

4

The Question of Phonology and Reading

Guy C. Van Orden and Heidi Kloos

Picture a 6-year-old puzzling out the printed word *island*. For this child, the printed form of the word is not entirely familiar and requires effortful decoding. First the child says /i . . . i/, then /iz/, then /land/, then /iz land/, and all of a sudden she gets it right: she correctly reads *island* aloud. English is notorious for words like *island*, with spellings that only partly reveal how they are spoken. Nonetheless, the way a word is spelled almost always reveals something of how to say it aloud, something of its phonology. This fact illustrates the alphabetic principle, which has long lead scientists to wonder about the role of phonology in reading. For instance, does a printed word's phonology play a role in deriving meaning from the word?

A central question of reading is whether understanding individual written words always depends upon prior access to their phonology: Do mental representations of phonology mediate the comprehension of written words? This question is based on the idea that reading can be taken apart as the links of a chain of mental events. Think again of the 6-year-old child and the word *island*. Reading *island* aloud, the effect, seems to depend on an intermediate representation of *island's* phonology, its cause. Presented with the visual stimulus *island*, the child's mind first forms a representation of the word's spelling. The spelling representation will be decoded, or puzzled out, to create a representation of phonology. In turn, the mental representation of *island's* phonology activates the motor program of its pronunciation, and the child says *island* aloud.

The child portrayed as reading *island* aloud illustrates the classic view of stimulus → mediating representations → response: the idea of a causal chain between the stimulus word *island* and its pronunciation. It is from this vantage point that the question of phonology and reading is posed. Both the printed word *island* and *island's* pronunciation appear in some sense to be outside the child's head, but she constructs phonology inside her head. Phonology is intermediate, both in the sense of a middle position in time or space and in the sense of a mental causal link between the printed word and its pronunciation (cf. Fodor, 1981; Kihlstrom, 1987; Markman & Dietrich, 2000).

Debate about whether comprehension of individual words requires intermediate phonology began in the nineteenth century and continues today. But decisive evidence, one way or the other, has eluded reading scientists. This chapter uses examples of the conflicting evidence to understand why the evidence has failed to settle the debate. The first half of this chapter describes homophone studies and priming studies. Both kinds of studies discover phonology effects that are contingent on task demands. They illustrate how task contingent evidence, instead of settling the debate, simply fuels it. The second half of the chapter describes evidence of a complex interaction between printed and spoken language. Complex interactions could explain why previous evidence was so sensitive to task demands.

How Evidence Fuels the Controversy

Do representations of phonology mediate comprehension in skilled reading? An influential theory – the dual-process theory – defined itself in an answer to this question. Proposed by Marshall and Newcombe (1973) and refined by Coltheart (1978), this important theory set the stage for a contemporary science of reading (see Van Orden, Pennington, & Stone, 1990, 2001, for reviews).

Dual-process theory

As the name suggests, dual-process theory includes two processes by which words can be identified in reading. One process involves rules for how to map elements of spelling onto elements of phonology. This process parses a word's string of letters into its elementary units of spelling, called graphemes. Then grapheme–phoneme rules are applied to map graphemes onto elementary units of phonology, called phonemes. Finally, the phonemes are combined to form a representation of the word's phonology, which can pick the word out in a phonological lexicon.

The other process is called direct access. Direct access takes a visual representation of the word as input and assigns it to an abstract placeholder in the mental lexicon. Identification of a word happens in one step going from a visual representation to an entry in the mental lexicon. It is called direct access because it creates a shortcut that bypasses the grapheme–phoneme rules. To link each word's visual representation to a lexical entry requires word-by-word associations. The links develop as a reader becomes familiar with words.

The two processes, grapheme–phoneme rules and direct access, should distinguish between skilled and unskilled readers. An unskilled reader should identify words by applying grapheme–phoneme rules, like a child who has just learned to read. Skilled readers should bypass grapheme–phoneme rules as direct access becomes available. Skilled readers should only use grapheme–phoneme rules when they confront an unfamiliar word such as *pharisee* (Doctor & Coltheart, 1980).

In this dual-process view, learning to read depends crucially on learning the alphabetic principle, of which grapheme–phoneme rules is one hypothetical approximation. Skilled

reading, in contrast, is predicted to occur without mediating phonology. Evidence from homophone errors and priming studies does not corroborate this clear-cut prediction, but neither is the prediction ruled out, as explained next.

Homophone errors

In a semantic categorization task, a homophone target such as *break* is sometimes miscategorized as *part of a car* (Van Orden, 1987). This error stems from the fact that *break* shares identical phonology with *brake*. At face value, such homophone errors clearly demonstrate mediating phonology; words are confused because they share the same phonology.

Dual-process theory predicted that skilled readers would read familiar words via direct access. Yet skilled readers make homophone errors to homophone words irrespective of their familiarity. Homophone errors are no less likely when frequently read homophones like *break* appear as targets than when targets are relatively unfamiliar homophones, like *peek* for the category *part of a mountain*. Insensitivity to word familiarity would appear to falsify the direct-access hypothesis of skilled reading, but the story is not that simple.

If a categorization task includes more broadly specified categories such as *object*, then familiar homophone targets such as *break* produce no more errors than control items (Jared & Seidenberg, 1991; Van Orden, Holden, Podgornik, & Aitchison, 1999). Familiarity now matters. Semantic categorization to broadly specified categories produces homophone errors to low frequency homophones like *peek* but not to high frequency homophones like *break*. This finding is inconsistent with the previous results, but consistent with the direct-access hypothesis. Direct-access should be available for words that are frequently read and direct access would preclude homophone errors.

