COGNITIVE NEUROPSYCHOLOGY, 2006, 23 (2), 278–307

Psychology Press
Taylor & Francis Group

Grapheme-to-lexeme feedback in the spelling system: Evidence from a dysgraphic patient

Michael McCloskey
Johns Hopkins University, Baltimore, USA

Paul Macaruso
Community College of Rhode Island, Warwick, and Haskins Laboratories, New Haven, USA

Brenda Rapp
Johns Hopkins University, Baltimore, USA

This article presents an argument for grapheme-to-lexeme feedback in the cognitive spelling system, based on the impaired spelling performance of dysgraphic patient CM. The argument relates two features of CM's spelling. First, letters from prior spelling responses intrude into subsequent responses at rates far greater than expected by chance. This letter persistence effect arises at a level of abstract grapheme representations, and apparently results from abnormal persistence of activation. Second, CM makes many formal lexical errors (e.g., *carpet* → *compute*). Analyses revealed that a large proportion of these errors are "true" lexical errors originating in lexical selection, rather than "chance" lexical errors that happen by chance to take the form of words. Additional analyses demonstrated that CM's true lexical errors exhibit the letter persistence effect. We argue that this finding can be understood only within a functional architecture in which activation from the grapheme level feeds back to the lexeme level, thereby influencing lexical selection.

INTRODUCTION

Like other forms of language processing, written word production implicates multiple levels of representation, including semantic, orthographic lexeme, grapheme, and allograph levels. In this article we explore how the levels interact: Does processing proceed in a strictly feedforward manner, from earlier to later levels of representation? Or are there feedback as well as feedforward connections between levels, so that later levels can influence processing at earlier levels? We describe the performance of CM, a brain-damaged patient with an acquired spelling deficit, arguing from his error pattern that the cognitive system for written word production includes feedback connections from grapheme representations to orthographic lexeme representations.

Figure 1 depicts a general theory of the cognitive mechanisms implicated in written word production (see Tainturier & Rapp, 2001, for a recent overview of this theoretical framework and relevant evidence; for a somewhat different perspective, see Graham, Patterson, & Hodges, 1997, 2000). The major assumptions of the

---

Correspondence concerning this article should be addressed to Michael McCloskey, Department of Cognitive Science, Krieger Hall, Johns Hopkins University, Baltimore, Maryland 21218, USA (Email: michael.mccloskey@jhu.edu).

This study was supported in part by NIH grant NS22201 and NIMH grant R29MH55758. We thank Donna Aliminosa for testing patient CM, and Marie-Josèphe Tainturier for helpful comments on the manuscript.

http://www.psypress.com/cogneuropsychology

DOI:10.1080/02643290442000518

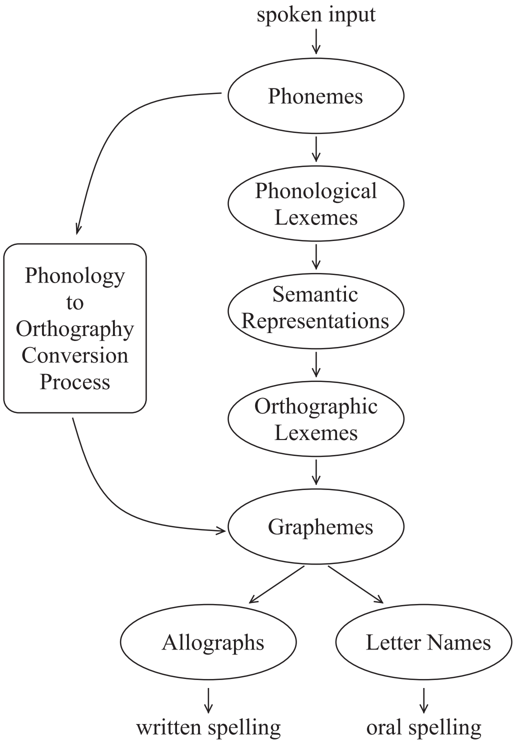

Figure 1. *Schematic depiction of the cognitive spelling mechanisms.*

theory may be described in the context of writing words to dictation. When a familiar word is dictated, its phonological lexeme is activated in a phonological lexicon. Activation of the lexeme leads to activation of a lexical-semantic representation within the semantic system, and the semantic representation in turn activates an orthographic lexeme in an orthographic lexicon. (Some theorists have proposed that orthographic lexemes may also be activated directly from phonological lexemes; see, e.g., Patterson, 1986; Roeltgen, Rothi, & Heilman, 1986; Romani, Olson, Ward, & Ercolani, 2002; but see Hillis & Caramazza, 1991b.)

From the orthographic lexeme, activation spreads to representations of the word's constituent graphemes. These abstract letter-identity representations provide the basis for activating allograph (letter-shape) representations appropriate for the desired form of written output (e.g., lower-case script, upper-case print), leading ultimately to production of writing movements. In the case of oral spelling, the abstract grapheme representations serve to activate letter-name representations.

When an unfamiliar word or pseudoword (e.g.,/floup/) is dictated, a somewhat different process is required, because no appropriate lexeme or semantic representations are available to be activated. According to the theory, graphemes corresponding to a plausible spelling (e.g., *flope* or *floap*) are activated by means of a sublexical phoneme–grapheme conversion process that exploits sound-to-spelling correspondences. As in the case of familiar word stimuli, the activated graphemes then activate letter shape or letter name representations.

Spelling processes may, of course, also be invoked by nonphonological stimuli. For example, when lexical-semantic representations are activated by pictures, objects, or thoughts the lexical item may be realised through written or oral spelling by activating orthographic lexemes, graphemes, and allographs just as in writing to dictation.

Although there are many issues regarding orthographic representation and processing that we have not discussed, the functional architecture we have sketched is sufficiently detailed to serve as a framework for the work reported in this article.

Bidirectional processing between segments and words

The present study concerns the connectivity, and hence the processing interactions, between orthographic lexeme and grapheme levels of representation. One possibility, illustrated in Figure 2a (and Figure 1), is that the flow of information between levels is strictly feedforward, such that semantic representations activate orthographic lexemes that, in turn, activate their constituent graphemes. An alternative, depicted in Figure 2b, is that in addition to the forward flow of activation, information from the grapheme

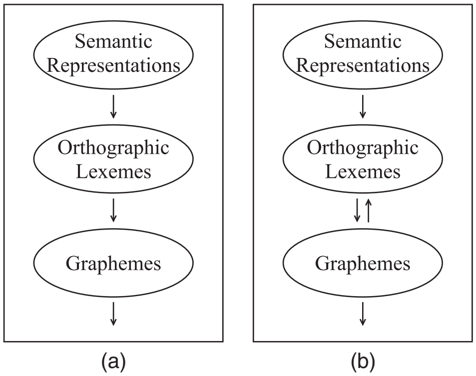

Figure 2. *(a) A strictly feedforward architecture. (b) A functional architecture with feedback between grapheme and orthographic lexeme levels.*

level influences lexical selection through the feeding-back of activation from grapheme to lexeme representations prior to lexical selection.

The question of unidirectional versus bidirectional connectivity between lexeme and segment levels has received little attention in studies of written word production (although see Romani et al., 2002, for a brief discussion). However, the issue has been intensively debated in research on spoken production. Evidence from the spoken production errors of both neurologically intact and neurologically impaired individuals has been put forward in support of phoneme-to-lexeme feedback. Several researchers have argued for feedback on the basis of results suggesting that form-based errors take the form of real words more often than expected by chance (the lexical bias effect; see, e.g., Baars, Motley, & MacKay, 1975; Best, 1996; Dell, 1986; Dell & Reich, 1980; Martin, Dell, Saffran, & Schwartz, 1994; Stemberger, 1985). Analyses demonstrating a significant influence of phonology on semantic errors have also been taken as evidence for phoneme-to-lexeme feedback (Dell, 1986; Dell & Reich, 1981; Harley, 1984; Martin, Gagnon, Schwartz, Dell, & Saffran, 1996). Although Levelt and colleagues (Levelt, Roelofs, & Meyer, 1999), as well as Garrett (1976), have claimed that the data can be interpreted within a strictly feedforward architecture, Dell and colleagues (Dell, Schwartz, Martin, Saffran, & Gagnon, 1997), Rapp and Goldrick (2000, 2003; Goldrick & Rapp, 2002) have argued that the evidence implies at least some influence of phoneme-level activation on the lexical selection process. In this paper we argue that the pattern of spelling errors produced by CM reveals an effect of grapheme-level activation on lexical selection, and hence provides strong evidence for feedback from grapheme to orthographic lexeme representations.

CASE HISTORY

CM is a right-handed man with a PhD in electrical engineering. He worked as a university professor until suffering a stroke in September 1986, at age 59. CT showed extensive cortical and subcortical damage in the distribution of the left middle cerebral artery. A premorbid writing sample, consisting of handwritten lecture notes, confirmed that CM's spelling was normal prior to his stroke. All 430 words in the sample were spelled correctly, including many long low-frequency words (e.g., *sinusoidal*).

General language assessment

Boston Diagnostic Aphasia Examination (BDAE; Goodglass & Kaplan, 1983)

CM performed well on the auditory comprehension tasks. He made no errors in word discrimination and body part identification, and scored 14/15 in following spoken commands and 10/12 in responding to questions involving complex ideational material. He also performed well in reading comprehension, obtaining high scores in symbol discrimination, word recognition, and word/picture matching. His comprehension score for reading sentences and paragraphs was 7/10.

In spoken language production CM showed impairment. He was able to recite automatised sequences and repeat single words, but scored only 2/16 in repeating phrases. The speech pathologist conducting the examination rated his speech at 3 (of a possible 7) for melodic line, and

4 for articulatory agility. Phrase length was two words. Confrontation and responsive naming were poor, and CM generated only 1 animal name in 90 seconds. He read aloud 3 of 10 words and no sentences.

In written language production the examiner rated CM at 5 (of 5) in writing mechanics (same as premorbid ability, with allowances made for use of nonpreferred hand). CM could write the alphabet in sequence, and had success writing numerals, letters, and short high-frequency words to dictation. Spelling to dictation was poor for long words and sentences. CM also performed poorly in written picture naming, and produced only 3 relevant words in a written narrative (rating 1 of a possible 5).

Peabody Picture Vocabulary Test-Revised (Dunn & Dunn, 1981)

In this test words were presented aurally one at a time, and CM indicated which of four pictures matched the word. On each of two administrations his score—147/175 and 161/175—fell within the average range.

Word–picture matching

On each trial in this task a single picture was first presented; a word was then presented aurally, and CM indicated whether the word matched the picture. Stimulus pictures were 144 line drawings of concrete objects (e.g., a stool). Each picture was presented three times, once paired with the correct name (*stool*), once with a semantic foil (*bench*), and once with a phonological foil (*stamp*). CM responded correctly on 98% (425/432) of the trials.

Spelling assessment

CM's written spelling was assessed with the Johns Hopkins University Dysgraphia Battery (Goodman & Caramazza, 1985). The battery probes spelling of dictated words and nonwords, as well as copy transcoding and written picture naming. Several factors (e.g., word length, concreteness) are varied systematically across the set of stimulus items. Normal control subjects perform all of the tasks with very low error rates.

