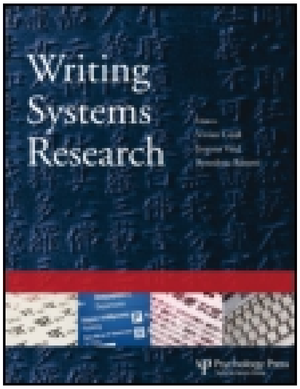

Writing Systems Research

Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/pwsr20

Statistical learning of conditional orthographic correspondences

Brett Kessler [a]

[a] Washington University in St. Louis, MO, USA

Version of record first published: 20 Dec 2011.

To cite this article: Brett Kessler (2009): Statistical learning of conditional orthographic correspondences, Writing Systems Research, 1:1, 19-34

To link to this article: http://dx.doi.org/10.1093/wsr/wsp004

Statistical learning of conditional orthographic correspondences

Brett Kessler
Washington University in St. Louis, MO, USA

Abstract

The English writing system deviates widely from the alphabetic ideal of uniform one-to-one correspondence between graphemes and phonemes, but its inconsistency is greatly reduced when conditional sound–spelling rules are applied. When reading or writing one part of a word, children and adults evinced knowledge of rules sensitive to the identity of other letters or phonemes, even those appearing much later in the word. Adults also showed sensitivity to the distinction between the basic and Romance subsystems of English (Albrow's Systems 1 and 2). Children as young as 6 years applied conditional rules that they were not taught, indicating implicit statistical learning of patterns observed in text. But learning is imperfect, and even adults did not match the frequency with which the patterns are found in English words.

Correspondence:
Brett Kessler,
Psychology Department,
Washington University
in St. Louis,
Campus Box 1125,
One Brookings Drive,
St. Louis, MO 63130-4899,
USA.
E-mail:
bkessler@wustl.edu

The English writing system is routinely criticized for falling far short of the alphabetic ideal (Caravolas, 2004) of one-to-one correspondence between phonemes and graphemes. Almost all its phonemes have multiple spellings—sound-to-letter, or spelling inconsistency—and almost all its letters have multiple pronunciations—letter-to-sound, or reading inconsistency. A study by Hanna *et al.* (1966) showed that only 73% of all phonemes would be spelt correctly if the writer picked the most common spelling for each phoneme: on average, a person who spells by the alphabetic principle would make a mistake on every fourth grapheme.

But arguably the speller's loss is the psycholinguist's gain. The fact that most mature spellers do much better than 73% raises questions as to how people read and write words, if applying one-to-one correspondences is not enough. Some theories that have been adduced are:

- Whole-word memorization: people memorize entire words as arbitrary sequences of letters. Such a theory accounts for the fact that people learn all sorts of inconsistent spellings, including one-offers like 'hiccough', where <gh>=/p/, and 'of', where <f>=/v/.[1] It also provides a straightforward explanation of how children can learn to read and write in classrooms whose curriculum is based on whole-word learning with no explanation of how sound–spelling correspondences work. But there is also much evidence that whole-word memorization cannot be the whole picture. All other factors being equal, people are slower and more error-prone at processing words with unusual sound–spelling correspondences (Balota *et al.*, 2004), which should not be a factor if they memorized all words as arbitrary letter sequences. Even more convincing is the fact that people can come up with plausible spellings when asked to write non-words, and other people can read such spellings. These feats would hardly be possible if the only route to reading and writing were memorizing words as unanalysable wholes.
- Learning large-unit sound–spelling correspondences: it has often been noted that the rimes (the vowel plus the following, coda, consonants) of riming words are often spelt alike. There may be many ways to spell /aɪ/ in English, but the

doi:10.1093/wsr/wsp004

spelling of the rime /aɪld/ is consistently <ild>, at least in simple monomorphemic words like ‘wild’ and ‘child’. The idea that people might learn the spelling of whole rimes as units is attractive because there is much evidence that preliterate children are good at breaking syllables down into onsets (the part before the vowel) and rimes—while breaking rimes down into vowels and codas is often very challenging for them (Goswami, 1993; Treiman and Kessler, 1995). Much work in literacy research has proceeded from the point of view that onsets and rimes are the effective units of reading and writing (Stanback, 1992; Ziegler *et al.*, 1997), which is not to claim that these researchers subscribe to the theory that literacy is nothing more that memorizing the spellings of whole onsets and rimes. One problem with such a theory is that many rimes have multiple spellings (‘bed’, ‘said’, ‘head’). Furthermore, Aronoff and Koch (1996) reported that there are only 12 rime spellings in English that are worth memorizing as wholes.

- Learning conditional sound–spelling correspondences: although the alphabetic ideal is typically described as having only unconditional rules like ‘the letter <n> spells the phoneme /n/’, full stop, researchers such as Carney (1994), Cummings (1988), and Venezky (1970, 1999) have shown that many correspondences have fewer exceptions when stated as rules that specify conditions—rules like ‘<n> spells /ŋ/ before <k>’. Conditional spelling rules can, in principle, refer not just to phonetic environments but also other word properties such as morpheme class, word length, or etymological stratum: ‘<ph> spells /f/ in words of Greek extraction’. Indeed, even one-off irregularities can be formulated as conditional rules: ‘<gh> spells /p/ in the word “hiccough” ’. Thus, conditional rules of orthography are very powerful but also potentially complex.

In this article, I report some recent psycholinguistic work in which my colleagues and I have explored whether people learn and use conditional sound–spelling rules, and if so, how. I do not mean that these are necessarily rules in the everyday sense of the term, as when one talks about the rules of chess: a set of linguistically encoded instructions that one consciously learns from a teacher and can easily formulate and pass on to one’s student. Even spelling researchers may have a hard time explaining that /k/ is spelt <ck> at the end of word-final stressed syllables when the rime of the uninflected lexeme would otherwise contain fewer than three letters (e.g. ‘sick’ versus ‘silk’, as discussed below), and most people who follow the rule have never been taught it overtly (Hayes *et al.*, 2006). An alternative hypothesis is that complex sound–spelling correspondences are learnt the same way we learn many other patterns in life: by observing and internalizing the relative frequency with which objects and events occur and co-occur. Perhaps through repeated exposure to text, people gradually pick up both unconditional and conditional correspondences, without necessarily formulating any conscious accounts of them. This type of learning can be called statistical learning, in the sense that it is based on implicit numeric analysis: orthographic patterns are essentially observations (or computations) that particular sound–spelling correspondences are more frequent than others in a particular context. It is important to keep in mind that most vocabulary about mental processes—words such as ‘rules’, ‘observation’, ‘computation’—carries an implication of conscious, explicit thought, but in the context of this article I do not intend any such implication.

With this proviso in mind, I address in this article such questions as whether people notice and use spelling patterns they have never been taught, and if so, how. Are people deterministic computers whose knowledge of letter frequencies perfectly reflects their experience? At what age do people begin to build up this knowledge? Do they use it for both reading and for writing? And what types of patterns can people pick up on?