On the basis of homophone errors' absence, one could argue that the readers are not using phonology (Jared & Seidenberg, 1991). That is, the null effect of homophone phonology, when familiar homophones are judged against broadly specified categories, could imply that no phonology link exists in this case. If so, then the link between phonology and comprehension of printed words is at best partial and certainly not obligatory for skilled readers.

This logic may seem too arbitrary or simplistic. Too many reasonable alternatives present themselves for how a change in task demands may eliminate a phonology effect but not eliminate phonology (Bosman & de Groot, 1996; Lesch & Pollatsek, 1993). For example, the effects of phonology as a cause of homophone word comprehension may be concealed in contexts where performance rises to ceiling, as it usually does to highly familiar words (Lukatela & Turvey, 1994a; Van Orden et al., 1999). Perhaps overly familiar homophone words are coded too efficiently to reveal a phonology effect under the conditions of the broadly specified categories (cf. Unsworth & Pexman, 2003). Or perhaps phonology interacts in complex ways with task demands and other sources of information and the question is altogether too simply framed (Van Orden, Holden, & Turvey, 2003).

Clearly, the evidence provided by homophone errors leaves the causal status of phonology as a mediator between print and meaning undecided. As a consequence, homophone

errors do not answer the question of phonology and reading to the satisfaction of all reading scientists. Special circumstances of task demands are required to produce homophone errors to familiar homophones. But a phonology effect based on special circumstance is not persuasive; it will not dissuade scientists who trust the direct-access hypothesis. In the same vein, special circumstances of task demands are required to make homophone errors go missing, and the consequent null phonology effect will not dissuade scientists who trust that reading includes mediating phonology.

Priming studies

The direct-access hypothesis has a constant traveling companion: the assumption that mediating phonology is delayed with respect to direct access (as Frost, 1998, points out; however, see Paap, Noel, & Johansen, 1992). According to this assumption, word identification via an assembly process of grapheme–phoneme rules takes more time than the direct visual associations of direct lexical access. Thus, for example, skilled readers may not base their response in a lexical decision task on phonology representations because direct access recognizes a familiar word, as a word, prior to assembly of phonology. Perhaps studies that address the delayed-phonology hypothesis may decide the status of phonology in reading.

How soon after seeing a printed word does phonology become available? One way to answer this question is with a combination of backward masking and priming. Masking concerns the length of time that items are visible. Priming concerns how one letter string may affect another, how a prime such as *REEZ* may affect identification of a target such as *rose*, for instance. The target word *rose* appears for a fraction of a second before it is replaced by the prime *REEZ*. *REEZ* serves as a mask of *rose* because it limits the amount of time available to derive *rose* phonology, and it serves as a prime because it shares partial phonology with *rose*, the consonants /r..z/. The prime *REEZ* itself is also briefly presented before being replaced by a visual pattern mask such as #####. The pattern mask ends visibility of *REEZ*. Backward masking strictly limits the time that *rose* and *REEZ* are visible, which limits the time available to derive phonology. If phonology becomes available rapidly, then the interaction of *REEZ* and *rose* phonology should benefit identification of *rose*, compared to a control condition.

The backward priming paradigm revealed that phonology is available very soon after seeing a word. Berent and Perfetti (1995) demonstrated that consonant phonology of pseudoword primes such as *REEZ* is available 20–40 ms after the pseudoword becomes visible (cf. Lee, Rayner, & Pollatsek, 2001; Perry & Ziegler, 2002; but cf. Lukatela & Turvey, 2000). Colombo, Zorzi, Cubelli, and Brivio (2003) established that both consonant and vowel phonology of printed Italian are available under the same conditions of brief visibility. And Lukatela, Frost, and Turvey (1998) demonstrated that the phonology of pseudohomophones such as *KLIP* is available within a 29 ms window of visibility (see also Berent & Van Orden, 2000, 2003; Lee, Rayner, & Pollatsek, 1999; Perfetti, Bell, & Delaney, 1988; Rayner, Sereno, Lesch, & Pollatsek, 1995; Xu & Perfetti, 1999).

Rapidly available phonology is at least consistent with the possibility that phonology is a mediating cause in word comprehension (Frost, 1998; see also Frost, this volume).

Yet it is one thing to demonstrate that phonology is rapidly available and another thing to demonstrate that phonology has priority over direct access. Ziegler, Ferrand, Jacobs, Rey, and Grainger (2000) conducted an incremental priming study to explore the latter issue. Incremental priming allows a continuous manipulation of how one letter-string may affect another, how a prime nonword may affect a target word, for instance. The beauty of incremental priming is its precise control over when primes are available. The duration or intensity of priming words can be changed incrementally from a range in which primes do not benefit the identification of target words to a range in which they do. This adds a dimension of control that is missing in most other priming studies (Jacobs, Grainger, & Ferrand, 1995).

Ziegler and his colleagues examined the relative priority of phonology versus direct access using forward masking and they conducted the experiment in French. Forward masking rearranges the order of events compared to backward masking. A forward-masking trial briefly presents a mask (#####), which is quickly replaced by a prime, which is in turn replaced quickly by a target word. The target word remains visible until the participant responds.

The experimental manipulation consisted of three priming conditions that differ in similarity between prime and target. In one condition the primes were similar to targets in spelling and identical to targets in French phonology (e.g., pseudohomophone *nerf* for target word *NERF*). This condition was called the O+P+ condition, O+ implying similar orthography between prime and target, and P+ implying similar phonology. Their second condition O[minus]P+ presented primes that were dissimilar in spelling but identical in phonology (e.g., pseudohomophone *nair* for target *NERF*). And their third condition O+P[minus] presented primes that were similar to targets in spelling but dissimilar in phonology (e.g., nonword *narf* for target *NERF*).