Writing words to dictation

As shown in Table 1, CM was only 39% correct (128/326) in writing words to dictation. His accuracy was higher for high-frequency words (44%) than for low-frequency words (33%), although this difference was only marginally significant, χ^2 (1, $N = 292) = 3.70$, $.05 < p < .10$. No significant effects were observed

Table 1. *CM's performance on the Johns Hopkins Dysgraphia Battery*

Task	*No. of stimuli*	*No. correct*	*% correct*
Writing to dictation			
Word/Nonword status			
Words	326	128	39
Nonwords	34	4	12
Word frequency			
High	146	64	44
Low	146	48	33
Grammatical word class			
Nouns	28	10	36
Verbs	28	8	29
Adjectives	28	4	14
Functors	20	11	55
Concreteness			
Concrete	21	8	38
Abstract	21	6	29
Phoneme–grapheme conversion Probability			
High	30	15	50
Low	80	42	53
Word length			
4 letters	14	6	43
5 letters	14	3	21
6 letters	14	6	43
7 letters	14	5	36
8 letters	14	3	21
Written picture naming	51	13	25
Copy transcoding			
Direct			
Words	84	84	100
Nonwords	40	40	100
Delayed			
Words	84	59	70
Nonwords	40	23	58

for grammatical word class, concreteness, or phoneme–grapheme conversion probability (the likelihood of generating a correct spelling through sublexical phonology-to-orthography conversion processes). The Dysgraphia Battery results also showed no reliable effect of word length. However, in subsequent testing with a larger set of stimuli that manipulated word length while controlling word frequency, CM's spelling accuracy was significantly higher for 3–4 letter words (69%) than for 5–8 letter words (50%), χ^2 (1, $N = 276$) = 10.16, $p < .01$.

One fourth of CM's errors on the Dysgraphia Battery (49/198) were lexical errors, in which he produced the spelling of a word other than the target word (e.g., *solid* → *solar, method* → *mother*). In the vast majority of these errors CM's response bore a clear (although not always close) orthographic relation to the target word (e.g., *cross* → *cough, vivid* → *given*). Only 3 of the lexical errors (6%) could be considered semantic or morphological (e.g., *weave* → *weaver*).

Three fourths of CM's errors (149/198) were nonlexical (e.g., *valley* → *valice, kitchen* → *kithen*). As in the case of the lexical errors, CM's response and the target word were nearly always orthographically related, although not always closely (e.g., *merge* → *mourer*). Only 8 of the 141 errors (4%) could be considered phonologically plausible (e.g., *open* → *opon*), suggesting that CM was not spelling through application of sublexical phonology-to-orthography conversion processes.

To characterise more specifically the orthographic relatedness between target and error responses, the lexical and nonlexical errors were submitted to a computer program that attempted to interpret each error as a letter substitution, deletion, transposition, insertion, or some combination of these error types. (For a description of the program, see McCloskey, Badecker, Goodman-Schulman, & Aliminosa, 1994.) An error was assigned an interpretation if it could be explained by one of the individual error types, or by a combination of two individual errors (e.g., *crisp* → *criosy*, interpreted as an *o* insertion plus a *p*-to-*y* substitution). Errors attributable to deletions or insertions of any number of letters were also interpreted as such (e.g., *palace* → *pa*, interpreted as a deletion of *l, a, c*, and *e*). If an error could not be interpreted in any of these ways, the program classified it as uninterpreted, on grounds that it was too complex for a meaningful interpretation to be assigned.

For the nonlexical errors the program placed 52% (77/149) in the uninterpreted category, reflecting the fact that CM's spellings were often substantially discrepant from the correct spellings (e.g., *future* → *furance, brisk* → *brouches*). Across the 72 interpreted nonlexical errors, CM made 59 letter substitutions (e.g., *ruin* → *roin*), 22 letter insertions (e.g., *myth* → *mytch*), 16 letter deletions (e.g., *journal* → *joural*), and 9 letter transpositions (e.g., *fluid* → *fliud*). (The total adds to more than 72 because some incorrect responses included multiple errors, as in *nuisance* → *nuestance.*)

The pattern was similar for the lexical errors. Sixty-five per cent (32/49) were too complex to be interpreted (e.g., *bring* → *better*). Among the interpreted errors, letter substitutions (18 errors) and insertions (8) were the most common error types, occurring either alone (e.g., *head* → *heat*) or in combination (e.g., *motel* → *mother*). Letter deletions (3) and transpositions (2) were less common.

Writing nonwords to dictation

In writing nonwords to dictation CM was only 12% correct (4/34), a level of performance significantly below the 39% correct he achieved for words, χ^2 (1, $N = 360$) = 8.88, $p < .01$. Nearly all of the errors (27/29) were nonlexical (e.g.,/tibl/→ *trerer*).

Written picture naming

CM was 25% correct (13/51) in the written picture naming task from the Dysgraphia Battery. On 16 trials he produced no response. Among the remaining 22 errors 8 were lexical (e.g., *tent* → *ticket*) and 14 were nonlexical (e.g., *orange* → *origage*). Almost all of the errors (20/22) were too complex for the error analysis program to interpret (e.g., *flag* → *hacket*).

Typing to dictation

The Dysgraphia Battery includes oral spelling tasks, but CM's spoken language production was too impaired for these tasks to be feasible, and typing to dictation tasks were administered as an alternative. A 92-word list was dictated twice (in separate test sessions), and CM spelled each word by typing on a computer keyboard. The list was also presented in three other sessions for writing to dictation. CM's performance was extremely similar in the typing and writing tasks. He was 61% correct (112/184) in typing and 59% (167/276) correct in writing, $\chi^2(1) < 1$. Lexical errors made up 36% of the total errors in typing (e.g., *diamond* → *diameter*), and 30% of the errors in writing (e.g., *pocket* → *rock*). The error-classifying program placed 35% of the typing errors (e.g., *perfume* → *pofuer*) and 39% of the writing errors (e.g., *perfume* → *roefer*) in the uninterpreted category, and among errors that were assigned an interpretation, letter substitutions were the most common error type in both typing (e.g., *pool* → *rool*) and writing (e.g., *woman* → *womar*).

Copy transcoding

In direct copy transcoding CM copied stimuli that remained in view, translating from upper case to lower case, and vice versa. For delayed copy transcoding CM looked at the stimulus for as long as he wished, after which it was covered and he wrote it from memory (again changing case). In the direct copy task he was 100% correct for both words (84/84) and nonwords (40/40). In delayed copying, however, his performance was significantly worse: 70% correct (59/84) for words and 58% (23/40) for nonwords, χ^2 $(1, N = 168) = 27.06$, $p < .01$, and χ^2 $(1, N = 80) = 19.12$, $p < .01$, respectively. Most of his errors (40/42) were nonlexical (e.g., *rather* → *rathen*; *talent* → *talbet*), and appeared similar to the errors observed in the writing to dictation tasks.

Intact and impaired processes

CM performed well on tasks requiring comprehension of spoken words, including the BDAE auditory comprehension tasks, the Peabody Picture Vocabulary Test, and the word–picture matching test. Further, he was able to repeat words successfully in the context of writing to dictation. In one writing-to-dictation task CM was asked to repeat the stimulus word both before and after writing it. On 115 of 120 trials (96%) both repetitions were correct (yet CM misspelled 61 of the 120 words). These results suggest that CM was capable of generating phonological and (in the case of words) semantic representations from dictated stimuli, and hence that his spelling deficit resulted from impairment at some later stage(s) of the spelling process.

CM's perfect performance on the direct copy transcoding task provides evidence that allographic conversion and more peripheral graphomotor processes were intact. Also, his spelling performance—in accuracy and error types—was remarkably similar in writing and typing to dictation, suggesting that his spelling deficit affected the processing of abstract lexeme and/or grapheme representations prior to the application of processes that convert the abstract grapheme representations to the allographic or keystroke representations required for writing and typing, respectively.

CM's spelling errors were predominantly phonologically implausible letter substitutions, insertions, deletions, and more complex misspellings. This error pattern, as well as the word length effect and CM's impaired performance in delayed copy transcoding, suggest a deficit affecting the processing of abstract grapheme representations at the level of the graphemic buffer (i.e., the *Graphemes* level in Figures 1 and 2). Below we offer a more specific characterisation of the deficit.

The marginal effect of word frequency, and CM's substantial number of lexical errors (e.g., *danger* → *dampen*), suggest that the lexeme level of representation may also be implicated in his spelling impairment. Finally, CM's very poor nonword spelling, the phonological implausibility of his spelling errors for words, and the finding that phoneme–grapheme conversion probability did not affect his spelling accuracy, suggest severe impairment of sublexical phonology-to-orthography conversion processes.

Intruded letters and the letter persistence effect

Most of CM's erroneous spelling responses contained what we will call intruded letters—that is, letters that did not match letters from the correct spelling of the stimulus word. For example, in the substitution error *head* → *heat* the *t* is an intruded letter, as is the *i* in the insertion error *spend* → *speind.* We also count as intruded letters any "extra" occurrences in the response of letters that appear in the correct spelling, such as the second *t* in *turkey* → *turget.*

Letters that occurred as intrusions in a response were often present in one or more of the several immediately preceding responses. Consider, for example, the error *value* → *valod*, in which the erroneous response contained two intruded letters, *o* and *d.* Table 2 illustrates the trial on which this error occurred (labelled E) and the five immediately preceding trials (E-1, E-2, . . . , E-5). As can be seen from the table, the letters intruded on trial E were present in many of the immediately preceding responses.

Table 2. *CM's error value* → valod *and the stimuli and responses from the five immediately preceding trials*

Trial	*Stimulus*	*Response*
E-5	drama	drama
E-4	avert	ovent
E-3	provide	provide
E-2	open	opon
E-1	leopard	leopord
E	value	valod

This pattern in CM's errors suggested that the letters occurring as intrusions in a response were often in some sense coming from preceding responses. Perhaps, in particular, grapheme representations activated during production of a response sometimes persisted abnormally in an activated state, rather than being deactivated upon completion of the response. If this were the case, then on any given trial one or more grapheme representations from previous trials might be abnormally active, leading potentially to letter intrusions. For example, the *o* and *d* intrusions in the *value* → *valod* error may have occurred because *o* and *d* grapheme representations were activated on preceding trials, and remained active beyond the point at which the stimulus *value* was presented.[1]

The present study explored the apparent letter persistence effect and its implications more thoroughly. In the following discussion we first demonstrate that the persistence effect is real by showing that the intruded letters in CM's responses were present in preceding responses far more often than expected by chance. Next, we present several results in support of the claim that CM's letter persistence effect arises at an abstract graphemic level of representation. With these initial conclusions as a foundation we then use the letter persistence effect to argue that the cognitive spelling system has an interactive architecture that incorporates grapheme-to-lexeme feedback.

EXPERIMENTAL STUDY

The letter persistence effect was explored primarily through writing-to-dictation tasks. Over a

---

[1] Cohen and Dehaene (1998) have suggested that perseverations may arise from normal persistence of activation, when a level of representation receives weaker-than-normal input from other levels. From this perspective, CM's letter intrusions might be interpreted by assuming not that grapheme activation persists to an abnormal extent, but rather that signals from the lexeme level to the grapheme level are abnormally weak, with the result that grapheme activation persisting normally from prior responses has a greater-than-normal influence on grapheme selection for the current response. However, the distinction between abnormally strong grapheme persistence and abnormally weak input from the lexeme level is not relevant for our purposes in this paper. The relevant assumption, made by both interpretations, is that persisting grapheme activation from prior responses is abnormally strong relative to input from the lexeme level. For convenience, we will therefore continue to describe CM's deficit as one of abnormally persisting grapheme activation.