1 Vocabulary Statistics

A first step in exploring this statistical learning hypothesis was to see what sort of usable spelling patterns could be found in English words. Kessler and Treiman (2001) were primarily interested in learning how the predictability of sound–spelling correspondences in one part of the syllable is improved when rules are sensitive to the contents of other parts of the syllable. Because the syllabification of intervocalic consonants is a moot issue in English, we found it expedient to restrict this first study to monosyllabic words. We gathered 3,117 words, basing our selection primarily on familiarity ratings obtained from American college students (Nusbaum *et al.*, 1984). Pronunciations were taken from the Random House Dictionary (Flexner, 1987). When multiple pronunciations were listed, we selected the first general American pronunciation, but homographic heterophones such as ‘read’ (/ɹid/, /ɹɛd/) were treated as separate words. After splitting the pronunciation into onset, vowel, and coda, we aligned those three parts with their spelling. For example, the word ‘crane’ aligned like this:<cr>=/kɹ/, <a_e>=/e/, <n>=/n/. ‘Sign’ aligned <s>=/s/, <ig>=/aɪ/, <n>=/n/, because we adopted Albrow’s (1972) criterion that all letters must align with some phoneme.

Table 1 Sound–spelling consistencies of the parts of monosyllabic words

Context[a]	**Onset**	**Vowel**	**Coda**
Reading[b]			
Unconditional	0.976	0.717	0.982
Given onset	–	0.807	0.992
Given vowel	0.993*	–	0.992
Given coda	0.988	0.920*	–
Writing[c]			
Unconditional	0.910	0.529	0.821
Given onset	–	0.649*	0.882
Given vowel	0.937*	–	0.925*
Given coda	0.942	0.737*	–

[a]‘Unconditional’ is unconditional consistency; the others are conditional, with the context noted after the word ‘Given’.

[b]Letter-to-sound consistency.

[c]Sound-to-letter consistency.

*Consistency gain over unconditional value is significantly greater than when the values in the context are randomized, $P < 0.001$.

With all these alignments in place, we were able to compute consistency measures for the various orthographic units in the three parts of the syllable. For example, in the onset of the words, the letter <c> has a consistency of 88.4%, because of the ninety-seven words that begin with <c>, ninety-one (93.8%) are pronounced with /k/ and the remaining six (6.2%) are pronounced with /s/; the weighted average of those two proportions is 88.4% = (91 × 93.8% + 6 × 6.2%) ÷ 97. (All counts are by word types, not weighted by their frequency in running text.) Further, by averaging such values for all letters in syllable onsets, we concluded that the average consistency of letter-to-sound mappings in onsets was 97.6%. The rows labelled ‘Unconditional’ in Table 1 report these unconditional consistency measures, as proportions ranging from 0 to 1, for all the three parts of the syllables. They quantify the impression that vowels are more difficult than consonants, and that writing is harder than reading.

Measurements of conditional correspondences were more directly relevant for our purposes. For example, we computed the consistency of initial <c> when the vowel is <a>, its consistency when the vowel is <i>, and so forth for all onsets and all vowels, then took the weighted average. Table 1 shows in the cell labelled ‘Onset Reading Given vowel’ that the letter-to-sound consistency of the onset rises to an almost perfect 99.3% when the vowel is taken into account. Similarly, all the other conditional consistencies are higher than the unconditional consistencies. It is particularly noteworthy how much more reliably one can read vowels when the coda is taken into account, and how much more reliably one can spell all parts of a syllable when the identity of the adjacent syllable position is taken into account. The mere fact that there are increases is not at all surprising: extra information can never hurt, even when the associated values are purely coincidental. In fact, in many cases it turned out that the conditional consistencies were comparably high even if we scrambled the contexts, arbitrarily swapping, for example, the coda of each word with the coda of another. But for some contexts, the ones denoted by asterisk in Table 1, the attested conditional consistencies were significantly higher than when the contexts were randomized. These contexts tend to be more systematic, typically raising the consistencies because of patterns that are due to conditional sound changes in the history of the language. For example, <i> is particularly likely to be pronounced /aɪ/ when the coda is <ld> because of a lengthening that occurred in Old English, and English has had many other vowel changes that were conditioned by specific codas. In contrast, very few sound changes in the onset were conditioned by the coda, or vice versa, which accounts for the non-significant connections between those syllable parts.

It bears remembering that these conditional consistencies were all derived by mathematically analysing the vocabulary itself: the data do not say that people take advantage of these patterns, but they do prove that the patterns are there (a full list of the individual patterns can be found at http://spell.psychology.wustl.edu/RelSoundLetMono). The data also give us some basis for addressing the large-unit hypothesis raised earlier: would it make sense to memorize the spellings of entire onsets and rimes? The fact that the consistencies of vowels are so much higher when the coda is taken into account, and vice versa, means that such a strategy would be much better than only learning unconditional phoneme-level correspondences. On the other hand, the data also show that taking the vowel into account can help one to read and spell the onset, and taking the onset into account can help one to spell the vowel. Just memorizing large units like onsets and rimes in an unconditional manner would ignore much information that could improve reading and spelling accuracy.

2 Spelling Vowels in Non-words Using Conditional Correspondences

The next question is whether humans are sensitive to, learn, and use conditional patterns. Treiman *et al.* (2002) explored how people spell vowels, which, as Table 1 shows, have the most to gain from conditional reference to context. Do adults take adjacent consonants into account when they spell vowels? We selected nine vowels for which our vocabulary analysis had revealed that a given spelling was much more likely in one context than another. These are listed in Table 2. For instance, it turns out that stressed /i/ is spelt <ee> in 64% of all words where it appears before word-final /d/ or /p/, but that <ee> is used only 15% of the time, on average, before several other consonants, such as /ð/ and /m/. The difference between those two numbers is presented as a proportion (0.49) in the table. There are more words like ‘bleed’ than ‘mead’ and more words like ‘dream’ than like ‘seem’, but even the minority spellings are not particularly unusual. Other patterns that were tested were less violable. For instance, /aʊ/ is quite reliably spelt as <ow> before coda /l/ and /n/ (‘clown’), although even in those contexts there are exceptions, such as ‘noun’.

Table 2 Increase in proportion of vowels that adults spelt with the conditioned spelling when a conditioning consonant is present

Vowel	Conditioned spelling	Context	Example	Vocabulary[a]	Responses[b]
Before coda					
/ɛ/	<ea>	/d/	head	0.40	0.06
/i/	<ee>	/d/, /p/	feed	0.49	0.27
/o/	<o> without <e>	/l/	roll	0.33	0.11
/ʊ/	<u>	/l/, /s/, /ʃ/, /tʃ/	push	0.73	0.11
/aʊ/	<ow>	/l/, /n/	clown	0.89	0.43
/aɪ/	<igh>	/t/	night	0.39	0.27
After onset					
/ɑ/	<a>	/w/	wasp	0.75	0.72
/ɚ/	<or>	/w/	word	0.52	0.28
/u/	<oo>	Non-coronals	food	0.57	0.58

[a]Based on standard spellings at the end or beginning of English words. Proportion of words with indicated vowel and consonant that have the indicated spelling, minus proportion of words with the vowel and another consonant that have the indicated spelling.