In all three prime conditions, lexical decisions to targets showed facilitation from priming compared to a no-prime control condition. A facilitation effect equals the degree to which the prime reduces the latency of the target "word" decision-time, compared to a baseline. In the facilitation calculus of Ziegler et al. (2000), the slower response time in the O[minus]P+ condition minus the faster response time in the O+P+ condition estimates the facilitation effect of spelling similarity – the direct access effect. Likewise, O+P[minus] minus O+P+ estimates facilitation due to similar phonology – the mediating phonology effect.

With a prime duration of 29 ms, the facilitation calculus revealed a greater magnitude of facilitation due to similar spelling compared to facilitation due to similar phonology. Similar spelling outdid similar phonology and so direct access must have priority over mediated access from phonology (see also Ferrand & Grainger, 1992, 1993, 1994). Other studies using masking and priming paradigms have found comparable patterns that sometimes include null effects of similar phonology, which seems to reinforce the case for priority of direct access (Brysbaert & Praet, 1992; Davis, Castles, & Iakovidis, 1998; Shen & Forster, 1999; Verstaen, Humphreys, Olson, & D'Ydewalle, 1995). Again however the story is not so simple; the pattern of facilitation changes if a different task is used.

In a comparable word naming study, Ziegler et al. (2000) observed results that contradict the pattern from lexical decision. In word naming, similar phonology appears to

outdo similar spelling at all prime durations and the pattern becomes statistically reliable at a prime duration of 42 ms (see also Montant & Ziegler, 2001). In this case, the results suggest that mediated access from phonology has priority over direct access. So which task demands are most comparable to the demands of natural skilled reading – forward masking or backward masking, 29 ms or 42 ms, lexical decision or naming, or none of the above? The story only gets murkier. How one may interpret Ziegler et al.'s (2000) lexical decision results rests on debatable assumptions about similarity and activation, and a possible confound, which is discussed next.

The facilitation logic depends on whether similarity has been straightforwardly added in or subtracted out of relations between primes and targets. This may not be the case for the O+P[minus] lexical decision primes. The O+P[minus] condition was supposed to entail a reduction in similar phonology between primes and targets compared to the O+P+ condition. The contrast between the conditions was meant to isolate the facilitation due to the more similar phonology of the O+P+ condition: O+P[minus] response times minus O+P+ response times estimated facilitation due to more similar phonology. However, in the O+P[minus] “priming condition the consonantal skeleton is typically maintained”; for example *n[mac]a[/mac]ɹf* – *N[mac]a[/mac]RF* (Ziegler et al., 2000, p. 687). This creates a confound whereby O+P[minus] and O+P+ priming can be almost identical at very short prime durations. As a consequence, the magnitude of facilitation due to similar phonology is systematically underestimated.

Consonant phonology is more quickly available than vowel phonology in languages with predominantly ambiguous vowel spellings (Berent & Perfetti, 1995; Lee et al., 2001; Perry & Ziegler, 2002). The earliest moments of activation emphasize reliable correspondences between consonant spelling and phonology, and the O+P[minus] primes share these reliable correspondences with their targets. This means that O+P[minus] primes are comparable to the O+P+ primes in their potential for facilitation in the earliest moments of activation. As a consequence, phonology priming is underestimated at the shortest prime durations, such as the 29 ms duration. The contrasted conditions O+P[minus] versus O+P+ could only be expected to diverge at longer prime durations, as vowel phonology comes into play. This confound undermines the contrast in the 29 ms condition that seemed to favor similar spelling over similar phonology. The confound renders the lexical decision outcome equivocal; it no longer favors direct access.

Priming manipulations that contrast degrees of similarity are often problematic. How does one discount the similarity in phonology that is inherent when items are similar in spelling? Some accounts claim that the first instants of word comprehension include multiply active patterns of phonology that, over time, settle into a single pattern (e.g., Kawamoto & Zemblidge, 1992; Van Orden et al., 1990; Van Orden & Goldinger, 1994). Consequently, items such as *plaid* and *plain* would activate virtually identical “clouds” of phonology in the first milliseconds, but they are not identical in spelling and do not settle into the same phonology. Other accounts assume that the first milliseconds of word comprehension include incompletely specified phonology. This assumption also allows that similar spellings may activate identical phonology at the outset of word comprehension (e.g., Berent & Perfetti, 1995; Frost, 1998).

Finally, how does one insure that similarity along a phonology dimension is ever comparable in magnitude to similarity along a spelling dimension? Do null effects of similar

phonology stem from weak manipulations of similarity? Sometimes yes; other times nobody knows (e.g., Frost, Ahissar, Gotesman, & Tayeb, 2003). Again, how one interprets the idiosyncratic task conditions that produce the evidence determines how one interprets the evidence, and there are inestimable degrees of freedom for interpretation of task demands. Like homophone errors, the evidence from masked priming studies leaves the causal status of phonology representations in skilled reading undecided. Some special task demands yield reliable effects of phonology variables, and others do not.

Giving Up Ether

Ideally, robust phonology effects would be found in all laboratory reading contexts. Ideally, laboratory methods should reveal a blueprint of reading that is independent of the laboratory tools used in the investigation. With respect to this ideal, a phonology effect that cuts across all reading contexts would satisfy the requirements (Jacobs & Grainger, 1994). Mediating phonology would then become an accepted component in the architecture of word comprehension. But so far no one has discovered a generally robust phonology effect in skilled reading, as the previous sections illustrate. The consequence is a perennial debate about phonology and reading.

Homophone errors and priming studies illustrate why the question of phonology has not found a satisfactory answer. Evidence that favors phonology is the product of special task demands, and evidence that favors direct access is also the product of special task demands. Idiosyncratic findings from idiosyncratic task conditions simply fuel the long-standing controversy about phonology and reading (cf. Frost, 1998; Pollatsek & Rayner, 2003; Van Orden et al., 1999; 2001). And yet, although no particular phonology effect can be found to familiar words in all contexts, a phonology effect of some kind can be found in most contexts. In some cases, the task contexts that yield a null effect of one phonology variable also yield a positive effect of a different phonology variable (Berent, 1997; Berent & Van Orden, 2003). Thus neither direct access nor mediating phonology can claim unequivocal empirical support.