5-year period, during which CM's performance remained stable, he wrote 3797 words to dictation (including the words from the Dysgraphia Battery and the other spelling tasks discussed above). Across the entire stimulus set his accuracy was 55% (2101/3797).

On 9 of the 1696 error trials CM produced no response. A further 86 errors were whole-response perseverations, in which CM's response matched an entire response or stimulus from earlier in the current list (e.g., *wagon → watch*, occurring 4 trials after he had spelled *watch* correctly). The remaining 1601 errors were nonperseverative misspellings, and these errors made up the corpus for subsequent analyses.[2] Within this corpus, 538 (34%) of the errors were lexical (e.g., *chin → chance*) and 1063 (66%) were nonlexical (e.g., *shatter → stach*).

The 1601 errors collectively included 3174 intruded letters. Among the 1483 errors that included letter intrusions, 574 contained 1 intruded letter (e.g., *sort → soft*), 432 errors had 2 intruded letters (e.g., *between → brethen*), and the remaining 477 had 3 or more intrusions (e.g., *bulb → bumber*).

Above-chance letter persistence

To determine whether CM's spelling errors showed a real letter persistence effect, we assessed whether intruded letters were present in preceding responses more often than expected by chance. For each occurrence of a letter intrusion we first scored whether the intruded letter was present in CM's response on each of the five immediately preceding trials. (Errors occurring in the first five trials of each test list were excluded from the analysis, because for these errors there were fewer than five preceding trials.) For instance, in the case of the *value → valod* error for the trial designated E in Table 2, CM's response on each of the five preceding trials (E-1 through E-5) was scored for the presence of an *o* and a *d* (the letters intruded on trial E). On trial E-1 (*leopard → leopord*) and E-3 (*provide → provide*) both *o* and *d* were present in the response, whereas on the remaining trials only one of the two intruded letters was present: *o* on trials E-2 (*open → opon*) and E-4 (*avert → ovent*), and *d* on trial E-5 (*drama → drama*).

A computer program carried out this tabulation for all of the letter intrusions, with the results shown in the upper curve of Figure 3. The .453 proportion shown for trial E-1 reflects the finding that 1342 of 2960 intruded letters (45.3%) were present in CM's response on the immediately preceding trial (E-1).[3] Similarly, the .344 proportion shown for trial E-5 indicates that 34.4% of the intruded letters were present in CM's response on the trial five back from that on which the intrusion occurred. It is evident from the figure that the likelihood of an intruded letter being present in a preceding response decreased monotonically with the number of trials intervening between the letter intrusion and the preceding response.

The computer program also estimated the likelihood of an intruded letter being present in a preceding response by chance. The rationale was as follows: If the occurrence of a letter intrusion is completely unrelated to the presence of the intruded letter in the immediately preceding responses, then the intruded letter should be just as likely to be present in responses that do not immediately precede the intrusion. For example, if the *d* intrusion on trial E in Table 2

---

[2] Even though the letters in whole-word perseverative responses are by definition present in preceding responses, these letters were excluded from tabulations of letter intrusions in letter persistence analyses, because the processes leading to whole-response perseverations may not be the same as those leading to persistence of individual letters. This conservative decision avoids the possibility of inflating estimates of the magnitude of the letter persistence effect by including errors that may not involve persistence of individual letters.

[3] The number of intrusion errors in this analysis was 2960 rather than 3174 (the total number of intrusions) largely because errors occurring in the first five trials of each test list were excluded, as described above. A small number of additional errors were excluded because for these errors no control items were available for the analysis described below, which estimated the probability of an intruded letter being present by chance in preceding responses.

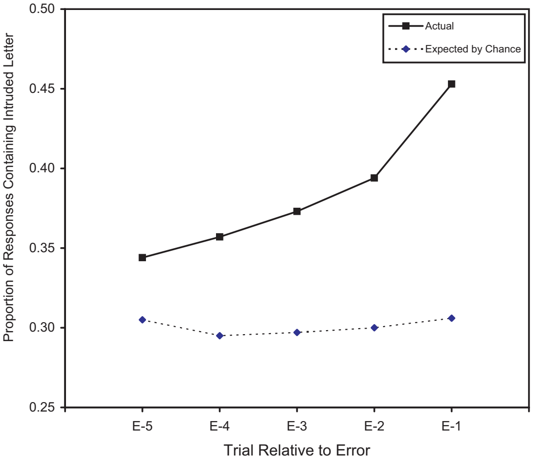

Figure 3. *The proportion of trials in which a letter intruded on trial E was present in the response on each of the five immediately preceding trials (E-1 through E-5), and the proportion expected by chance.*

(*value* → *valod*) had nothing to do with whether *d*s were present in the responses on trials E-1 through E-5, then *d*s should be just as likely to be present in responses that did not immediately precede the *valod* response—for example, in the response on trial E-20 or E+20. Put another way, tabulating the likelihood of a *d* being present in responses not proximal to trial E provides an estimate of the likelihood of a *d* being present by chance on the trials immediately preceding trial E.

The chance estimates were obtained by tabulating, for each letter intrusion, the presence and absence of the intruded letter in responses on trials not in the range E-5 through E+5. For example, given a 70-word list and an intrusion occurring on trial 50, chance estimates were obtained for the intrusion by considering responses on trials 1−44, and 56−70. We refer to the trials used for generating chance estimates as the control trials. (Trials E-5 through E+5 were excluded on grounds that the presence of the intruded letter on these trials could have been nonaccidentally related to the occurrence of the letter intrusion. The presence of the intruded letter in responses on trials E-1 through E-5 might have influenced the occurrence of the intrusion; and the presence of the intruded letter on trial E could have led to its occurrence in responses on trials E+1 through E+5.)

The specific procedure for generating chance estimates may be illustrated by referring once again to the *d*-intrusion in *value* → *valod*. The likelihood of a *d* being present by chance on trial E-1 was assessed by first identifying all control-trial responses of the same length as the response on trial E-1. Given that CM's response on the E-1 trial (*leopord*) had 7 letters, all 7-letter control responses were identified. The analysis program then randomly selected one of these control responses, and determined whether or not it contained a *d*. The same operations were then performed for trials E-2 through E-5. For example, the response on trial E-2 (*opon*) was 4 letters in length, and so the analysis identified all 4-letter control responses, randomly selected one of these, and determined whether it contained a *d*.

This procedure was carried out for all of CM's letter intrusions. Thus, for each of the 2960 E-1 responses, a control response of the same length was selected and checked for the intruded letter; and similarly for each of the E-2, E-3, E-4, and E-5 responses. For each of the five sets of control responses the proportion that contained the intruded letter was calculated. For example, on one run of the analysis 924 of the 2960 control responses for trial E-1 contained the intruded letter, for a proportion of .312. This value represents an estimate of the proportion of E-1 responses that would be expected by chance to contain the intruded letter.

The entire process of generating estimated chance proportions for trials E-1 through E-5 was carried out repeatedly. Because the control responses entering into the estimates were selected randomly from the control-trial responses with the appropriate length, each repetition of the analysis produced a new sample of control responses, and hence a new estimated chance proportion. A total of 1,000,000 repetitions were carried out, yielding a distribution of estimated chance proportions for each of the trials E-1 through E-5.[4]

The distribution for trial E-5 is presented as an example in Figure 4. The mean of the distribution was .305, and this value was taken as the proportion of E-5 trials expected by chance to include the intruded letter. The distribution also provided a basis for assessing the statistical reliability of the difference between the observed proportion of E-5 responses that contained the intruded letter (.344) and the proportion expected by chance (.305). As the figure illustrates, none of the estimated chance proportions obtained in 1,000,000 repetitions of the analysis was as high as the observed proportion of .344 (shown as the tall vertical bar on the far right side of the figure), indicating that the probability of obtaining the observed proportion by chance was less than one in a million ($p < .000001$).

The estimated chance proportion (i.e., the mean of the frequency distribution) for each of the trials E-1 through E-5 is shown in the lower curve of Figure 3. For each of these trials the observed proportion of responses containing the intruded letter was reliably higher than the proportion expected by chance (all $p < .000001$). Accordingly, we conclude that letters occurring as intrusions in CM's responses were present in the several immediately preceding responses at rates considerably greater than chance, and hence that his spelling performance shows a letter persistence effect.

In discussing the results of subsequent analyses we report the magnitude of the persistence effects in terms of the difference between observed and chance proportions, rather than presenting the observed and chance values separately. For example, in the analysis of all intrusions, the magnitude of the persistence effect for trial E-1 was calculated as .453 (observed proportion) minus .306 (proportion expected by chance), yielding .147. Figure 5 plots the magnitude of the persistence effect for trials E-1 through E-5.

For some analyses we will report the magnitude of the persistence effect as a single value, calculated as the average of the effects for trials E-1 through E-5. (For these analyses we generated a distribution of expected average chance proportions, for use in significance testing.) In the analysis of all intrusions the average magnitude of the persistence effect was .084 ($p < .000001$). This value reflects the fact that when results were averaged over trials E-1 through E-5, the observed proportion of responses containing the intruded letter was .384, whereas the average proportion expected by chance was .300.

Our method for determining that CM's intruded letters were present in preceding responses more often than expected by chance is very similar in rationale to a method described by Cohen and Dehaene (1998) for evaluating whether perseverations occur at above-chance rates. However, our specific procedures for estimating rates expected by chance are somewhat different from those of Cohen and Dehaene, the most notable difference being that our criteria for selection of control responses are more stringent and therefore may yield more accurate estimates of chance rates.

Functional locus of the letter persistence effect

In this section we present several results and analyses aimed at clarifying the level(s) of representation giving rise to CM's letter persistence effect.

Orthographic or phonological persistence

In the preceding discussion we have assumed that the persisting representations were grapheme

---

[4] The mean number of available control items for a source item was 20.3. On each repetition of the analysis for each source position (E-1, E-2, ...) a control item was sampled for each of the approximately 3000 intrusion errors. As a consequence the number of possible control item samples was astronomically large: An average of approximately 20 options were available for each of approximately 3000 choices. (An attempt to calculate the exact number of possible samples in Excel failed due to overflow errors when the value reached approximately 4×10^{307}.) Hence, in carrying out 1,000,000 repetitions of the analysis we were not simply generating the same control-item samples repeatedly.