[b]Results from experimental trials, measured the same way as the vocabulary statistics

Three of the patterns involved the influence of the onset. For example, we set up stimuli that tested whether people were sensitive to the pattern by which General American /ɑ/ (this would be /ɒ/ in some other accents) is spelt as <a> 86% of the time it appears after /w/ (‘wand’), but only 11% of the time after other consonants (not ‘*fand’ but ‘fond’; the difference 0.75 appears in the table). Thus, we investigated a wide variety of patterns.

To test whether spellers were sensitive to a particular pattern, we designed ten pairs of non-words like /θɛd/ and /θɛk/ and asked adults to spell them as if they were English. The pairs were split up and interleaved with similar questions for the other eight patterns, so that the participants would not deduce any patterns from the stimuli in the experiment itself. In this particular case, we tallied whether they used <ea> significantly more often in non-words like /θɛd/ than in those like /θɛk/. And that proved to be true. In this case, the percentage of non-words like /θɛd/ that were spelt with <ea> was 11%, whereas only 5% of the stimuli like /θɛk/ had <ea>; the difference of 6% is what is reported in the ‘Responses’ column in Table 2. This is a fairly small effect, but it was statistically significant (the cutoff $P<0.05$ will be used throughout this article) and, as the table shows, the other cases we tested had effects of larger magnitudes.

Thus, adult spellers have learnt conditional patterns, and they do make active use of them in their own productions. This knowledge is statistical in several senses of the word. The patterns are rarely if ever taught explicitly or recalled consciously, but appear to be the product of long experience seeing different vowel spellings being used in different contexts. They are not exceptionless rules but statements of frequency distributions. Indeed, a few of the rules are, at least by themselves, rather useless in spelling: it is never advisable to guess that the spelling of /ɛ/ is <ea>, because even before /d/ that spelling is less frequent than <e> (‘bed’). Nevertheless, <ea> is in the speller’s repertoire, and it has a higher probability of being selected before some codas than before others. The experiment demonstrated that statistical patterns of conditional probabilities do play an important role in phonological spelling.

The experiment also answered the question as to whether apparently conditional spellings like /ɛ/=<ea> before /d/ could be due to a different mechanism whereby spellers memorize spellings for whole rimes (/ɛd/=<ead>) and onsets. In this experiment, that possibility is contradicted by the fact that the same spellers who drew on the identity of the coda to spell vowels (intra-rime effects) also took into account the identity of the onset. For example, they spelt /ɑ/ as <a> the great majority of the time after /w/, but comparatively rarely after other consonants. Thus, a whole-rime spelling strategy appears to be not only an insufficient explanation, but also an unnecessary one: the same type of (implicit) awareness of conditional patterns that is needed to explain the influence of the onset can be used to explain the influence of the coda.

3 Conditional Spelling of Vowels in Real Words

Non-words can be ideal stimuli for orthographic experiments, because they do not carry the same sort of baggage that words do. Factors such as familiarity, frequency, age of learning, and semantic factors such as imageability can differentially affect how people read and spell words (Balota *et al*., 2004). There is always the concern that when we think people are reacting distinctively to a group of words because of one property, such as high inconsistency, they are really reacting to some other property, such as lower familiarity, which we were unable to control for or might not even have known about. In the previous experiment, if people were asked to spell real words, such as ‘dead’ versus ‘deck’, the finding that people used more <ea> before /d/ would be compromised by the obvious objection that the results could be explained by whole-word memorization. Because we used non-words such as /θɛd/ versus /θɛk/, which people had never seen spelt before, such an objection is untenable, and the implication that the participants used conditional patterns is much more believable.

Nevertheless, there is a sense in which studying the spelling of words via non-words is somewhat artificial. One might reasonably object that non-word spelling is an imperfect window on natural spelling, because people might write them using different strategies than when writing familiar words. Consequently, we configured another version of the experiment to use actual words (Treiman *et al*., 2002). We asked the participants to spell pairs like ‘shred’ and ‘fleck’. We then performed error analyses to see how often they misspelt the vowel as <ea>, testing the hypothesis that such an error would be more likely in the context (coda /d/) that is more strongly associated with the /ɛ/ = <ea> correspondence in English orthography. Other examples included looking for <ee> in ‘reap’ versus ‘ream’; <o> without silent <e> in ‘shoal’ versus ‘croak’; <u> in ‘swoosh’ versus ‘nook’; <ow> in ‘pronoun’ versus ‘slouch’; <igh> in ‘contrite’ versus ‘confide’; <a> in ‘wombat’ versus ‘possum’; <o> in ‘whir’ versus ‘blurt’; and <oo> in ‘scuba’ versus ‘frugal’. Because errors of any type would be more numerous on longer or less familiar words, the stimuli were matched for length and familiarity. An attempt was made to select words that college students would know, but not so well that they would spell them perfectly.

The errors went in the expected direction: 27% of the errors in the conditioning context constituted use of the conditioned vowel, but only 8% of the errors in the other contexts were of that nature. That is, we often saw errors like <reep> instead of <reap>, but much less often did we get errors like <reem> for <ream>. The only tested pattern that was not reflected in the participants’ misspellings was the use of <u> before /ʃ/, as in ‘push’, and that may be due to the fact that the only item tested, ‘swoosh’, would have unusual graphotactics, the sequence <wu>, if the contextual rule were applied to it.

Thus, people do use statistical conditional rules in a naturalistic spelling task, when trying to recall the spelling of moderately familiar words that they do not encounter every day. Granted, the way the experiment was framed, it sounds like context sensitivity is a bad thing, because we investigated how contexts like /p/ after /i/ could mislead spellers into making errors. One must keep in mind, however, that more errors on words like ‘reap’ mean fewer errors on words like ‘ream’. It is easy to discount /m/ after /i/ as being a conditioning context, because the spelling it conditions, <ea>, is the more common one, but from a statistical perspective /m/ is just as much a context as /p/ is. Another thing to keep in mind is that we intentionally selected the critical stimuli from the minority of words that go against the general pattern. Even if sensitivity to context leads some spellers astray on a few words like ‘wombat’, that same sensitivity should help them out on the much larger number of words in which /ɑ/ is spelt <a> after <w>, such as ‘waffle’ and ‘wallow’.

4 Reading Vowels Using Conditional Rules

Much more research has studied how people read than how they spell. It has often been reported that readers’ pronunciation of vowels varies depending on the coda (e.g. Johnson and Venezky, 1976; Glushko, 1979; Ryder and Pearson, 1980; Andrews and Scarratt, 1998). However, the prior literature left several issues open. As discussed above in the case of spelling, one interpretation of a coda effect is that people memorize letter-to-sound correspondences at the level of whole onsets and rimes: if a person reads <nind> as /naɪnd/, perhaps they have broken the word down into <n> and <ind>, and applied the correspondences <n> = /n/ and <ind> = /aɪnd/, treating the whole rime as a unit. A rime-unit strategy could be distinguished from a conditional strategy (<i> = /aɪ/ before <nd>) if vowel pronunciations were found to be influenced by onset consonants as well as by coda consonants. Several studies (e.g. Treiman and Zukowski, 1988; Treiman *et al*., 1995; Andrews and Scarratt, 1998) reported that the onset has negligible influence on vowel reading, thus leaving open the possibility that readers do follow a rime-unit strategy.