The dilemma that phonology effects present leaves the contest between mediating phonology versus direct access in stalemate. As a consequence, the question of whether phonology mediates comprehension in reading is never answered to the satisfaction of reading scientists. Capricious phonology effects do not dissuade scientists who believe that skilled reading is a visual process of direct access (e.g., Coltheart, 2000; Daneman & Reingold, 2000; Davis et al., 1998; Shen & Forster, 1999; Verstaen et al., 1995). Null effects do not dissuade scientists who believe that reading is a linguistic process that includes mediating phonology (e.g., Frost, 1998; Liberman, 1992, 1999; Lukatela & Turvey, 1998; Rayner, Pollatsek, & Binder, 1998).

One idea about how to resolve the dilemma would be to add more factors into the equation. Plausibly, reading scientists could isolate factors such as task demands or participant strategies. However, simply manipulating more factors to collect more data is likely to further complicate matters, as more and higher-order interactions with context arise (Gibbs & Van Orden, 1998; Stone & Van Orden, 1993; Van Orden et al., 1999,

2001). If all factors interact, then attempts to isolate any single factor will meet the same fate as phonology factors.

Perhaps another way can be found around the stalemate, one that circumvents the exclusively reductive logic that guides mechanistic explanations. It may help to look elsewhere, to other science that has successfully abandoned mechanistic explanation, science that has already confronted and moved past a comparable dilemma. For example, the present dilemma of reading scientists is analogous to the dilemma of nineteenth-century physicists concerning the concept of ether. Ether, at that time, defined an absolute frame of reference for all movement. To sustain this belief physicists accumulated many ad hoc assumptions. To explain why ether's motion does not disturb matter, they assumed that ether does not interact with matter. But to explain why the velocity of light changes when it passes through glass or water, they were forced to assume that ether does interact with matter. "In other words, there is an interaction between ether and matter in optical phenomena, but none in mechanical phenomena! This is certainly a very paradoxical conclusion!" (Einstein & Infeld, 1938, p. 120).

For reading scientists a context independent architecture of reading takes the role of ether, and task demands take the role of matter. When reliable phonology effects are observed, pro-phonology scientists assume the architecture of reading did not interact with task demands. But to explain null effects of phonology, the architecture must have interacted with task demands, which concealed the effects of phonology. Anti-phonology reading scientists accumulate inverted assumptions. When null phonology effects are observed, then the architecture of direct access did not interact with task demands. But to explain positive effects of phonology, the architecture of reading must have interacted with task demands, which produced the sham phonology effects. Both camps resemble nineteenth-century physicists in their paradoxical conclusions about reading and task demands.

One hundred years of stalemate warrants rethinking the traditional assumptions about reading. The assumptions are in conflict with actual reading phenomena. Phonology factors interact with other word factors and their pattern of interaction changes with changing task demands. All effects of all reading factors are in motion, so to speak, with respect to each other and with respect to the task contexts in which they are observed. One may take this flexible catalog of interaction effects at face value – no absolute frame of reference presents itself. Reading scientists may follow the lead of physicists and give up altogether a baseless idea of an absolute frame of reference.

Summary and conclusion

The role of phonology in skilled reading remains unclear. Looking across this now vast literature, the appearance and disappearance of phonology effects does not divide neatly among processes of word comprehension. The appearance and disappearance of phonology effects are only captured in reliable high-order interactions among reading histories of participants, word factors, and the special circumstances of task demands. Apparent phonology effects and other reading effects are context sensitive by their very nature. Such context sensitivity is symptomatic of complex interactive systems. The next sections

review evidence that more pointedly suggests that phonology and other processes always join in a complex interaction.

Spelling and Phonology in an Interactive System

The role of phonology in skilled reading can be considered in terms of interactive processes. The question is now asked with respect to variables that modulate the performance of an interactive system, rather than with respect to isolated causal factors. One prominent variable is ambiguity.

Notice how many ways the same ambiguous phoneme /*[mac]a[/mac]*/ can be spelled in *Kay*, *weigh*, *made*, and *pail*, or how many ways the same ambiguous vowel spelling *ai* can be pronounced in *plaid*, *raid*, *said*, and *aisle*. Such ambiguity has consequences for performance. For instance, if a presented spelling can be pronounced in more than one way, then it yields a slower naming time compared to an unambiguous spelling, all other things equal.

Ambiguity is not your standard causal factor. Ambiguity effects cannot be localized in spelling or phonology taken separately. Ambiguity is only defined in a relation between the two. Thus empirical tests for ambiguity effects are tests about how phonology is related to other aspects of language, such as spelling. The next sections of this chapter describe empirical findings that demonstrate ambiguity effects.

Simulations of interactive processes

Before turning to the experiments, briefly consider some previous simulations of interactive processes. Simulations of interactive processes among spelling, phonology, and semantics have changed the way scientists look at the structure of language (e.g., Grossberg & Stone, 1986; Jacobs, Rey, Ziegler, & Grainger, 1998; Kawamoto & Zemplidge, 1992; Masson, 1995; McClelland & Rumelhart, 1981; Plaut, McClelland, Seidenberg, & Patterson, 1996). They have focused scientists' attention on ambiguity and the statistical structure of language (Plaut et al., 1996; Saffran, 2003; Van Orden et al., 1990), and they introduced the possibility of feedback in word comprehension.