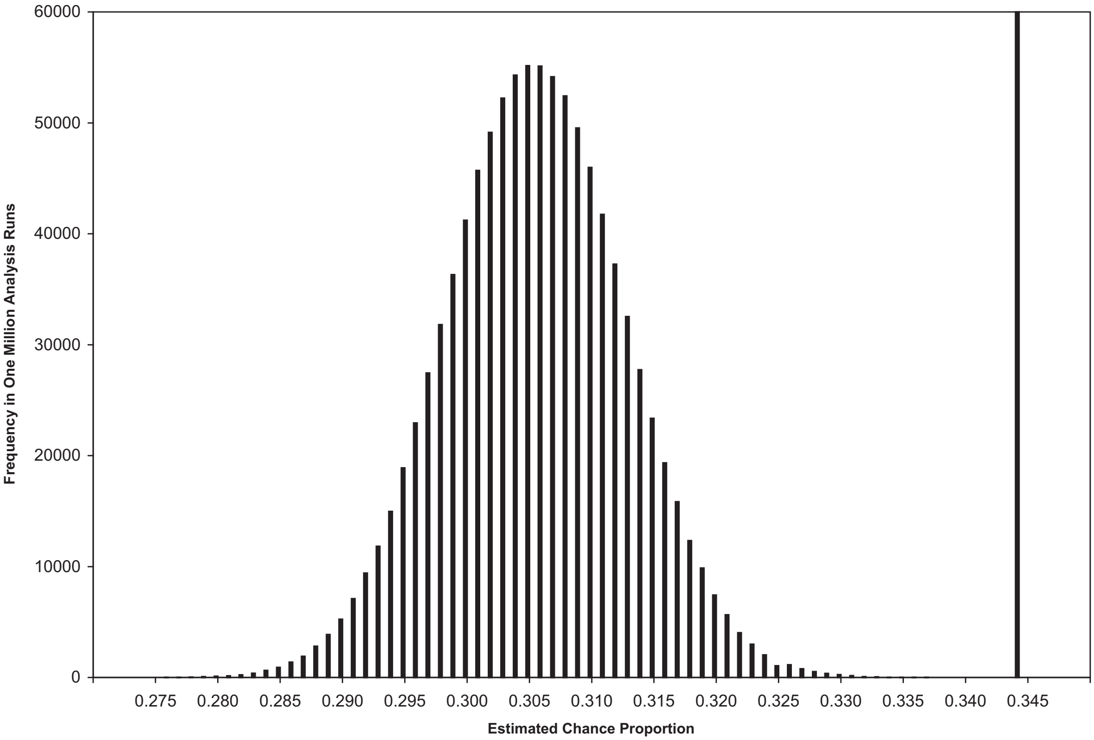

Figure 4. *Frequency distribution for proportion of E-5 trials expected by chance to contain the intruded letter. Vertical bar at far right indicates the observed proportion.*

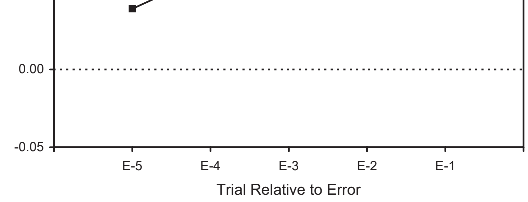

Figure 5. *Letter persistence effect (difference between actual and chance proportions of preceding responses containing the intruded letter) for each of the five trials preceding the letter intrusion.*

representations activated in the course of producing responses. However, an alternative possibility is that the persistence effect stems from abnormal persistence of activation among phoneme representations that became active when stimulus words were dictated. For example, the *d*-intrusion in the *value* → *valod* error (see Table 2) may have occurred not because a representation of the letter *d* remained active after the production of the response on the preceding trial (*leopard* → *leopord*), but rather because a representation of the phoneme /d/ remained abnormally active, leading eventually to activation of the grapheme *d* via lexical or sublexical processes.

To explore this possibility we presented CM with two lists for writing to dictation. Each list contained 24 sets of three consecutive words. In each three-word set the first word was a prime that contained the phoneme /f/, and the next two words were targets that did not contain /f/. In the Prime-with-F list the /f/ in each prime word was spelled with the letter *f* (e.g., *afraid*), whereas in the Prime-without-F list the /f/ was spelled with other letters (e.g., *sphere*). The target words were identical in both lists. For example, as illustrated in Table 3, the Prime-with-F list contained the prime-target-target sequence *afraid-direct-sudden*, whereas the Prime-without-F list contained *sphere-direct-sudden.* Prime words were matched in frequency and length across the two lists.

The question of interest was how often the letter *f* would occur as an intruded letter in CM's responses to the target words. If the persistence effect were phonologically based, we would expect *f* intrusions to occur equally often in the Prime-with-F and Prime-without-F lists, because the prime words in both lists contained the phoneme /f/. In contrast, if the persistence effect were due to persisting activation of grapheme representations, we would expect *f* intrusions to be more frequent in the Prime-with-F list than in the Prime-without-F list.

The results were unequivocal. For the Prime-with-F list CM's response to 12 of the 48 target words contained the letter *f*, but for the Prime-without-F list CM produced no *f* intrusions in spelling the target words, χ^2 (1, $N = 96$) = 11.52, $p < .001$. Table 3 illustrates this pattern of results, showing that CM intruded *f*s into his spellings for the target words *direct* and *sudden* when the prime word was *afraid*, but not when the prime was *sphere.* We conclude that CM's persistence effect arises from abnormally persisting activation of grapheme and not phoneme representations.

Table 3. *Example of a trial sequence from the Prime-with-F and Prime-without-F lists*

Trial sequence	*Prime-with-F list*		*Prime-without-F list*	
	Stimulus	*Response*	*Stimulus*	*Response*
Prime	afraid	aflersuil	sphere	sphere
Target	direct	diffent	direct	diriction
Target	sudden	suffent	sudden	suddion

Stimulus- or response-based persistence

Further support for the interpretation of persisting grapheme representations came from analyses carried out over the full set of letter intrusions. Two predictions of the interpretation were tested: First, letters appearing in the correct spelling of a stimulus word but not in CM's response should show little or no tendency to intrude into subsequent responses. For example, in the error *dial* → *dian* the correct spelling of the stimulus word contains the letter *l*, but this letter did not appear in CM's response. If we assume that CM failed to produce the *l* because an internal representation of this letter did not become sufficiently activated, then we would not expect persisting activation of the letter representation to be strong enough to cause *l* intrusions in subsequent responses.

The second prediction was that letters appearing in CM's response but not in the correct spelling should intrude into subsequent responses at above-chance rates. For example, the presence of the *n* in the *dial* → *dian* response implies that a representation of the letter *n* became activated during the production of the response. Persisting activation of this letter representation could then lead to *n* intrusions on subsequent trials.

Note that on a persisting-phonemes interpretation we might expect the opposite pattern of results. Given that CM was able to generate accurate phonological representations of dictated stimulus words, the persisting phonemes should be those present in the stimulus words, and as a consequence persistence effects should be stimulus-based.

Two analyses were carried out to evaluate the contrasting predictions. The stimulus-not-response analysis tabulated, for each letter intrusion, whether the intruded letter was present on each of the five preceding trials among letters that appeared in the correct spelling of the stimulus word but not in CM's response. Consider the following sequence of trials:

Trial	*Stimulus*	*Response*
E-1	chap	chorlon
E	thick	thirk

In making the tabulations for the *r* intrusion on trial E, the analysis considered for trial E-1 whether an *r* was present in the set of letters {*a, p*}. Chance probabilities were determined by calculating the likelihood of the intruded letter being present on control trials among letters present in the correct spelling of the stimulus word but not in CM's response.

The response-not-stimulus analysis was the same, except that for trials E-1 through E-5, and the corresponding control trials, the letters considered were those that appeared in CM's response but not in the correct spelling of the stimulus word. Thus, for the *r*-intrusion in trial E above the analysis considered for trial E-1 (*chap* → *chorlon*) whether an *r* was present in the set of letters {*o, r, l, n*}.

The results, presented in Figure 6, were entirely as predicted by the interpretation of persisting letter representations. The Stimulus-not-Response analysis revealed no hint of a letter persistence effect. In the Response-not-Stimulus analysis, however, the likelihood of an intruded letter being present among letters that appeared in CM's response but not in the correct spelling was reliably greater than chance for each of the trials E-1 through E-5 (all $p < .0001$).[5]

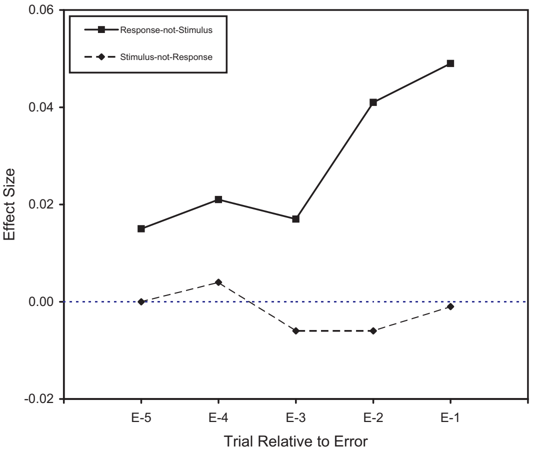

Figure 6. *Letter persistence effect for the Response-not-Stimulus and Stimulus-not-Response analyses.*

Visual feedback as the cause of the persistence effect?

We have assumed that the letter representations subject to abnormally persisting activation were localised within the cognitive spelling system, and were activated by processes intrinsic to the production of spelling responses. In this section we consider a potential alternative interpretation—that visual feedback occurring as CM looked at what he was writing activated visual (or more abstract) letter representations specific to reading, and these representations persisted abnormally in an activated state. Perhaps, for example, when CM looked at his written response *leopord*, reading-specific representations for *d, o*, and other letters became activated, and remained activated during subsequent spelling trials (even though each written response was covered before the next trial began). The reading-specific letter representations may then have directly or indirectly (e.g., via lexeme representations) activated letter representations implicated in spelling, leading to letter intrusions on the subsequent trials (e.g., *value* → *valod*).

To evaluate this possibility we presented CM with two 70-word lists for writing to dictation

---

[5] The magnitude of the persistence effect (i.e., the difference between observed and chance rates) was somewhat smaller in the Response-not-Stimulus analysis than in the original analysis that considered all letters in CM's responses. This result is exactly as expected, given that some of CM's letter intrusions were presumably caused by persisting activation of representations for letters that were present on preceding trials in both the correct spelling and in CM's response. For these intrusions the Response-not-Stimulus analysis would not have found the intruded letter among letters present in the response but not in the stimulus on the preceding trial. As a result, we would expect the persistence effect to be smaller in the Response-not-Stimulus analysis than in the main analysis considering all letters in preceding responses.

under each of two conditions. In the eyes open condition the normal writing-to-dictation procedures were followed: CM was allowed to look at what he was writing (and each response was covered before the next stimulus was dictated). In the eyes closed condition, however, CM kept his eyes closed throughout administration of the stimulus lists, so that he received no visual feedback from his written responses. If the letter persistence effect were due to activation of reading-specific letter representations as CM looked at what he was writing, then little or no persistence effect should be observed in the eyes closed condition.

CM's spelling accuracy was comparable in the two conditions: 41% (57/140) in the eyes closed condition, and 50% (70/140) in the eyes open condition, χ^2 (1, $N = 280$) = 2.44, $p > .05$. Table 4 presents the magnitude of the letter persistence effect for trials E-1 through E-5 in each of the two conditions. A highly reliable persistence effect was obtained in the eyes closed condition ($p < .0001$). If anything, the effect was somewhat larger in this condition than in the eyes open condition. We conclude that the letter persistence effect was not caused by visual feedback occurring as CM looked at what he was writing.

Persistence of abstract graphemic or more peripheral representations

The results presented thus far indicate that CM's persistence effect arises from abnormally persisting activation of letter representations in the cognitive spelling system. In this section we consider whether the persisting letter representations were abstract grapheme representations—that is, abstract representations of letter identity—or more peripheral representations, such as representations of allograph shapes (e.g., the shape of a lower-case cursive *g*) or even motor plans for production of writing movements.