Treiman *et al.* (2003) addressed this issue with an experiment structurally similar to that of the spelling experiment reported earlier, but geared towards conditional patterns that are effective in the reading direction. Adults were asked to read, interleaved among other trials, pairs of non-words such as <spange> and <spance>. If they pronounced words of the former type with /e/ more than they did the latter, that would show sensitivity to the fact that the coda spelling <nge> conditions the /e/ pronunciation (words like ‘range’) whereas <nce> does not (words like ‘prance’). Two of the patterns tested included onset conditioning. Table 3 shows the patterns that were tested, along with their reliability in text, measured as in Table 2. Most of the patterns were quite strong, but not completely exceptionless (e.g. ‘flange’ is pronounced with /æ/).

The last column in Table 3 shows that the readers were sensitive to the patterns. The numbers are the difference between the proportion of pronunciations in the conditioning context (e.g. <nge>) that have the conditioned vowel (/e/) and the proportion of those in the non-conditioning context (<nce>) that have that vowel. Thus, any positive proportion in this column shows a trend towards sensitivity to the context; its maximum is 1.0. In all cases, the differences between the two contexts were statistically significant.

It is important to note that the two sets of items that involved an onset context had just as strong a response as the items that had only a coda context. This similarity is expected under the theory that people learn conditional patterns, but it would not be expected if people rely on rime-unit reading. In fairness to previous researchers who found no reliable effect of onset, there are few cases in English where the pronunciation of a vowel letter is strongly sensitive to the identity of the onset; indeed, the only truly important cases may be those involving /w/. Thus, it is true that the overall impact of the onset is small. But from a psychological point of view, that makes it all the more striking that people pick up on the few onset–vowel patterns that do exist. Because rimes are particularly salient for English speakers, and because the great majority of contextual effects on vowel pronunciation come from within the rime (i.e. the coda), one might expect that people would be locked in to look for patterns within rimes, or at least learn and apply them much more robustly. Instead, it appears that people learn statistical patterns wherever they exist.

5 Magnitudes and Computer Modelling

A curious fact about the experiments with non-word stimuli is that the participants’ rate of use of statistical patterns was lower than that found in the general vocabulary. For example, while 40% of English words ending in stressed /ɛd/ are spelt with <ea>, the participants in our experiment used <ea> for only 11% of the non-words that ended in /ɛd/ (Table 2). Across both experiments, participants approached the frequency attested in the lexicon in only three of the seventeen patterns we examined. If readers and writers were perfect learners of the statistics of the writing system, and if the process of spelling consisted entirely of retrieving the most probable relevant patterns, we would expect the attested frequencies to be closer to those that people observe in text (Brown, 1998). Although we have successfully predicted that people will learn and apply certain patterns, we have no good model for predicting the magnitude with which they will apply those patterns.

Table 3 Increase in proportion of vowels read with conditioned pronunciation when a conditioning consonant is present

Vowel	Conditioned pronunciation	Context	Example	Vocabulary	Responses
Before coda					
<a>	/e/	<nge>	range	0.95	0.54
<a>	/ɔ/	<ld>, <lt>	bald	0.79	0.86
<ea>	/ɛ/	<d>	dread	0.78	0.12
<i>	/aɪ/	<nd>, <ld>	find	0.76	0.33
<o>	/o/	<ld>, <lt>	bold	0.96	0.83
<oo>	/ʊ/	<k>	look	0.96	0.70
After onset					
<a>	/ɑ/	<u>, <w>	wand	0.93	0.58
<a>	/ɔ/	<u>, <w>, before <r>	warm	0.96	0.16

Numbers are differences between proportions, analogous to those in Table 2.

It is always a challenge to predict how strongly people will react to environmental stimuli, and I will not conclude this section with a bold new model that accounts for why the vocabulary numbers are so much greater than the participants'. However, it was instructive to look at computer models of reading to see how their numbers compared with those of the vocabulary and of our participants. We examined the output of twelve different simulators and looked for implementation characteristics that provided the closest match with the human behaviour (Treiman *et al*., 2003).

The most prominent of the current spelling models is DRC (the dual-route, cascaded, model of Coltheart *et al*., 2001). This model pronounces non-words by applying a series of letter-to-sound rules that are almost all unconditional. Not surprisingly, DRC scored 0.00 on all patterns: that is, it was no more likely to use a conditioned vowel pronunciation in a conditioning consonant context than when the vowel letter appeared in another context. Clearly, unconditional rules are not a good model of human behaviour for these stimuli.

Most other reading models pronounce non-words by applying connectionist principles. Broadly speaking, a connectionist model simulates a neural network: it has units that represent neurons and connections between units that represent synapses. Processing, such as reading, is a matter of setting the activation levels of a certain set of units to represent the input (letters); spreading the activation through the net along the connections between the units; then reading off the activation levels at the set of units that represent the output (phonemes). Learning is a matter of adjusting the strengths of connections until the processing of a training set—in our case, a word list containing spellings and pronunciations—is optimized (Rumelhart and McClelland, 1986). At no point do classical connectionist models formulate explicit rules or even store specific frequency information at a particular locus: information is distributed throughout the network of units.

In many respects, the design of connectionist models is redolent of the way people gradually and implicitly learn spelling patterns through long exposure to and use of the spelling system. We might, therefore, expect connectionist models to do a good job simulating human reading of vowels whose pronunciation is conditional, and this is partly true. Like our research participants, the connectionist models did learn the conditional patterns, and they did tend to use the conditioned pronunciations more in the strongly conditioning environments than in other environments—e.g. they pronounced <a> as /e/ more before <nge> than before <nce>. They also matched human performance quite closely on non-words that have environments that only weakly condition pronunciations, e.g. stimuli that end in <nce>, with agreement levels as high as 94–96% for some models (those of Plaut *et al*., 1996; Zorzi *et al*., 1998; Powell *et al*., 2001; Harm and Seidenberg, 2004). These are all notable achievements. However, the connectionist models did not match the magnitudes of human performance very closely on the strongly conditioning environments like <a> before <nge>. The best performance on such non-words was by the model of Norris (1994), which matched human performance only 68% of the time; some other models performed as badly as 43% on such stimuli (Zorzi *et al*., 1998).

A closer error analysis reveals a few reasons for the models' performance. The model of Zorzi *et al*. (1998) was intentionally designed to severely restrict the complexity of patterns that the system could learn, in that it lacked hidden units and its connections were set up so that patterns could only be learnt between vowels and codas, not between vowels and onsets. At the other extreme, the model of Norris (1994) specialized in contextual effects: vowels were set up to be read in a way that explicitly took into account how they were pronounced in words that shared the same onset or coda. This strategy worked moderately though not spectacularly well for non-words that have a strongly conditioning onset or coda, but the strategy apparently only got in the way of reading non-words whose vowel has a typical unconditional pronunciation. Most of the remaining models applied the conditioned pronunciations in conditioning environments (e.g. using /e/ in <a>+<nge> words) more than the humans did, often at rates approaching or even exceeding the patterns' frequency in actual vocabulary.