Consider the ambiguous spelling of the homograph word *wind*. *Wind* has two legitimate pronunciations: it can rhyme with *pinned* or *find*. In an interactive model, spelling nodes representing *wind*'s spelling activate nodes that represent the two pronunciations of *wind*, and these two pronunciations both feed back activation to their common spelling. This creates two competing feedback loops, which characterizes how ambiguity is expressed in an interactive model. Ambiguity breeds competition between multiple potential outcomes, which takes time to resolve (see also Lupker, this volume).

Kawamoto and Zemplidge (1992) simulated the competition between homograph pronunciations as it unfolds across a naming trial. The model included feedforward and feedback connections among letter, phoneme, and semantic node families. The connections modulate node activity very roughly as synapses may modulate the activity of

neurons. In the Kawamoto and Zemblidge model, connections were excitatory between node families but mostly inhibitory within node families. For instance, letter nodes excite phoneme nodes and phoneme nodes excite letter nodes, but competing phoneme nodes inhibit each other. Consequently, phoneme nodes compete directly with other phoneme nodes and indirectly with letter or semantic nodes. A phoneme node competes indirectly by activating some particular letter or semantic node that can compete directly. Thus every node interacts with every other node, either directly or indirectly.

Simulations have been successful as guides for how to look at language. They are less successful as models of actual psychological processes. Despite highly unintuitive and yet reliable predictions, actual simulations are perpetually challenged by the details of human performance (e.g., Spieler & Balota, 1997; Treiman, Kessler, & Bick, 2003). It is the assumptions behind the simulations that seem to capture a reliable picture of language, but painted in somewhat broad strokes.

Ambiguity at the scale of whole words

Homographs like *wind* have a dominant pronunciation (the more regular pronunciation that rhymes with *pinned*) and a subordinate pronunciation (the less regular pronunciation that rhymes with *find*). In an actual word naming experiment, some readers will produce the dominant pronunciation and some will produce the subordinate. Also, when the dominant pronunciation is produced, it yields faster naming times, on average, than the subordinate pronunciation. One way to think about this pattern is that the two pronunciations compete in the course of a word naming trial prior to an observed pronunciation.

In a simulated naming trial, *wind*'s subordinate pronunciation is less strongly activated, at least initially, but nevertheless can win the competition. To do so it must accrue sufficient activation, within the time course of the trial, to overcome activation of the dominant pronunciation. This implies an on-line qualitative change from dominant to subordinate phonology. The qualitative change occurs at an exchange point in what is called a bifurcation. Kawamoto and Zemblidge (1992) simulated the bifurcation of a homograph pronunciation, from statistically dominant to subordinate, as a transcritical bifurcation.

The dominant pronunciation of *wind* has a stronger feedback loop between letter and phoneme nodes, a stronger and more stable local attractor. The subordinate pronunciation has the weaker or less stable attractor between letter and phoneme nodes, but has the more stable attractor between phoneme and semantic nodes. The feedback loop between phoneme and semantic nodes takes some time to grow in strength and lend sufficient support to *wind*'s subordinate pronunciation. Enough support makes *wind*'s subordinate pronunciation a winner. This outcome occurs when a reader or model is sufficiently more familiar with the subordinate pronunciation's semantic variants, which counters the inherent disadvantage of the subordinate pronunciation's less-regular relation between spelling and phonology.

Initially *wind* activates the two pronunciation patterns, and the dominant pattern is initially favored. However, slowly accruing activation in a semantic and phoneme feed-

back loop lends increasing support to the subordinate pronunciation. Within the time of a naming trial, activation in the phoneme-semantic feedback loop grows to a sufficient degree that it turns the tide in the competition. The tide turns at the bifurcation point. Within the time between the appearance of *wind* and a pronunciation, semantic-phoneme activation and the subordinate's letter-phoneme activation overtake the otherwise dominant pronunciation. At the bifurcation point, semantic-phoneme feedback puts *wind*'s subordinate pronunciation over the top, and the dominant pronunciation exchanges stability with the less-regular subordinate pronunciation. Subsequently, the model produces the subordinate pronunciation.

So why do some readers produce the dominant pronunciation and others the subordinate? Different readers, or models, may sample language differently. Each reader has a unique history of covariation among words' spellings, phonology, and semantics. Pronunciations can have strong or weak ties to semantics based on different readers' different familiarity with different words. At any particular time, some readers will quickly produce the dominant more regular pronunciation, and other readers, sufficiently more familiar with subordinate variants, will more slowly produce the subordinate pronunciation.

Wind's homograph spelling is one ambiguous spelling, one pocket of ambiguity, within a reader's accumulated sample of English. Yet *wind* is only ambiguous if that reader's history includes samples of both interpretations of *wind*. A reader's sample of a language delimits the potential for ambiguous or unambiguous relations. The aggregate statistical pattern of relations that makes up a reader's language is specific to the reader's history and changes throughout a lifetime of reading.

Multiple scales of ambiguity

Homograph ambiguity exists at multiple scales. In a homograph, every letter has associations with different pronunciations. For example, the homograph *wind* is ambiguous at a micro scale because its grapheme *i* is ambiguous. This ambiguity is amplified at a meso scale of *wind*'s ambiguous spelling-body *-ind*, and is further enlarged at a macroscale of the ambiguous whole word. In this way of thinking, local ambiguity is infectious, in a manner of speaking. A local ambiguity, like *wind*'s ambiguous grapheme *i*, infects every larger scale of spelling that has a history of multiple pronunciations.