Table 4. *Size of the letter persistence effect in the eyes open and eyes closed conditions*

	Condition	
Trial	*Eyes closed*	*Eyes open*
E-5	−.044	−.016
E-4	.048	.013
E-3	.014	.061
E-2	.142	.006
E-1	.153	.080

The issue was explored with a 60-trial writing-to-dictation task in which CM wrote his response in upper case (e.g., *BEAN*) on odd-numbered trials, and lower case (e.g., *fat*) on even-numbered trials. The rationale was as follows: If the letter persistence effect arises at an abstract graphemic level of representation, then intruded letters should be present at above-chance rates in the responses on all of the immediately preceding trials (E-1 through E-5), independent of whether the preceding response was written in the same case as the intruded letter, or in the alternative case. Consider the following sequence of trials:

Trial	*Stimulus*	*Response*
E-5	door	DOOR
E-4	fad	lage
E-3	toil	TOLE
E-2	doom	dome
E-1	dim	DIMOT
E	mob	mot

For the lower-case *t* intrusion on trial E, we would expect the letter *t* to be present with above-chance likelihood not only in preceding lower-case responses (trials E-2 and E-4), but also in preceding upper-case responses (trials E-1, E-3, and E-5). Put another way, production of upper-case *T*s as well as lower-case *t*s on trials in the E-1 through E-5 range should influence the occurrence of lower-case *t* intrusions on trial E.

On the other hand, if the persistence effect arises at a level of representation more peripheral than that of abstract graphemes (e.g., a level of shape-specific allograph representations) then intruded letters should more often be present on preceding trials with responses in the same case, than on trials with responses in the alternative case. For the intruded lower-case *t* in *mob* → *mot*, we would expect the letter *t* to be present with above-chance likelihood on preceding lower-case trials (E-2, E-4), but not (or to a lesser extent) on preceding upper-case trials (E-1, E-3, E-5).

The results are presented in Table 5. As usual, the letter persistence effect was largest for trial E-1, even though the E-1 response (e.g., *DIMOT*) was written in a different case to the E response (e.g., *mot*). The persistence effect was reliable for trial E-1 ($p < .01$), as well as for the three different-case trials (E-1, E-3, E-5) taken together ($p < .01$). Thus, letters written in one case intruded at above-chance rates into subsequent responses written in the alternative case. Furthermore, the data do not give any indication that the persistence effect was diminished for different-case trials relative to same-case trials. These results strongly suggest that the letter persistence effect has its locus in a level of representation shared by upper- and lower-case writing, and hence that the letter representations giving rise to the effect are abstract grapheme representations.

Another relevant finding is that CM showed the letter persistence effect not only in writing, but also in typing: For the typing to dictation task described above (see pp. 10-11) the mean effect size across trials E-1 through E-5—difference between observed and chance likelihoods of an intruded letter being present in a preceding response—was .085 ($p < .000001$), virtually identical to the effect of .084 observed in the analysis of CM's writing to dictation performance. This result suggests that the letter persistence effect arises at a level of letter representation shared by writing and typing, and hence that the persisting representations are abstract grapheme representations.

Table 5. *Letter persistence effect in the alternating-case writing-to-dictation task*

Trial	*Case relative to trial E*	*Size of persistence effect*
E-5	Different	.011
E-4	Same	.096
E-3	Different	.069
E-2	Same	.060
E-1	Different	.138

Having presented evidence that the letter persistence effect is real, and that the effect arises at an abstract graphemic level of representation, we next use the persistence effect as part of the basis for arguing that the cognitive spelling system has an interactive architecture.

LEXICAL ERRORS, THE PERSISTENCE EFFECT, AND INTERACTIVITY

As we have noted, 34% of CM's misspellings (538/1601) were lexical errors (e.g., *dignify* → *define*, *arm* → *amber*, *carpet* → *compute*). Such errors could arise systematically, as a result of activating the wrong word representation at the lexeme level. For example, the *dignify* → *define* error may have occurred when the incorrect lexeme *DEFINE* became activated, leading to activation of the grapheme sequence *D-E-F-I-N-E* at the grapheme level. However, lexical errors could also arise by chance, for reasons having nothing to do with the fact that the incorrect spelling happens to be a word—that is, without the lexeme representation of the incorrect word playing any role in the activation of the incorrect graphemes. For example, the error *bean* → *bead* might have occurred when the grapheme D was activated instead of the grapheme N, without the *BEAD* lexeme representation playing any role in the activation of that incorrect grapheme. We will refer to lexical errors resulting at least in part from activation of an incorrect lexeme representation as *true lexical errors*, and to lexical errors occurring wholly for other reasons as *chance lexical errors*.

The distinction between true and chance lexical errors is important for an argument we develop in the following discussion. Specifically, we argue that CM's lexical errors, in conjunction with the letter persistence effect, motivate the postulation of grapheme-to-lexeme feedback in the spelling system. The argument has three major steps:

1. At least some of CM's lexical errors are true lexical errors.
2. The true lexical errors show the letter persistence effect.
3. The presence of the persistence effect in the true lexical errors implies grapheme-to-lexeme feedback.

Establishing the first two steps of the argument, and demonstrating that the third step then follows, are far from trivial exercises, and will require substantial discussion.

Step 1: Some observed lexical errors are true lexical errors

The first task is to demonstrate that at least some of CM's lexical errors are true lexical errors, resulting from activation of incorrect word representations at the lexeme level. To establish this point we need to show that CM made lexical errors more often than expected by chance, so that at least some of the errors must be true lexical errors.

How, though, can we estimate the rate at which lexical errors should occur by chance, in order to show that CM's rate of lexical errors exceeds the chance rate? If we consider CM's entire corpus of errors, the problem appears intractable. Given that CM's erroneous spellings were often grossly different from the correct spellings (e.g., *umbrella → ublerrer, string → slint, rope → shop*), it is entirely unclear how we could estimate the rate at which his misspellings should happen by chance to come out as words.

Faced with this difficulty, we decided to limit our analyses to a subset of CM's misspellings for which estimation of the chance rate of lexical errors appeared tractable (although not simple). In particular, we limited the analysis to errors taking the form of letter substitutions. More specifically still, we considered those errors in which the incorrect spelling differed from the correct spelling only in that one or two letters of the correct spelling were replaced by incorrect letters; examples are shown in Table 6. The aim of the analysis was to determine whether lexical errors occurred more often than expected by chance among the one- and two-letter substitutions.

Restricting the analysis to one- and two-letter substitution errors is a very conservative strategy, because these errors are certainly among those most likely to come out as words by chance. For example, *bell → ball* could presumably occur for reasons having nothing to do with the fact that *ball* is a word, far more readily than *carpet → compute* could occur for reasons unrelated to the fact that *compute* is a word. In limiting our analysis to the simple substitution errors, we are therefore stacking the deck against ourselves, making it difficult to show that lexical errors occurred more often than expected by chance. (Romani et al., 2002, took a similar approach in assessing whether lexical errors occurred at an above-chance rate in their patient DW.)

Among CM's total of 1601 errors, 469 took the form of one- or two-letter substitutions (including both lexical errors such as *bell → ball*, and nonlexical errors such as *mate → mape*). Twenty of these errors were excluded from the analysis because they occurred too early in a stimulus list to be considered in persistence effect computations (required for a later step in the argument for grapheme-to-lexeme feedback). Of the remaining 449 errors, 191 were lexical errors and 258 were nonlexical errors. The question we attempted to answer was whether the observed proportion of lexical errors (.43) was

Table 6. *Examples of CM's one- and two-letter substitution errors*

Lexical errors		*Nonlexical errors*	
Stimulus	*Response*	*Stimulus*	*Response*
sold	cold	jeep	jeek
square	squire	hint	rint
log	bag	fled	flug
blind	blank	water	waton

reliably greater than the proportion expected by chance.

The methods we developed for answering this question were similar, although not identical, to procedures previously used by other researchers for assessing whether lexical errors occurred more often than expected by chance in spoken production (e.g., Best, 1996; Martin et al., 1996) and writing (e.g., Romani et al., 2002). Our strategy was to consider what letter substitutions might have occurred by chance in each of the 449 errors, and then to determine from this set of candidate substitutions what proportion of lexical errors would be expected by chance. For example, in the error *take* → *make* an *m* was substituted for the initial *t*, yielding a lexical error. However, other letters could conceivably have been substituted for the *t*, some of which would have yielded lexical errors (e.g., *bake*), and others of which would not (e.g., *pake*).

For the *take* → *make* error we decided (through methods described below) that in addition to *m*—the letter that actually occurred as a substitution—the letters *b, c, d, f, n*, and *p* could also potentially have been substituted for the *t*, so that the potential substitution errors were *bake, cake, dake, fake, make, nake*, and *pake*. Four of these seven potential errors are words; accordingly, we estimated the probability that a substitution for the letter *t* would produce a word by chance as 4/7, or .57.

Obviously, the validity of such estimates hinges on the method for deciding what letters to consider as possible substitutions. Clearly it will not do simply to assume that any letter of the alphabet could have been substituted. For example, some letters (e.g., low-frequency letters, such as *q* or *z*) may be very unlikely to occur as substitution errors.

For each actual substitution error we decided to consider as candidate letters for substitution any letter that occurred in the responses to the five immediately preceding stimuli, plus any letter that occurred as a substitution error in the response under consideration.

Consider, for example, the error *list* → *lift*, in which an *f* was substituted for the *s*. This error, and the stimuli and responses from the five preceding trials, are shown below:

Trial	*Stimulus*	*Response*
E-5	beef	beef
E-4	felt	felt
E-3	mood	moof
E-2	dirt	dirt
E-1	wall	wall
E	list	lift

The preceding responses contain the letters *a, b, d, e, f, i, l, m, o, r, t*, and *w*. These letters were taken as the preliminary candidates for substitution errors. (The letter *f* would have been considered a candidate substitution even if it had not occurred in the five preceding responses, because it was the letter actually substituted.)

This method of selecting candidate substitutions directly takes into account the possibility of letter persistence at the grapheme level; the letters most likely to occur as persistence errors are treated as candidates for substitution. The selection of letters occurring in preceding responses also takes into account letter frequency in CM's responses—letters that CM rarely produced are unlikely to be treated as substitution candidates.

The selection method may not sample all of the letters that might have occurred by chance as substitution errors. However, what is important is not that all possible substitutions be considered, but rather that the candidates considered be representative of the full set with respect to the likelihood of producing a word. As far as we can see, the method of taking candidates from previous responses is unlikely to introduce systematic biases.

After generating the preliminary set of candidate letters for substitutions, we imposed an additional constraint, corresponding to a regularity apparent in CM's errors. When CM made substitution errors, he nearly always substituted consonants for consonants, and vowels for vowels (e.g., *s* for *c* in *clash* → *slash*, and *o* for *u* in *drum* → *drom*). (This phenomenon has also been observed in other dysgraphic patients; see, e.g., Caramazza & Miceli, 1990; McCloskey et al., 1994.) Therefore, whatever the processes leading to CM's substitution errors, these processes did not freely substitute vowels for consonants, or vice versa. Accordingly, we considered as candidate letters for substitution only those having the same consonant-vowel status as the letter CM actually substituted. For example, in the *list* → *lift* error, the letter CM substituted (*f*) is a consonant, and therefore only consonants were allowed in the candidate substitution set. Eliminating the vowels from the initial set {*a, b, d, e, f, i, l, m, o, r, t, w*} yields {*b, d, f, l, m, r, t*}.