A couple of models (Plaut and McClelland, 1993; Powell *et al*., 2001) were trained not only on English vocabulary, but also on individual letter-to-sound rules like <a>=/æ/, because many children are heavily exposed to these unconditional rules in school, and so may tend to over-apply them, therefore under-applying conditional rules. Although that is an ingenious idea, it did not have a noticeable effect in bringing down the programs' rate of application to human levels.

The models also differed in many other details, any of which could account for substantial differences among their outputs. Therefore, it is perhaps best to think of their outputs not as rigorous tests of any particular theory of statistical learning, but as sources of ideas for future research. Norris's (1984) way of incorporating context seems promising. Instead of learning only the highly discriminating contextual effects that we tested on—and which most connectionist models would learn most thoroughly—his model also explicitly uses information about non-discriminating contexts. For example, in reading a non-word like <chead>, we have only been considering the fact that when <d> is present, the probability of the correspondence <ea>=/ɛ/ jumps to 78% (it normally spells /i/ in other contexts). We, like most connectionist models, have been ignoring the onset <ch>, because it has no effect on how the following <ea> is pronounced; in 'cheat', for example, <ea> has the same pronunciation it has in the vast majority of words, /i/. However, if humans do learn that <ea> is always pronounced /i/ after <ch>—a useless but true fact—and if they apply that pattern when reading non-words like 'chead', the conflict between the two patterns could, in part, account for why people do not reliably read such words with the conditioned pronunciation, /ɛ/, as often as we would expect: other factors besides the coda /d/ pull them towards other pronunciations. Such a scenario seems very likely and definitely worthy of further investigation.

6 Course of Acquisition and Use of Context in Reading

The previous studies all examined reading and writing behaviour in adults, or in computer simulations that had been trained to adult levels. But literacy researchers pay at least as much attention to how children read and write words. Children who have not yet mastered a writing system are the best sources of information on how literacy skills are acquired. There is also an important practical benefit to understanding the development of orthographic knowledge, in that educators can take into account what sort of information children are capable of dealing with at various ages, and in what order they naturally acquire different types of knowledge. See, for example, Cook (2004) for an introduction to the voluminous literature on theories of literacy acquisition. With respect to the question of conditional orthographic patterns, the central concerns here are how old children are when they acquire sensitivity to context; at what rate their application of conditional rules increases over the course of their life; whether they have conditional rules properly speaking, or whether they use rime-unit strategies; and whether some types of sensitivity to context are acquired earlier than others.

Treiman *et al.* (2006) replicated the above study of vowel reading, but instead of recruiting college-age adults, we studied children and young adults ranging in age from 6 to 17 years. Because children vary appreciably in their rate of development, we analysed the results as a function of the children's grade-equivalent reading ability as determined by a standardized test (Wilkinson, 1993). We looked at how often they produced conditioned vowel pronunciations when they read vowel letters in conditioning environments as opposed to other environments, e.g. how often <a> was read as /e/ before <nge> as opposed to <nce>.

All the groups that we studied had a positive score on this measure: e.g. /e/ was used significantly more often in the conditioning environment. This sensitivity to context was found even among the first-graders that we studied. However, when we grouped the first-graders by reading ability, we found that those reading at a kindergarten level did not show sensitivity to context; indeed, most of their pronunciations did not match how the letters are pronounced in any English word. Those reading at the first-grade level, however, did evince sensitivity to context.

As for the specific types of patterns, children reading at first-grade level were sensitive to only three of the patterns we investigated. They pronounced <a> as /ɑ/ more often after <w> and <qu> than after other onset letters; they pronounced <o> as /o/ more often before <ld> and <lt> than before other coda letters in the absence of final <e>; and they pronounced <oo> as /ʊ/ more often before <k> than before other coda letters. At higher reading levels, children were sensitive to all eight patterns. The only exception was that the high-school students in our sample showed an atypical reluctance to read <ea> as /ɛ/ before <d>.

As for the magnitude of the effect, first-grade readers already showed an appreciable difference of about fourteen percentage points between how often they used a conditioned vowel in conditioning versus non-conditioning contexts. As reading grade improved, so did this magnitude. This change was due both to increases in children's production of the conditioned vowel where appropriate (e.g. producing /ʊ/ for <oo> before <k>) and to decreases in producing that vowel where not appropriate (e.g. for /ʊ/ <oo> before <n> or <m>). With the exception of a small dip at the high-school level due to that rejection of <ea>=/ɛ/, both trajectories continued throughout our sample and into the college-age population described earlier. However, after about Grade 5, the improvements were small and not statistically significant.

Thus, children acquire and apply conditional orthographic patterns from the very beginning of their careers as readers. Presumably these rather obscure patterns are not taught overtly but are picked up through general statistical learning, although it is, of course, impossible to prove that children never have heard anyone else mention the rule. These results do more damage to the rime-unit theory of reading, because the children were at least as sensitive to patterns where the onset influences the vowel pronunciation (e.g. <a> as /ɑ/ after <w> or <qu>) as to those where the influence is the coda; indeed, the onset pattern appears to be one of the first ones acquired. Thus, children do not simply read rimes as whole units in an unconditional fashion.

If children learn and apply contextual patterns, why does the magnitude of the effect remain below adult levels for the first several grades of school? One reasonable hypothesis is that the vocabulary patterns children are exposed to changes as they progress through school. After all, the highly context-dependent patterns we have been investigating mostly constitute what are often considered irregularities in the orthographic system; perhaps textbooks intentionally hold off on irregular words until children get older. An examination of the words that appear in a graded corpus (Zeno *et al.*, 1995) showed that this is not the case: The proportion of such words in books remains remarkably constant from first grade on. Therefore, we cannot hypothesize that first-graders are less likely to read <ea> before <d> as /ɛ/ because they are not exposed to so-called irregular words like 'head' and 'dead'. Alternatively, perhaps children change their reading strategies as they get older or more mature, gradually becoming more and more convinced that context is a good thing. A simpler explanation is that children get better at patterns as they get more experience and practice with them. Given the complexity of English, it would not be surprising that the course of improvement might take five years or more. We tested this hypothesis by running an implementation of the connectionist reading model described by Harm and Seidenberg (2004). It was repeatedly exposed to 3,102 monosyllabic words, and its training was periodically interrupted to test how it performed on our reading task. Importantly, the vocabulary and reading strategy of the simulation never changed, but the trajectory of its output on the reading task as its exposure to vocabulary increased was very similar to the differences between children of different reading levels. It started off in early trials with a profile similar to that of our first-graders, got progressively and rapidly better at discriminating between the conditioning and non-conditioning environments, then began improving much more gradually after about 500,000 epochs of training (full details are at http://spell.psychology.wustl.edu/InflCContxtOnPronV). While it is true that no model perfectly emulates human behaviour, the simulation did show that a progression like that of our grade-school students can be accounted for entirely by gradual learning over time, using constant vocabulary and no change in learning styles or reading strategy.