Words infected with more ambiguity have slower naming times. Compare the homograph *wind* with the word *pint*. *Pint*'s spelling is also ambiguous but not to the same degree as *wind*. *Pint* is infected with ambiguity up to the scale of its spelling-body *-int*, but *pint* does not entail whole-word ambiguity. The difference explains why homograph pronunciations are slower than pronunciations to ambiguous control items that are not homographs (Gottlob, Goldinger, Stone, & Van Orden, 1999; but cf. Hino, Lupker, & Pexman, 2002). This outcome would be observed even if every letter of *wind*, taken one at a time, were no more ambiguous than the individual letters of *pint*. *Wind* has a slower naming time even when contrasted with precisely constructed *mint* and *pint* controls equated for spelling body ambiguity (Holden, 2002).

Connectionist models track in the same matrix all the scales at which spelling relates to phonology. They illustrate how all these relations can be co-instantiated; different levels

of representation are not necessary for the different scales to be effective. Models with recurrent feedback connections, in addition, track multiple scale relations in all directions. Stronger feedback loops like those of dominant relations correspond to relatively more stable attractors in the network. One can find dominant and subordinate relations at each scale, which means that dominant and subordinate relations may be nested across scales. In other words, there are relations within relations, attractors within attractors.

Now everything is in place to discuss feedforward ambiguity effects, and then feedback ambiguity effects, that have been demonstrated empirically. Ambiguity effects can be identified at the scale of spelling-bodies and graphemes and feedback effects can be identified at all the same scales.

Feedforward ambiguity at the scale of spelling-bodies

The more regular dominant pronunciation of the spelling body *-int* rhymes with *mint* (consider *lint*, *tint*, and *hint*). The subordinate pronunciation of *-int* rhymes with *pint*. *Pint* takes longer to name than *mint* because *pint*'s rime is the subordinate pronunciation of the body *-int*. When *pint* is the word to be named, a mispronunciation of *-int* to rhyme with *mint* strongly competes with *pint*'s correct pronunciation. This competition is so close that a mispronunciation of *pint* can be elicited even from skilled readers. For example, participants can be trained to respond rapidly, in time with a beat, in a word naming task, but in doing so they commit errors of pronunciation including the kind of error in which *pint* is mispronounced to rhyme with *mint* (Kello & Plaut, 2000). More slowly emerging semantic features must combine with *pint*'s correct pronunciation to counter the dominant rhyme with *mint* (Farrar & Van Orden, 2001). In this case, *pint*'s rhyme with *mint* is not a word and would not have coherent semantic associations (cf. Lesch & Pollatsek, 1998, however).

When *mint* is the word to be named, the subordinate mispronunciation that would rhyme with *pint* competes with *mint*'s correct pronunciation. Just as for homographs, the two pronunciations compete in the course of a naming trial prior to an observed pronunciation, and the competition takes time to resolve. Thus naming times to *mint* should be slower than to words with unambiguous body-rime relations. Compare the spelling body *-int* with *-uck*, the spelling body of *duck*. *Duck*'s spelling body is unambiguous; it supports only one pronunciation (consider *luck*, *buck*, *muck*, and *puck*). The /uk/ rime also reliably covaries with the *-uck* body. Together they form an invariant relation between body and rime, and rime and body. Indeed, words like *mint* are more slowly named than words like *duck* (Glushko, 1979). A word like *mint* is more widely infected with ambiguity than a word like *duck*.

Feedforward ambiguity at the scale of graphemes

Pockets of more or less ambiguity are also found at the microscale of graphemes and phonemes (compare Zorzi, Houghton, & Butterworth, 1998). English vowel spellings are almost always ambiguous. But some English consonants have invariant relations with phonology. The consonant grapheme *d*, at the beginning of a word, is always associated

with the phoneme /d/, and the /d/ phoneme is always spelled *d*. Overall, in English, consonant spellings covary more reliably with their pronunciations than do vowel spellings. Consequently, in English, consonant phonology is resolved earlier than vowel phonology. For example, the relative ambiguity of consonant and vowel spellings predicts when their phonology will become available in masked priming experiments: consonant phonology coheres before vowel phonology (Berent & Perfetti, 1995; Lee et al., 2001; Perry & Ziegler, 2002).

For a visually presented word, the mapping from spelling to phonology is the feed-forward relation and the mapping from phonology to spelling is the feedback relation. For auditory presentations this is reversed. The mapping from phonology to spelling is feedforward and the mapping from spelling to phonology is feedback. The previous examples all concerned ambiguity from spelling to phonology. Ambiguity effects also generalize to the inverted mapping from phonology to spelling, as feedback in visually presented homophones for instance.

Feedback ambiguity at the scale of whole words

Consider the homophone phonology /braik/ and the corresponding spellings *break* and *brake*. Just as the homograph *wind* supports two pronunciations, the homophone /braik/ supports two spellings. Homophone words produce slower visual lexical decision times than control words that are not homophones (Ferrand & Grainger, 2003; Pexman, Lupker, & Jared, 2001; Pexman, Lupker, & Reggin, 2002). Homophones have slower lexical decision times even when contrasted with precisely constructed controls equated for rime-body ambiguity – feedback effects accrue across scales (Holden, 2002).

Notice that homophone effects in visual lexical decision are unintuitive. From the traditional view, activation should always flow forward from a cause to an effect, as from spelling to phonology in a visual lexical decision task. It should not matter for visual lexical decisions that *break's* pronunciation /braik/ may have more than one spelling, unless there exists feedback from phonology to spelling. Consequently, slower visual lexical decision times to homophone words imply feedback from phonology to spelling.

Feedback ambiguity at the scale of spelling-bodies and pronunciation-rimes

A feedback ambiguity effect at the scale of pronunciation-rimes and spelling-bodies is found in visual lexical decision. For instance, the English word *hurt* has an ambiguous rime (/ûrt/) that is linked to more than one spelling body (-urt, -ert, -irt). As one might expect, by now, words like *hurt* with ambiguous rimes yield slower visual lexical decision times and more errors than words with unambiguous rimes (Stone, Vanhoy, & Van Orden, 1997; Ziegler, Montant, & Jacobs, 1997; see also Seidenberg & Tannenhaus, 1979).