After developing the candidate letter set, we generated candidate substitution errors by inserting each candidate letter into the word in place of the letter CM actually substituted. Thus, for the *list* → *lift* error, with candidate letter set {*b, d, f, l, m, r, t*} we generated the candidate substitution errors *libt, lidt, lift, lilt, limt, lirt*, and *litt.*

For errors in which two different letters were substituted, we developed a candidate letter set for each substituted letter, and generated a candidate substitution error for all combinations of letter candidates. In the case of the error *desk* → *task*, for example, the set of candidate letters for substitution into the *d*'s position was {*b, c, g, h, n, r, s, t*} whereas for the *e*'s position the candidate set was {*a, i, o, u*}. As a consequence, there were 32 candidate substitution errors, each representing 1 of the combinations of the 8 candidates for the *d*'s position and the 4 candidates for the *e*'s position: *bask, bisk, bosk*, . . . , *tosk, tusk.*

Many of the candidate errors generated in this way were orthographically unacceptable letter sequences, such as *libt* or *tlash.* However, CM's actual errors, whether words or nonwords, were almost always orthographically acceptable sequences of letters, such as *toble* or *dorn.* This result suggested that orthographically unacceptable sequences should not be considered potential errors. Accordingly, we eliminated such sequences from the lists of candidate errors.[6] For the *list* → *lift* error, *libt, lidt, limt*, and *litt* were eliminated, leaving *lift, lilt, lirt.* Across the entire set of 449 one- and two-letter substitution errors, 3431 of the 8303 candidate substitution errors (41%) were eliminated as orthographically unacceptable.

Imposing an orthographic acceptability constraint is a conservative measure, in that the constraint works against our attempts to demonstrate that CM's proportion of lexical errors is greater than expected by chance. Orthographically unacceptable letter sequences are, of course, quite unlikely to be words, and eliminating these sequences therefore increases the proportion of words in the set of possible substitution errors.

The procedures described thus far generated, for each of CM's substitution errors, a set of errors (including the actual error) that could plausibly have occurred via processes that were entirely insensitive to the word vs. nonword status of letter strings (such as processes whereby incorrect grapheme representations become activated by means other than signals from incorrect lexeme representations).

For each of CM's substitution errors, we then calculated the proportion of candidate errors that were English words. For the *list* → *lift* error, two of the three potential errors—*lift* and *lilt*—were words, whereas the third—*lirt*—was not. The proportion of words in the candidate set was therefore .67.

The likelihood of obtaining a word by chance was estimated in this way for each of the 449 substitution errors made by CM, including not only the lexical errors like *list* → *lift*, but also the nonlexical errors like *half* → *halp.* (For this latter error, 2 of the 6 candidate errors—*hall, halt*—were words, yielding .33 as the estimated probability of obtaining a word by chance in a letter substitution at word-final position.) Across all 449 substitution errors, the mean likelihood of obtaining a word by chance was .226.

---

[6] Orthographic acceptability judgments were made by a computer program that assessed whether each candidate substitution error could be parsed into a sequence of one or more orthographically acceptable syllables. Acceptable syllables were those with an orthographically acceptable nucleus, and (optionally) an acceptable onset and coda.

Estimating the number of lexical errors expected by chance

Using the chance probabilities, we conducted Monte Carlo simulations to estimate the total number of lexical errors expected by chance in the set of 449 substitution errors. The simulation program assumed that all of the errors occurred via processes insensitive to the lexical status of letter strings. On each run the program decided probabilistically, for each of the 449 errors, whether the error had come out by chance as a word or a nonword. Consider, for example, the *half* → *halp* error. As described above, we estimated the probability of a word being produced by chance through a word-final letter substitution to be .33. Accordingly, on each run of the simulation program this error had a .33 probability of coming out as a lexical error, and a .67 probability of coming out as a nonlexical error. (The determination was made each time by generating a random number between 0 and 1, and counting the error as lexical if the generated number was less than or equal to .33.)

On each run the simulation program carried out this process for each of the 449 errors, and then tallied the number of errors that came out by chance as lexical errors. Because of the random element in each lexical/nonlexical determination, the total number of simulated lexical errors was different on different runs of the simulation. By running the simulation multiple times we generated a distribution of numbers of lexical errors expected by chance. More specifically, we ran the simulation one million times, and constructed from these runs the frequency distribution presented in Figure 7. The mean of the distribution was 101.7, indicating that whereas the actual number of lexical errors was 191, the mean number expected by chance in the set of 449 substitutions was about 102. Ninety-five percent of the simulated lexical error frequencies fell below 115, indicating that the probability of obtaining 115 or more lexical errors by chance was less than .05; 99% of the simulated values were below 120, indicating a probability of less than .01 of obtaining 120 or more lexical errors

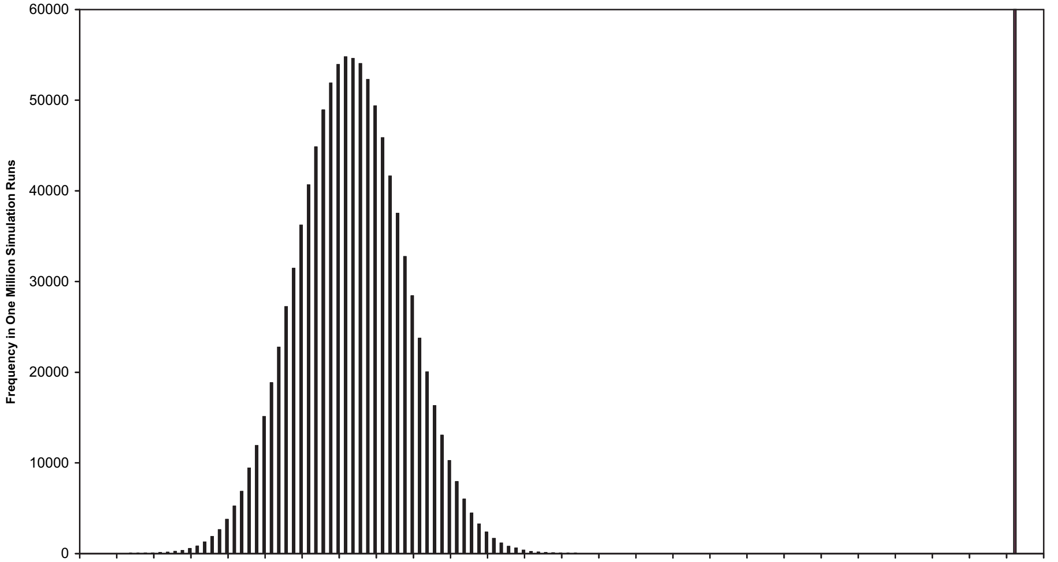

Figure 7. *Frequency distribution for number of lexical errors expected by chance among 449 one- and two-letter substitution errors.*

by chance. The highest number of simulated chance lexical errors obtained in one million runs of the simulation was 140, indicating that the observed number of 191 had a probability (far) lower than one in a million of occurring by chance.

This result provides a basis for concluding that among CM's one- and two-letter substitution errors, the proportion of lexical errors was reliably higher than expected by chance, and therefore that at least some of the lexical errors were true lexical errors, resulting from activation of the wrong lexical representation at the lexeme level.

Estimating the number of true lexical errors

We can use the results of the preceding analysis not only to conclude that the observed lexical errors include some true lexical errors, but also to estimate the number of true lexical errors. It might appear that we could simply take 191, the number of observed lexical errors, and subtract 102, our best estimate of the number of lexical errors expected by chance, to yield 89 as the estimated number of true lexical errors. However, this procedure would be erroneous. The analysis that yielded 102 as the expected number of chance lexical errors was aimed at determining how many lexical errors would be expected *if the error corpus included no true lexical errors.* Accordingly, the analysis assumed that all of the 449 substitution errors were grapheme-level errors (i.e., errors not resulting from activation of an incorrect lexeme representation), and hence that each of the errors had an opportunity to come out as a lexical error by chance.

For present purposes we need to estimate the number of chance lexical errors in light of the conclusion that some of the 449 substitution errors were true lexical errors. By definition, true lexical errors do not have the opportunity to come out as lexical errors *by chance*, and this complication must be considered.

Fortunately, the problem can be overcome by taking from the preceding analysis not an estimated number of chance lexical errors, but rather an estimated percentage. The analysis indicated that if all 449 substitution errors had been grapheme-level errors, then about 102, or 23%, would have come out as lexical errors by chance.

According to this estimate, grapheme-level errors should by chance produce lexical errors about 23% of the time, and nonlexical errors about 77% of the time. The observed number of nonlexical errors—258—should therefore represent about 77% of the total number of grapheme-level errors in the set of 449 substitutions. By means of simple algebra ($258 = .77g$, where g is the total number of grapheme-level errors) we can estimate the total number of grapheme-level errors to be about 335, with about 77 of these (335 total−258 nonlexical) being chance lexical errors.

Given that the total number of lexical errors was 191, we can then estimate the number of true lexical errors as 191−77, or 114. According to this estimate, about 60% of the observed lexical errors (114/191) were true lexical errors.

In these computations we used the mean number of lexical errors expected by chance (102) to estimate the proportion of grapheme-level errors that should happen by chance to be words. We can perform the same computations using the number of lexical errors at the 99th percentile of the chance frequency distribution: 120. From this number we obtain .27 (120/449) as the estimated proportion of grapheme errors that yield words by chance. This value can be taken as the highest proportion of lexical errors that could be reasonably expected to occur by chance in the set of 449 substitution errors. Using the .27 value, we estimate that the observed number of nonlexical errors (258) represents about 73% of the total number of grapheme level errors, which yields 353 as the estimated total number of grapheme-level errors, and 95 (353 − 258) as the number of chance lexical errors. The estimated number of true lexical errors then becomes 96 (191 − 95), and the estimated proportion of true lexical errors in the set of observed lexical errors is .50 (96/191). This value can be taken as the minimum proportion of true lexical errors in the set of observed lexical errors. Therefore, our analyses suggest that at the very least, 50% of the observed lexical errors were true lexical errors, and our best estimate is that about 60% of the observed lexical errors were true lexical errors.

This discussion concludes the first step of our three-step argument: At least some of CM's lexical errors are true lexical errors.

Step 2: True lexical errors show the persistence effect

In this section we argue that CM's true lexical errors show the letter persistence effect. The solid line in Figure 8 shows the results of a persistence effect analysis applied to the 191 lexical errors within the set of 449 one- and two-letter substitution errors. As is evident from the figure, the lexical errors showed a strong persistence effect ($p < .000001$).

However, this result alone is not sufficient to establish that *true* lexical errors exhibit the persistence effect. Although the arguments presented in the previous section demonstrated that at least some of the observed lexical errors are true lexical errors, the set of observed lexical errors may (and almost certainly does) include some chance lexical errors as well. If the chance lexical errors show a persistence effect, then the full set of observed lexical errors might show the effect even if there is no persistence effect among the true lexical errors.