7 Course of Acquisition and Use of Context in Writing

Treiman and Kessler (2006) looked at children's acquisition of conditional patterns from the other angle, that of spelling. We used a testing procedure very similar to the one that had been used for adults: testing how people would spell vowels for non-words that have different types of consonantal contexts, as per Table 2. However, because spelling large number of words can tax the patience of young children, we used a fill-in-the-blank procedure, asking them, for example, to complete the spelling of the word dictated as /glɛd/ by filling in what was missing on the answer form, which had <gl__d> already written. Aside from making the task easier, this procedure also made the results somewhat cleaner, because the participant had no opportunity to select unintended (and non-conditioning) consonants. One may object that filling in a blank is an unnatural way to write except in crossword puzzles and classroom exercises and so may draw on different skills than spelling a whole word, but this objection is largely countered by the observation that the oldest participants in this study, high-schoolers, behaved very similarly to the adults who had to spell entire words in the previous study. The most important drawback was that one pattern from the adult study had to be omitted: the test of how /o/ is spelt before /l/. The basic choices are between <ole> and <oll>, so providing the coda consonant or consonants in advance could have strongly biased the results either towards or away from using final <e>.

For analysing the results, the participants were grouped by spelling-grade levels (two or three grades in a group) as determined by a standard spelling test (Jastak and Wilkinson, 1993). In many ways the results were the same as for the child reading experiment. The kindergarten/first-grade group showed only a slight preference for using conditioned vowels in conditioning contexts, but the magnitude of the effect increased through grade school. There was no evidence for rime-unit spelling; in fact, children began observing a conditioning effect of an onset on vowels (that /ɚ/ is mostly <or> after /w/) at least three grade levels earlier than they observed the first effects of codas on vowels. Also as in the reading study, analysis of children's texts show that the relevant vocabulary statistics remain constant from first grade on, so differences in children's textual environment cannot explain why children's spelling changes as they progress through grade school.

In other respects, the course of spelling development was different from that of reading. Children were slower to use context, with only very small effects in the kindergarten/first-grade group and with acquisition not levelling off until the sixth-to-eighth-grade group, as compared with fifth grade in the reading study. Although the sound-to-letter patterns we chose for the spelling task were a little more difficult than the letter-to-sound patterns used in the reading task, as suggested by the Vocabulary statistics in Tables 2 and 3, it seems more likely that the nature of the two tasks is primarily responsible for the differences. If children learn patterns primarily by being exposed to them in print, then it is reasonable to assume that access to the patterns would be activated more readily when seeing patterns already printed, in the reading task. In writing, the stimuli are auditory, and young children not yet used to visualizing their spelling productions in advance may find it hard to recognize that they are violating a spelling pattern until they have already written a word, presumably with its most common, unconditional, vowel spelling.

The slow emergence of context sensitivity in spelling affords us the opportunity to observe the order in which different patterns emerge. Table 4 lists for all eight patterns the spelling-grade level at which use of the pattern is first reliably attested. A comparison with Table 2 shows that the age of first use does not correlate in any obvious way with how strongly the context conditions the vowel spelling (the Vocabulary statistic) nor with how often adults observe the context in their own spellings (Responses). However, Table 4 shows that the age of first use correlates with the frequency of the most common unconditional spelling for the vowel in question. Although one must be very cautious about *post hoc* speculation on just eight data points, an enticing theory is that young children are more likely to pay attention to context when there is no

Table 4 Conditional vowel spelling patterns arranged by frequency of modal spelling

Vowel	Modal correspondence[a]		Conditional correspondence	
	Spelling	Proportion[b]	Pattern	Grade attested
/ɚ/	<ir>	0.25	<or> after /w/	Kindergarten–1
/u/	<oo>	0.36	<oo> after non-coronals	2–3
/i/	<ea>	0.41	<ee> before /d/, /p/	4–5
/ʊ/	<oo>	0.47	<oo> before /k/	6–8
/ɑ/	<o>	0.51	<a> after /w/	4–5
/aʊ/	<ou>	0.63	<ow> before /l/, /n/	4–5
/aɪ/	<i>	0.66	<igh> before /t/	6–8
/ɛ/	<e>	0.74	<ea> before /d/	6–8

[a]Most common spelling of the vowel across all contexts in monosyllabic English words.

[b]Proportion of instances of the vowel that use the modal spelling.

candidate that clearly dominates across contexts. To put what must be a complex, implicit process into crude terms, it is as if they quickly learn that their best unconditional guess for spelling /ɚ/ will be correct only 25% of the time, so they had better pay attention to conditional rules. In contrast, guessing <e> for /ɛ/ works 74% of the time, so the payoff for learning to use conditional rules is not very great.

8 Consonant Spelling

This article has focused up till now on vowels, primarily because of the disproportionately high number of important contexts that condition their sound–spelling correspondences. But some English consonant spellings are conditioned by the vowel, and these have inspired some informative studies as well. Hayes *et al.* (2006) looked at how children and adults spell word-final consonants. As a first approximation, this rule states that <f>, <l>, <k>, and <ch> are extended to <ff>, <ll>, <ck>, and <tch>, respectively, lexeme-finally after a stressed vowel if the rime would otherwise be spelt with just two letters: <stiff>, <pill>, <back>, <patch>, not *<stif>, *<pil>, etc.; as opposed to <lift>, <life>, <loaf>, and so forth, where the extension is not needed. We asked participants from the age of 7 to 22 years to spell pairs of non-words like /səl/ versus /sul/, or to choose which of two spellings like <sul> versus <sull> or <sool> versus <sooll> looks better. In both of these tasks, children as young as second grade were sensitive to some of these patterns, generating or preferring spellings like <sull> and <sool>. The magnitude of the effect increased with each age group considered. Magnitudes of the effect were also much greater in the recognition task than in the writing task, across all age groups. Finally, the effect was much stronger for <(c)k> and <(l)l> than for <(f)f> and <(t)ch>, in both tasks and across all age groups; this point probably reflects the fact that the general rule has some very high-frequency exceptions in the latter two cases, namely, ‘if’, ‘of’, ‘much’, and ‘such’.

The higher robustness of the effect in this study as compared with the vowel spelling study may be due to the fact that here, the conditional rule is entirely graphic. That is, <ll> would simply be illegal following <soo>, regardless of what sound one were spelling. The fact that these graphic rules are applied at younger grade levels than the conditional sound-to-letter rules studied earlier suggests again that children find purely visual patterns easier to work with. The conclusion that there is a visual advantage is strongly reinforced by the fact that the same data show the effect much more strongly when presented in a purely visual task (‘which looks better, <sul> or <sull>?’) than when the participants had to generate spellings based on an auditory stimulus (‘spell /səl/’).

Hayes *et al.* (2006) also studied one conditional rule affecting the spelling of onsets: that /k/ cannot be spelt <c> but is normally spelt <k> before <e>, <i>, or <y>. We found that children as young as second grade were more likely to spell /k/ as <k> in non-words before those letters than before <a>, <o>, or <u>.