The rime ambiguity effect in visual lexical decision has yielded controversy, including claims that the original studies did not truly produce ambiguous rime effects (e.g., Peere-man, Content, & Bonin, 1998). One can understand why feedback effects stir up the nest. Nevertheless, new studies with increasingly precise control continue to find reliable

feedback ambiguity effects (Holden, 2002). Also, once feedback ambiguity is taken into account, reliable feedforward ambiguity effects emerge in both visual and auditory lexical decision performance, effects previously thought to be unreliable (Stone et al., 1997).

Another feedback effect at the same scale is found in auditory lexical decision. The word *pint* in an auditory lexical decision is spoken, but ambiguity in how *pint*'s spelling might be pronounced slows down the lexical decision time – even though no spellings ever appear in the experiment! What is feedforward for a visually presented word is feedback for an auditory presentation of the same word. The feedback relation for an auditory presentation runs from the spelling body to the pronunciation rime. Thus the fact that auditory lexical decisions are slower to words such as *pint*, with ambiguous body spellings, is a fact that corroborates feedback (Ziegler & Ferrand, 1998; Ziegler, Muneaux, & Grainger, 2003).

The feedback effect of *pint*'s ambiguous spelling body in auditory lexical decision is extremely counterintuitive from a traditional perspective. According to that perspective, spoken word recognition should be independent of the spelling of a word. Yet printed language interacts with spoken language in situations where it could just as well leave spoken language alone, as the feedback effect demonstrates (see also chapters by Lupker and Morais, current volume).

Feedback ambiguity at the scale of phonemes

Damian and Bowers (2003) examined phoneme ambiguity as a feedback phenomenon in a speech production task. The words *coffee* and *cushion* share the same initial phoneme and grapheme. If *coffee* is a cue to, say, *cushion*, then *cushion*'s voice onset will be faster than to control items. The sequencing of words that share initial phoneme and grapheme creates a benefit for saying *cushion* aloud.

Contrast *camel* and *kidney* with the pair *coffee* and *cushion*. *Camel* and *kidney* share an ambiguous phoneme spelled *c-* in *camel* and *k-* in *kidney*. If *camel* is the cue to say *kidney*, then *kidney*'s voice onset is slower than *cushion*'s and no faster than to control items. The repeated phoneme with different graphemes erases the previous benefit of sequencing the same phoneme with identical graphemes.

The feedback effect of the phoneme's two different spellings is found when the cues are printed words and when they are spoken words. When the cues are spoken words no spellings appear in the experiment. Why should it matter for spoken word production that a repeated phoneme has a different spelling in each repetition? It only matters because spoken word production includes a feedforward and feedback interaction between phonology and spelling (Dijkstra, Roelofs, & Fieuws, 1995).

Reminders

The previous sections reviewed ambiguity findings, findings that corroborate a complex interaction among spelling and phonology in visual and spoken language. Several other

effects lend themselves to this framework, but do not fit so neatly into the previous story about scales of ambiguity. These additional findings concern interactions between semantics and surface forms, letter perception, and the possibility that relations are the source of perceived lexicality. These findings will be reviewed next. Also, there are theoretical and methodological loose ends that have not been discussed previously. One loose end concerns languages other than English, another concerns whether dual-process theory has been falsified, and a final point, for future reference, concerns how response times from reading tasks should be viewed.

Ambiguity between semantics and surface forms

The relations among words' spellings, phonology, and their semantics all matter for sorting out ambiguity effects. For instance, ambiguous feedback from semantics can affect lexical decision and naming performance. More ambiguous semantic features are associated with more spellings and pronunciations and words can be more or less ambiguous in this relation between semantics and surface forms. More ambiguous words produce slower lexical decision and naming times (Pecher, 2001; Pexman, Lupker, & Hino, 2002). The effect is a striking parallel to ambiguity effects in the relations between spelling and phonology. Semantic ambiguity effects may prove to be the most important effects mentioned so far. After all they concern relations with words' meanings and it is the pursuit of meaning that drives word comprehension in reading.

Letter perception

A briefly presented pseudohomophone such as *brane* can induce the false impression that a pre-designated letter *i* was seen (Ziegler, Van Orden, & Jacobs, 1997). Participants report that an *i* appeared in the presented spelling *brane*, but only if the letter is contained in *brane*'s sound-alike base-word *brain*. The flip side of this effect is also observed. Pseudohomophones such as *taip* may induce the false impression that a pre-designated letter *i* did not appear, but only if the letter is missing from *taip*'s sound-alike base-word *tape*. These phenomena were first demonstrated in German (Ziegler & Jacobs, 1995), then later in English (Ziegler et al., 1997) and French (Lange, 2002). Such phenomena appear quirky within a conventional framework where they may suggest postlexical inferences about which letters were seen. They are expected, however, if feedback from base-word phonology activates *brain*'s letters or inhibits letters that are not present in the base-word *tape*.

Perceived lexicality

Relations between spelling and phonology are sources of perceived lexical structure (Vanhooy & Van Orden, 2001). Wordlike body-rimes actually add "word-ness" to letter strings that are not words. For example, it is widely reported that lexical decisions to

pseudohomophones such as *jale* are slower and more likely to end in a false 'word' response than are control items. Also, correct 'word' responses to actual words are slower when pseudohomophones appear as foils.

It is not simply that pseudohomophones mimic word phonology; it also matters that they are composed of body-rime relations like those found in actual words. *Jale* is constructed on an extant body-rime that appears in the words *bale*, *sale*, and *tale*. The pseudohomophone *stahp*, which sounds like *stop* in American English, is constructed on a novel body-rime that does not appear in an actual word. In lexical decision, pseudohomophones like *jale* produce reliable pseudohomophone effects; pseudohomophones like *stahp* do not.