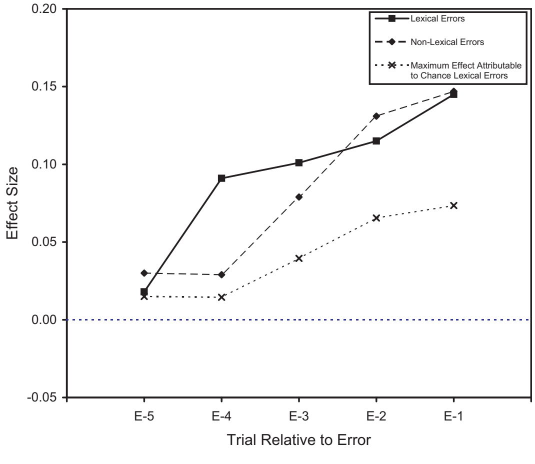

Figure 8. *Letter persistence effect for the lexical and nonlexical one- and two-letter substitution errors. Also shown is maximum effect size expected for the lexical errors if the persistence effect were due entirely to chance lexical errors.*

To demonstrate that the true lexical errors show a persistence effect we must show that the persistence effect for the set of observed lexical errors is too large to be due entirely to whatever chance lexical errors are included in the set. Think of the 449 substitution errors as a mix of (1) true lexical errors resulting from activation of an incorrect word representation at the grapheme level, and (2) grapheme-level errors in which one or two incorrect graphemes were activated at the grapheme level for reasons other than activation of the wrong lexeme representation. The *observed* lexical errors will then be a mixture of true lexical errors, and grapheme-level errors that happened by chance to produce words (i.e., chance lexical errors). The observed nonlexical errors will then be the grapheme-level errors that happened by chance not to produce words. Presumably the letter persistence effect should be the same for grapheme-level errors that did and did not happen to produce words—that is, the persistence effect should be the same for the chance lexical errors and the nonlexical errors. Therefore, if true lexical errors do not show a persistence effect, the persistence effect for observed lexical errors should be weaker than the effect for the observed non-lexical errors. For example, if the set of observed lexical errors is made up of half true lexical errors (which show no persistence effect) and half chance lexical errors (which show the same persistence effect as the nonlexical errors), then the persistence effect for the observed lexical errors should be half as large as that for the nonlexical errors.

According to the arguments developed in the previous section, our best estimate is that about 60% of the observed lexical errors were true lexical errors, and the remaining 40% were chance lexical errors. Accordingly, if the true lexical errors show no persistence effect, we should expect the magnitude of the persistence effect for the observed lexical errors to be only 40% of that for the nonlexical errors. Our estimate of the maximum proportion of observed lexical errors that could be chance lexical errors was 50%, so at the very most the persistence effect for observed lexical errors should be half that for nonlexical errors.

As shown in Figure 8, however, the persistence effect was at least as large for the observed lexical errors (.094) as for the nonlexical errors (.083). The figure also shows the persistence effect expected for the observed lexical errors if (1) half of these errors were chance lexical errors (the maximum proportion according to our estimates), and (2) only the chance lexical errors contributed to the persistence effect, with true lexical errors making no contribution. Obviously, the obtained persistence effect for lexical errors was much too large to be attributed solely to the chance lexical errors. We conclude that CM's true lexical errors show the letter persistence effect.

This conclusion constitutes the second step in our three-step argument:

1. At least some of CM's lexical errors are true lexical errors.
2. The true lexical errors show the letter persistence effect.
3. The presence of the persistence effect in the true lexical errors implies grapheme-to-lexeme feedback.

Despite stacking the deck against ourselves in various ways—for example, by limiting the analysis to CM's one- and two-letter substitution errors, surely among those most likely to have come out as words by chance—we were able to make a strong case for the first two steps of the three-step argument. Nevertheless, the case we developed rests upon a substantial number of assumptions, and a complicated chain of reasoning. In the following section we strengthen the case by offering an independent motivation for steps 1 and 2. Based on a different subset of CM's errors, the alternative foundation is less formal, but equally convincing.

An alternative foundation for steps 1 and 2

Whereas one- and two-letter substitutions could occasionally have produced lexical errors by chance, many of CM's lexical errors were much more complex (e.g., *carpet → compute, grass → craft, feather → fielder*). These complex lexical errors are, we suggest, very unlikely to have occurred by chance—it seems extremely improbable that letter substitutions, deletions, insertions, and transpositions occurring at a post-lexical level without the influence of incorrect lexeme representations would by coincidence produce such errors. For example, to interpret the *carpet → compute* error as a chance lexical error one would have to assume that three letter substitutions (*a → o, r → m, e → u*) and a letter insertion (adding *e* at the end of the word) just happened by coincidence to produce a letter string corresponding to an English word.

We assert, therefore, that CM's complex lexical errors—or, at the very least, the vast majority of these errors—are true lexical errors, reflecting activation of incorrect lexeme representations. Therefore, if the complex lexical errors show a robust persistence effect, this result would constitute another piece of evidence that CM's true lexical errors showed the persistence effect.

We defined complex lexical errors as those the computer program for classifying errors (see p. 282) considered too complex to interpret—that is, the errors placed by the program in the uninterpreted category. One hundred and eighty-nine of CM's lexical errors (35%) fell into this category. Table 7 presents some examples.

Presumably, the vast majority of CM's complex lexical errors are true lexical errors. Therefore, if true lexical errors do not show a letter persistence effect, little or no persistence effect should be seen in the complex lexical errors. In contrast, a substantial persistence effect would make a strong case that true lexical errors show the effect.

Table 7. *Examples of CM's complex lexical errors*

Stimulus	*Response*
shook	floor
stem	chasten
machine	improve
bulk	brush
grew	glove
few	feature

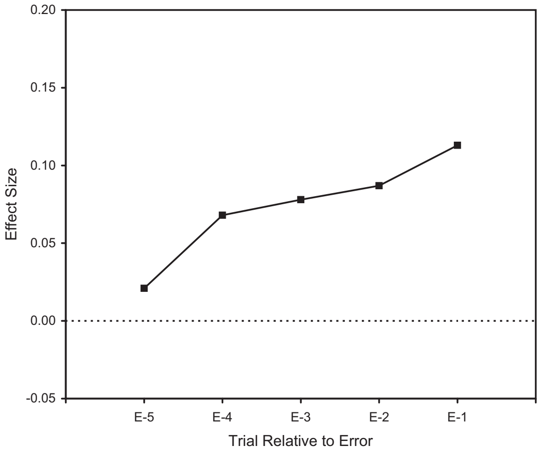

Figure 9. *Letter persistence effect for the complex lexical errors.*

Figure 9 presents the results of the persistence analysis for the complex lexical errors. It is evident from the figure that these errors showed a robust persistence effect ($p < .000001$). Like the results from the one- and two-letter substitution errors, the findings for the complex lexical errors argue strongly that CM's true lexical errors show the letter persistence effect.

Step 3: The persistence effect for true lexical errors implies grapheme-to-lexeme feedback

The finding of a letter persistence effect in CM's true lexical errors cannot be explained by a purely feedforward process in which activation flows from the lexeme to the grapheme level, but does not flow back from the grapheme to the lexeme level. To explain the result it is necessary to postulate grapheme-to-lexeme feedback. Consider the following two-trial sequence:

Trial	*Stimulus*	*Response*
E-1	bench	bench
E	arm	amber

Amber, the incorrect response on trial E, includes two intruded letters (*b* and *e*), both of which were present in the immediately preceding response (*bench*).

We might try to explain the error by assuming that only the correct lexeme—ARM—was activated, but that persisting activation of the B and E graphemes at the grapheme level led to production of *amber*. However, errors occurring in this way are chance lexical errors; what we are trying to explain is the occurrence of the letter persistence effect for true lexical errors, in which the wrong lexeme representation is activated.

Assume, then, that the *arm* → *amber* error was a true lexical error, resulting at least in part from activation of the AMBER representation at the lexeme level. What needs to be explained is why, given that the lexeme representation for an incorrect word was activated, the word was one that shared letters with the preceding response *bench*. Why, for example, was the AMBER lexeme activated, and not, say, the lexeme ARMOR or ARMY or RAMP? We cannot assume that the incorrect word (e.g., AMBER) shared letters with preceding responses simply by chance—the letter persistence effect is the *nonchance* occurrence in a response of letters appearing in preceding responses.

The question, then, is why incorrect lexemes for words (e.g., *amber*) that systematically shared letters with preceding responses (e.g., *bench*) came to be activated. In a purely feedforward process, lexemes are activated from semantic and/or phonological information. Such a process could lead to semantic or phonological errors. However, *arm* and *amber* are not semantically related; more generally, CM's lexical errors rarely showed semantic similarity between stimulus and response. Phonological errors could lead to an apparent letter persistence effect—persisting activation of phoneme representations from previous trials could lead to activation of lexemes for words sharing phonemes (and therefore letters) with previous responses, leading to a letter persistence effect. However, we have already shown that CM's letter persistence effect is not phonologically based (see pp. 287–289).

To account for the letter persistence effect in true lexical errors, we need to assume that persisting activation of letter representations at the grapheme level influences the activation of word representations at the lexeme level. For example, in the case of the *arm* → *amber* error we need to assume that persisting activation of the B and E graphemes (activated during generation of the preceding response *bench*) led to some activation of the AMBER representation at the lexeme level, and that this activation contributed to the selection of the wrong representation at the lexeme level during the processing of the stimulus *arm.*

The most straightforward way for activation at the grapheme level to affect activation on the lexeme level is to assume that connections between lexeme and grapheme nodes are bidirectional, such that activation spreads not only "forward" from lexeme nodes to the corresponding grapheme nodes, but also "backward" from grapheme nodes to lexeme nodes for words containing the graphemes (see Figure 2b). According to this hypothesis, activation of the BENCH lexeme activates the B, E, N, C, and H grapheme nodes, and activation then feeds back from these nodes to the lexemes for all words containing these letters (including, among many others, AMBER). Then, when the stimulus "arm" is dictated, activated semantic and/or phonological nodes send activation to the ARM lexeme, which in turn activates A, R, and M nodes at the grapheme level. These nodes in turn feed back activation to lexemes for words containing the letters *a, r*, or *m.* By this means the activation of the AMBER lexeme is increased and this lexeme, which is also receiving activation from the persisting grapheme nodes B and E, may become more activated than the correct lexeme ARM (which receives no support from the persisting B, E, N, C, and H graphemes). This computational architecture, in which activation spreads forward and then backward over multiple iterations, is an instance of what has been referred to as an attractor system (e.g., Plaut & Shallice, 1993).[7] In such a system the interconnected levels constrain one another as the system settles into a stable state. In the case of CM, graphemes that remained active from preceding trials feed activation back to the lexeme level, contributing to selection of an incorrect lexeme, and leading in turn to the activation of that lexeme's corresponding graphemes.

Two potential alternative interpretations merit brief consideration. First, one might suggest that the letter persistence effect for true lexical errors reflects *external* rather than *internal* grapheme-to-lexeme feedback. In particular, one might imagine that visual feedback occurring when CM looked at what he was writing (e.g., *bench*) activated letter representations (e.g., *b, e*, etc.), which in turn activated lexeme representations for words sharing those letters (e.g., *amber*). However, we have shown that the letter persistence effect does not result from visual feedback (see pp. 290–291).