9 Consonant Reading

One last question concerns the reading of consonants. Treiman *et al.* (2007) looked at how conditional rules affect the way adults read onsets. We were interested in the fact that the pronunciations of word-initial <c> and <g> are not only conditioned by the following vowel, but also by the different subsystems of English orthography. In the terminology of Albrow (1972), in System 1, the basic system, word-initial prevocalic <g> is always /g/, and <c> is always /k/ but is used only before the historically back vowel letters, <a>, <o>, and <u>. In System 2, the Romance system, both consonants have two pronunciations, depending on whether they occur before the historically back vowels, where <g> is /g/ and <c> is /k/, as in the basic system, or before the historically front vowels, <e>, <i>, and <y>, where <g> is /dʒ/ and <c> is /s/. (For brevity, I will henceforth talk about front and back vowels and the front (/dʒ/ and /s/) and back (/g/ and /k/) pronunciations of the consonants.) Thus, the pronunciation of these two letters is conditioned in English by the interaction of two separate factors: the identity of the following vowel and the orthographic system to which the word or word component belongs. In a series of experiments, we asked college-age adults to read aloud a variety of words and non-words that begin with <c> or <g>, to see what conditions would influence them to select front or back pronunciations for those consonants.

In all the experiments, the participants showed sensitivity to the identity of the following vowel. Before back vowels, the consonants were almost always given a back pronunciation, in accordance with an almost exceptionless rule that applies in all systems. Before front vowels, both front and back pronunciations were given for the consonant. All factors being equal, this is not surprising for <g>, which does appear with both pronunciations before front vowels in many words ('get', 'gem'), but it is perhaps surprising that our participants pronounced <c> as /k/ a substantial number of times when it appeared in non-words before <e> or <i>, a pattern virtually unknown in English. For example, when asked to pronounce monosyllables like <cersh>, they produced /k/ 16% of the time, but for <garsh>, nobody used /dʒ/. This imbalance appears to be another instance of people's tendency to never quite let go of the most common unconditional pronunciation of a letter, even when a specific context does not support it. The same type of error has been reported in the Romance languages, even though the rules for front and back pronunciations of <c> and <g> apply in a very large number of words and are virtually exceptionless (Content and Peereman, 1992, for French; Job *et al.*, 1998, for Italian).

We also conducted several experiments where the vocabulary stratum of the word was indicated by graphic properties of the word's spelling. In one experiment, we constructed triplets such that a nonsense syllable was presented either by itself (<geb>), or with a Romance suffix (<gebic>), or with a basic suffix (<gebful>). The stimulus <gebic> is overtly Romance, and so the pronunciation /dʒ/ is expected, while <geb> and <gebful> have no explicit markers for any spelling system (because basic suffixes are common even on Romance bases), and so could be pronounced either way. In another experiment, we constructed pairs of non-words that were as equivalent as possible, except that one member of the pair, such as <gireak>, used spelling patterns that are typical in the basic system, and the other, such as <girec>, used typically Romance spelling patterns. Here, one would expect to hear more /g/ in the former case and more /dʒ/ in the latter case. We found that our participants did tend to use front pronunciations before front vowels, the Romance pattern, more often when the word had overt clues that it was Romance than when it did not. This difference was significant, though it was small. For example, <g> before <e> was pronounced /dʒ/ 15% of the time in stimuli that had a basic suffix, and 24% of the time in stimuli that had a Romance suffix.

Our experiments showed that adult readers have learnt a fairly abstract conditional rule. Whereas most rules discussed up to now have only required the reader or speller to consult the adjacent letter or sound, this rule requires the reader to evaluate a variety of clues to guess the appropriate spelling system. Some of these clues involve looking ahead quite a few letters: sometimes only the end of a word tells how the initial letter should be pronounced. Lookahead is a challenge for models of reading such as the DRC (Coltheart *et al*., 2001), which assume that processing is strictly sequential.

10 Conclusion

Although English orthography appears chaotic, there are many conditional rules that significantly increase its consistency. Such rules are somewhat complicated, but a large body of evidence shows that even beginners process English in a way that cannot be accounted for by unconditional strategies such as whole-word or rime-unit memorization. Children take into account the identity of adjacent units when spelling consonants or vowels, and adults, at least, take into account relatively long-distance contexts, looking at the spelling or morphology of the rest of the word in order to decide how to pronounce the first letter. At present, we are investigating to what extent children can do this as well.

Our main motivation in this work is to understand the processes by which people learn and process complex writing systems, but it is tempting to speculate briefly on paedagogical applications. The obscurity of many of the patterns we have investigated, combined with the testimony of teachers (Hayes *et al*., 2006), strongly suggests that children are not explicitly taught these patterns but pick them on their own through general processes of statistical learning: observing what letters, sounds, and sound–spelling correspondences are repeatedly found together in text. To a certain extent, therefore, these findings validate whole-language or text-based approaches to literacy instruction. On the other hand, we have also shown that several useful spelling patterns are not acquired until fairly late in grade school, and in virtually all cases, the patterns that are acquired are not applied nearly as often as they should be. This suggests that there may be a place for increased attention to conditional patterns in literacy instruction after all.

Funding

This work was supported by the (US) National Institutes of Health [R01HD051610-02].

References

Albrow, K. H. (1972). *The English Writing System: Notes Toward a Description*. London: Longman Publishing Group.

Andrews, S. and Scarratt, D. R. (1998). Rule and analogy mechanisms in reading nonwords: hough dou peapel rede gnew wirds? *Journal of Experimental Psychology: Human Perception and Performance*, **24**(4): 1052–86.

Aronoff, M. and Koch, E. (1996). Context-sensitive regularities in English vowel spelling. *Reading and Writing: An Interdisciplinary Journal*, **8**(3): 251–65.

Balota, D. A., Cortese, M. J., Sergent-Marshall, S. D., Spieler, D. H., and Yap, M. J. (2004). Visual word recognition of single-syllable words. *Journal of Experimental Psychology: General*, **133**(2): 283–316.

Brown, **G. D. A.** (1998). The endpoint of skilled word recognition: the ROAR model. In Metsala, J. L. and Ehri, L. C. (eds), *Word Recognition in Beginning Literacy*. Mahwah, NJ: Lawrence Erlbaum Associates, pp. 121–38.

Caravolas, M. (2004). Spelling development in alphabetic writing systems: a cross-linguistic perspective. *European Psychologist*, **9**(1): 3–14.

Carney, **E.** (1994). *A Survey of English Spelling*. London: Routledge.

Coltheart, M., Rastle, K., Perry, C., Langdon, R., and Ziegler, J. (2001). DRC: a dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*, **108**(1): 204–56.

Content, A. and Peereman, R. (1992). Single and multiple process models of print to speech conversion. In Alegria, J., Holender, D., Morais, J., and Radeau, M (eds), *Analytic Approaches to Human Cognition*. Amsterdam: Elsevier Science Publishers, pp. 213–36.

Cook, V. (2004). *The English Writing System.* London: Arnold.

Cummings, D. W. (1988). *American English Spelling: An Informal Description*. Baltimore, MD: Johns Hopkins University Press.