Natural variation across languages

Each language presents a unique compilation of ambiguity that will be uniquely sampled by each reader. Hebrew includes mostly homographs, and Chinese includes very many homophones. Dutch, Spanish, German, and Italian minimize or eliminate ambiguity between phonology and spelling by staying closer to a system of grapheme–phoneme rules. French is more like English. French has ambiguities at multiple scales of correspondence between phonology and spelling. Serbo-Croatian has two alphabets that sometimes contradict each other in their relation to phonology, and other times not. Clearly, the consequences of ambiguity for complex interactions must be worked out carefully one language at a time (e.g., Frost, this volume; Colombo et al., 2003; Frost, Katz, & Bentin, 1987; Goswami, Ziegler, Dalton, & Schneider, 2003; Lukatela & Turvey, 1998; Ziegler, Perry, Jacobs, & Braun, 2001; Ziegler et al., 2000; and many other publications not cited here). Different languages exaggerate or reduce different sources of ambiguity and all sources interact in performance (Bosman & Van Orden, 1997; Lukatela & Turvey, 1998; Van Orden & Goldinger, 1994).

Is dual-process theory false?

The spectrum of ambiguity effects and feedback effects that experiments demonstrate would not likely have been anticipated with dual-process theory as the guide. However, this does not mean that dual-process theory is false. Findings that contradict dual-process theory simply reveal that grapheme–phoneme rules were not the best compass to discover salient structure between phonology and spelling (Paap et al., 1992). The theory itself can be reconstituted indefinitely to absorb new contradictory findings (e.g., Coltheart, Curtis, Atkins, & Haller, 1993; Norris, 1994; Zorzi et al., 1998). Ad hoc changes create alternative ways to see the contradictory data and can be useful for that fact (Feyerabend, 1993). Nonetheless, it is a bit hard to imagine how scientists in the exclusive pursuit of mechanistic causal chains would have stumbled on these effects. The discovery of feedback effects as predicted by feedback models is a remarkable discovery of basic reading science with profound implications for all cognitive science.

What is the nature of response time?

The last point is a caveat that concerns how one should look at the data from all these experiments. The previous discussion has emphasized mean effects, differences between average response times or accuracy, as did almost all of the cited authors. This will prove in time to have been misleading. Ambiguity effects are not so simply expressed; they do not simply reflect shifts in average response times. Rather, they largely reflect increases in the proportion of very slow responses. They reflect redistribution of response times and changes in the shapes of response time distributions (Holden, 2002). This general observation about effects and response times is not new to reading science (Andrews & Heathcote, 2001; Balota & Spieler, 1999), but its implications have not been widely acknowledged.

Redistributions of response times often appear as changes in so-called power laws – equations in which the probability of a particular response-time is a function of the response-time itself (Holden, 2002; Van Orden, Moreno, & Holden, 2003). Power laws may suggest a complex interdependence in which the processes that compose a system change each other as they interact (Jensen, 1998). Consequently, co-instantiated relations between phonology and spelling, for example, become causally entwined and interdependent (Van Orden & Holden, 2002; Van Orden et al., 2003). It is the nature of living systems that they comprise entwined processes and do not reduce to causal elements (e.g., Rosen, 2000; Wilson, 2003).

Power law behavior could imply a radical suggestion that separate representations of phonemes and letters, for example, need not be posited. Relations between a word's spelling and its phonology, its body and rime, and its graphemes and phonemes become mutually reinforcing relations with neither being causally prior to the other. Yet there remains a useful way to think about cause in the sense of a basis or foundation for reading. Unless a child becomes attuned to the alphabetic principle in relations between spelling and phonology, learning to read does not occur or occurs with great difficulty (Rayner, Foorman, Perfetti, Pesetsky, & Seidenberg, 2001). In this sense of cause, the alphabetic principle has a causal priority in the development of skilled readers.

Summary and Conclusions

The first half of this chapter ended on the horns of the dilemma concerning phonology and skilled reading. Over 100 years of reading research failed to decide whether skilled reading involves mediating phonology, or whether it does not. The question of mediating phonology hinges on the discovery of a task independent phonology effect for skilled readers reading familiar words. This discovery could possibly situate phonology in the architecture of word comprehension, part of cognition's larger absolute frame of reference. However, despite the plausibility that such a phonology effect could exist, all phonology factors, like all other word factors, change the pattern of their effects across the variety of task conditions.

The second half of this chapter reviewed ambiguity effects at multiple scales of relations between spelling and phonology. The reviewed findings present snapshots of a complex structure that relates phonology and spelling. In the contemporary picture of English, this relation appears as a context sensitive, bidirectional, statistical structure that changes on multiple scales and in each instance of reading – a statistical structure in perpetual motion, one might say. The complex structure of ambiguity effects intertwines written and spoken English in feedback. Some prominent intertwined relations are readily discernible, relations like those between bodies and rimes, or graphemes and phonemes. Nonetheless, the intention is not to propose a pretty hierarchy, and it would soon sprout weeds in any case. Letters and groups of letters change their relation to phonemes and groups of phonemes according to the contexts in which they appear.

Feedback models of interacting processes predict ambiguity and feedback effects. Context sensitivity within these models is useful to explain the context sensitivity of relations between spelling and phonology. It is a natural extension of this view to expect context sensitivity at all levels of a system, including sensitivity to the laboratory contexts of task demands. Until now context sensitivity has been a reason not to take some other scientist's data as conclusive. Now context sensitivity is the likely key to understand reading, the paradigmatic cognitive performance.¹

Note

1. We acknowledge support from the National Science Foundation. The views expressed in this chapter are the authors' and do not necessarily represent those of the National Science Foundation.