Flexner, S. B. (ed.) (1987). *Random House Dictionary*, 2nd edn. New York: Random House.

Glushko, R. J. (1979). The organization and synthesis of orthographic knowledge in reading aloud. *Journal of Experimental Psychology: Human Perception and Performance*, **5**: 674–91.

Goswami, U. (1993). Phonological skills and learning to read. In Tallal, P., Galaburda, A. M., Llinas, R. R., and von Euler, C. (eds), *Temporal Information Processing in the Nervous System: Special Reference to Dyslexia and Dysphasia*. New York: New York Academy of Sciences, pp. 296–311.

Hanna, P. R., Hanna, J. S., Hodges, R. E., and Rudorf, E. H. (1966). *Phoneme–grapheme Correspondences as Cues to Spelling Improvement.* Washington, DC: U.S. Department of Health, Education, and Welfare.

Harm, M. W. and Seidenberg, M. S. (2004). Computing the meanings of words in reading: cooperative division of labor between phonological and visual processes. *Psychological Review*, **111**(3): 662–720.

Hayes, H., Treiman, R., and Kessler, B. (2006). Children use vowels to help them spell consonants. *Journal of Experimental Child Psychology*, **94**(1): 27–42.

International Phonetic Association (1999). *Handbook of the International Phonetic Association: A Guide to the Use of the International Phonetic Alphabet.* Cambridge, England: Cambridge University Press.

Jastak, S. and Wilkinson, G. S. (1993). *Wide Range Achievement Test*, 3rd edn. Wilmington, DE: Jastak Associates.

Job, R., Peressotti, F., and Cusinato, A. (1998). Lexical effects in naming pseudowords in shallow orthographies: further empirical data. *Journal of Experimental Psychology: Human Perception & Performance*, **24**(2): 622–30.

Johnson, D. D. and Venezky, R. L. (1976). Models for predicting how adults pronounce vowel digraph spellings in unfamiliar words. *Visible Language*, **10**(3): 257–68.

Kessler, B. and Treiman, R. (2001). Relationships between sounds and letters in English monosyllables. *Journal of Memory and Language*, **44**(4): 592–617.

Norris, D. (1994). A quantitative multiple-levels model of reading aloud. *Journal of Experimental Psychology: Human Perception and Performance*, **20**, 1212–32.

Nusbaum, H. C., **Pisoni**, **D. B.**, **and Davis**, **C. K.** (1984). *Sizing Up the Hoosier Mental Lexicon: Measuring the Familiarity of 20,000 Words*. Bloomington: Indiana University Press.

Plaut, D. C. and McClelland, J. L. (1993). Generalization with componential attractors: word and nonword reading in an attractor network. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 824–9.

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

Powell, **D., Plaut, D. C., and Funnell, E.** (2001). *A Developmental Evaluation of the Plaut, McClelland, Seidenberg, & Patterson (1996) Connectionist Model of Single Word Reading* [Abstract]. Worcester, UK: Meeting of the British Psychological Society, Developmental and Education Sections.

Rumelhart, D. E. and McClelland, J. L. (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. Cambridge, MA: MIT Press.

Ryder, R. J. and Pearson, P. D. (1980). Influence of type-token frequencies and final consonants on adults' internalization of vowel digraphs. *Journal of Educational Psychology*, **72**(5): 618–24.

Stanback, M. L. (1992). Syllable and rime patterns for teaching reading: analysis of a frequency-based vocabulary of 17,602 words. *Annals of Dyslexia*, **42**: 196–221.

Treiman, R. and Kessler, B. (1995). In defense of an onset–rhyme syllable structure for English. *Language and Speech*, **38**(2): 127–42.

Treiman, R. and Kessler, B. (2006). Spelling as statistical learning: using consonantal context to spell vowels. *Journal of Educational Psychology*, **98**(3): 642–52.

Treiman, R., Kessler, B., and Bick, S. (2002). Context sensitivity in the spelling of English vowels. *Journal of Memory and Language*, **47**(3): 448–68.

Treiman, R., Kessler, B., and Bick, S. (2003). Influence of consonantal context on the pronunciation of vowels: a comparison of human readers and computational models. *Cognition*, **88**(1): 49–78.

Treiman, R., Kessler, B., and Evans, R. (2007). Anticipatory conditioning of spelling-to-sound translation. *Journal of Memory and Language*, **56**(2): 229–45.

Treiman, R., Kessler, B., Zevin, J. D., Bick, S., and Davis, M. (2006). Influence of consonantal context on the reading of vowels: evidence from children. *Journal of Experimental Child Psychology*, **93**(1): 1–24.

Treiman, R., Mullennix, J., Bijeljac-Babic, R., and Richmond-Welty, E. D. (1995). The special role of rimes in the description, use, and acquisition of English orthography. *Journal of Experimental Psychology: General*, **124**: 107–36.

Treiman, R. and Zukowki, A. (1988). Units in reading and spelling. *Journal of Memory and Language*, **27**: 466–77.

Venezky, R. L. (1970). *The Structure of English Orthography*. The Hague: Mouton.

Venezky, R. L. (1999). *The American Way of Spelling: The Structure and Origins of American English Orthography*. New York: Guilford.

Wilkinson, G. S. (1993). *Wide Range Achievement Test 3*. Wilmington, DE: Wide Range, Inc.

Zeno, S. M., Ivenz, S. H., Millard, R. T., and Duvvuri, R. (1995). *Educator's Word Frequency Guide*. Brewster, NY: Touchstone Applied Science Associates.

Ziegler, J. C., Stone, G. O., and Jacobs, A. M. (1997). What is the pronunciation for -ough and the spelling for /u/?: a database for computing feedforward and feedback consistency in English. *Behavior Research Methods, Instruments & Computers*, **29**: 600–18.

Zorzi, M., Houghton, G., and Butterworth, B. (1998). Two routes or one in reading aloud?: a connectionist dual process model. *Journal of Experimental Psychology: Human Perception and Performance*, **24**(4): 1131–61.

Notes

1 Graphemes of the orthography are enclosed in angled brackets. Phonemes are represented by symbols of the International Phonetic Association (1999), and are enclosed in slash marks. The Appendix gives keywords illustrating how I applied the symbols of the IPA to American English.

Appendix

Table A1 Keywords for use of International Phonetic Alphabet symbols for American English

IPA	Keyword
aɪ	**ice**
aʊ	**ou**t
æ	**a**dd
ɑ	**o**dd
d	**d**ig
dʒ	**j**ump
ð	**th**is
e	**ace**
ɛ	**e**bb
ə	**u**p
ɚ	**ear**n
f	**f**ine
ɡ	**g**ood
i	**ea**t
k	**k**iss
l	**l**og
m	**m**ouse
n	**n**ice
ŋ	si**ng**
o	**oa**k
ɔ	**ough**t
p	**p**ie
ɹ	**r**ed
s	**s**un
ʃ	**sh**ip
t	**t**oy
tʃ	**ch**eese
u	**oo**ze
ʊ	b**oo**k
v	**v**oice
w	**w**ind
θ	**th**ink