Second, it is conceivable that our results could be explained by some form of attractor architecture that does not involve feedback from a grapheme to a lexeme level. This category of potential explanations cannot be ruled out as a class. However, specific proposals could be formulated and evaluated. Here we briefly consider one example. It might be suggested that the persistence effect for true lexical errors resulted from connections among orthographic lexeme nodes. Specifically, one might posit excitatory connections between lexemes based on orthographic similarity, such that activation spreads from an activated lexeme representation to the representations for orthographically similar words. For example, activation of BENCH would result in some activation of the AMBER lexeme, due to the shared letters *b* and *e.* Activation of ARM on the following trial would also lead to some activation of AMBER, which, with persisting activation from the previous trial, might lead AMBER to become more activated than ARM itself. This interpretation could account for a letter persistence effect in true lexical errors; however, the assumption of orthographically based excitatory

---

[7] Although the notion of an attractor network is often associated with the assumption of distributed representations, this association is by no means necessary. Systems with either local or distributed representations for lexemes and graphemes can have an attractor architecture. Although our figures show lexemes and graphemes as individual nodes, our hypothesis is neutral with respect to whether the representations are local or distributed.

connections among lexeme nodes is problematic for several reasons. Most notably, positing excitatory lexeme-to-lexeme connections is implausible in light of processing considerations. At the orthographic lexeme level the goal of processing is to arrive at a state in which the correct lexeme node is strongly activated and all other lexeme nodes are inactivated (or at most weakly activated). Excitatory connections between lexeme nodes would not contribute in any obvious way to this goal, and in fact would probably work against it by promoting activation of incorrect lexemes. Indeed, computational models often assume *inhibitory* connections among the nodes in a set when the goal of processing is to activate only one of the nodes. In these models the goal state is achieved through an iterative "winner-takes-all" process in which a single node gradually emerges as the winner, in part by inhibiting the competing nodes (e.g., McClelland & Rumelhart, 1981). For example, the computational model of spelling proposed recently by Houghton and Zorzi (2003) assumes inhibitory connections and competitive interactions among orthographic lexeme nodes (although this aspect of the model is not implemented in their simulations).

Whereas excitatory lexeme-to-lexeme connections would probably be counterproductive, excitatory grapheme-to-lexeme feedback could facilitate selection of the correct lexeme. Although feedback connections would send activation not only to the correct lexeme but also to lexemes for orthographically similar words, the correct lexeme would receive the strongest activation, contributing to the "victory" of this lexeme in a competitive activation process. (See the General Discussion for elaboration of this point.)

On the basis of the preceding evidence and arguments, we conclude that the letter persistence effect for CM's true lexical errors constitutes evidence of grapheme-to-lexeme feedback in the spelling system.

GENERAL DISCUSSION

Through extensive analyses of the impaired spelling performance exhibited by patient CM we have developed an argument for grapheme-to-lexeme feedback in the cognitive spelling system. Two features of CM's spelling were central to the argument. First, letters from prior spelling responses intruded into subsequent responses at rates far greater than expected by chance. We presented evidence that this letter persistence effect resulted from abnormal persistence of activation at the level of grapheme representations. Second, CM made many formal lexical errors (e.g., *arm* → *amber*). We demonstrated that a large proportion of these errors were "true" lexical errors originating in lexical selection, rather than "chance" lexical errors that happened by chance to take the form of words. Crucially, analyses revealed that CM's true lexical errors exhibited the letter persistence effect. This finding, we argued, can be understood only within an architecture in which activation from the grapheme level feeds back to the lexeme level, in this way exerting an influence on lexical selection.

Other evidence for grapheme-to-lexeme feedback?

As we noted in the Introduction, little attention has previously been directed toward questions concerning feedback in the cognitive spelling system. Folk, Rapp, and Goldrick (2002) discussed several findings that were consistent with grapheme-to-lexeme feedback, but concluded that none constituted clear-cut evidence for feedback. Feedback is also discussed, albeit very briefly, in a recent paper by Romani and colleagues (2002). These authors studied a dysgraphic patient (DW) whose spelling errors consisted predominantly of words formally related to the target (e.g., *table* → *cable*). The principal aim of the study was to investigate alternative hypotheses regarding the source of formal lexical errors and, in this regard, Romani et al. concluded that DW suffered from "a lexical impairment that makes selection between competing representations difficult (possibly because of reduced inhibition-increased activation)" (p. 329).

In interpreting their results Romani et al. also posited grapheme-to-lexeme feedback, suggesting that "DW's errors arise because of confusions among a cohort of lexical neighbors activated top-down from a phonological input *and bottom-up from shared letters*" (p. 325, emphasis added). However, feedback was not a focus of their investigation, and the evidence presented was not strong. Analysing 252 nonmorphological lexical errors, Romani and colleagues found that most of the errors were both phonologically and orthographically similar to the target word. However, 29 of the errors were more phonologically than orthographically similar to their targets, and 15 were more orthographically than phonologically similar. Romani et al. (2002) interpreted the former subset of errors as evidence for top-down activation of lexical representations from phonology, and the latter as evidence for bottom-up activation of lexical representations via feedback from letter representations.

However, these conclusions are not entirely well founded. In particular, the fact that some errors were more orthographically than phonologically related to the target does not imply that orthographic influences (such as grapheme-to-lexeme feedback) played any role in the genesis of the errors. The phonological neighbours of a target word will vary in their orthographic similarity to the target, and occasionally the orthographic similarity may just happen to be greater than the phonological similarity. For example, *steak* and *steal* are phonological neighbours by virtue of sharing the first two phonemes, but actually have greater orthographic than phonological overlap. Therefore, even if lexical representations were activated in a purely top-down fashion from phonology (and semantics), with no feedback from grapheme representations, activation of the target's phonological neighbours could lead to selection of a word with greater orthographic than phonological overlap with the target. Thus, in DW's error *curbs* → *rubs* (cited by Romani et al. as an example of greater orthographic than phonological target-error overlap), *rubs* may have been activated solely because of its phonological similarity to *curbs*. Hence, the finding that DW occasionally made errors of this sort does not constitute evidence for grapheme-to-lexeme feedback. (See Folk et al., 2002, for a method of evaluating orthographic and phonological influences on formal lexical errors.) In contrast, the evidence from CM specifically supports feedback by demonstrating the influence of persisting grapheme activation on lexical selection.

Feedback and graphemic buffer deficits

From our assumptions about grapheme-to-lexeme feedback, and the role of the feedback in CM's lexical errors, a question arises about interpretation of results from other dysgraphic patients: Does our position imply that any patient with a deficit affecting the grapheme level should make substantial numbers of formal lexical errors? Consider in particular patients characterised as having graphemic buffer deficits, who frequently produce letter substitution and insertion errors, presumably reflecting activation of incorrect letter representations at the grapheme level (e.g., Caramazza & Miceli, 1990; Caramazza, Miceli, Villa, & Romani, 1987). In positing grapheme-to-lexeme feedback, are we committed to predicting that such patients should show high rates of formal lexical errors, due to feedback from erroneously activated graphemes to incorrect lexemes? If so, the prediction represents a significant problem for the grapheme-to-lexeme feedback hypothesis, because results from several well-studied graphemic buffer patients show very low rates of lexical errors. For example, Caramazza et al. (1987) reported that lexical errors accounted for only 16 of patient LB's 305 errors in spelling words, and most if not all of these errors could have occurred by chance. (See also McCloskey, et al., 1994.)

Fortunately, the grapheme-to-lexeme feedback hypothesis is not inconsistent with the results from graphemic buffer patients. In patients with graphemic buffer deficits grapheme-level activation may be abnormally weak—buffer patients are typically assumed to have difficulty activating, or maintaining activation of, grapheme representations. If grapheme-level activation were abnormally weak, then grapheme-to-lexeme feedback would also be weak, and might therefore have little effect on selection of an orthographic lexeme. For CM, in contrast, we have assumed that grapheme-level activation is abnormally strong, because activation persists abnormally over time.

Another potentially relevant observation is that feedback from erroneously activated graphemes may be more likely to cause errors in selecting orthographic lexemes when there is some impairment at the lexeme level, such that activation at this level is abnormally weak or noisy. Thus, patients with selective graphemic buffer deficits may make few lexical errors not only because of weak activation at the grapheme level, but also because processing at the lexical level is intact. (As we noted in an earlier section, CM may, in contrast, have mild impairment at the orthographic lexeme level.)

Potential functions of feedback

Our arguments and evidence for grapheme-to-lexeme feedback lead naturally to questions about the computational functions of such feedback. Generally speaking, feedback can serve two related functions. The first is to bring additional constraints to bear on a decision or process—in this case the lexical selection process. The second function is to strengthen, stabilise, or bias certain representations or responses—in this case, the target lexeme and grapheme representations, and hence the correct spelling response.

The lexical selection process chooses lexemes to express meanings that an individual has in mind. Lexical selection should therefore be constrained primarily by semantic input. However, the lexical selection process might also usefully be informed by information at subsequent representational levels (see, e.g., Dell et al., 1997; Rapp & Goldrick, 2000). In the case of spelling, grapheme-to-lexeme feedback allows lexical selection to be influenced by any factor that can exert an effect on grapheme activation. For example, if we posit a functional architecture in which both lexical and sublexical processes activate the same set of graphemes (as in Figure 1; see also Houghton & Zorzi, 2003), grapheme activation will be determined in part by the output of the sublexical phoneme–grapheme conversion system (in tasks such as spelling to dictation). Through grapheme-to-lexeme feedback, lexical selection would then be influenced by the plausible spellings of the sounds of the target word (Folk et al., 2002), providing a mechanism for the "summation" of lexical and sublexical information proposed by Hillis and Caramazza (1991a). One might wonder if the contribution of the sublexical phoneme–grapheme conversion system would be detrimental in the case of irregular words. Although this question would certainly benefit from computational analysis, it is important to note (as pointed out by Houghton & Zorzi, 2003) that the majority of phoneme–grapheme correspondences, even in irregular words, are highly predictable. As a result, input from the sublexical system should primarily serve to strengthen a target word relative to its competitors (see Folk et al., 2002, for an empirical test of this claim).

In addition to providing multiple constraints on selection, feedback provides a means by which particular representations or responses can be strengthened, stabilised, or favoured. Dell (1986) pointed out that cyclic feedforward and feedback processing can create positive feedback loops between lexical representations and their segments—that is, a lexeme activates the corresponding graphemes, which further activate the lexeme via feedback, increasing the feedforward activation sent to the graphemes, and so forth. Under normal circumstances, positive feedback loops would favour the target lexeme and its corresponding graphemes, stabilising the correct representations and biasing the system toward the correct response. (See Dell & Gordon, 2003, for computational evidence regarding some of the specific benefits of feedback.)

These suggestions about potential functions of grapheme-to-lexeme feedback do not constitute, and are not put forward as, definitive logical arguments for the necessity of feedback in the cognitive spelling system. (If such arguments could be made, we would not need the empirical evidence from CM.) Rather, our point is simply that the conclusions about grapheme-to-lexeme feedback we drew from analyses of CM's performance are not implausible in light of potential computational functions of feedback.

Feedback as a mechanism for "summation" effects

We conclude by elaborating the above-mentioned point that feedback provides a mechanism for the interaction or summation of lexical and sublexical information in lexical selection. A number of researchers (e.g., Hillis & Caramazza, 1991a, 1995; Hillis, Rapp, & Caramazza, 1999; Miceli, Benvegnù, Capasso, & Caramazza, 1997; Miceli, Capasso, & Caramazza, 1994) have proposed that selection of a lexical representation is constrained not only by activation from semantic representations, but also by the output of sublexical phonology-to-orthography conversion processes (in the case of spelling to dictation) and orthography-to-phonology conversion processes (in the case of reading). For example, Hillis and colleagues (1999) tested a dysgraphic patient (RCM) in the first week after she suffered a stroke, and again 2 weeks later. Results from the first phase of testing pointed to deficits affecting the orthographic output lexicon and the sublexical phonology–orthography conversion process. At this time RCM exhibited high rates of semantic errors in spelling to dictation. However, by the time of the second testing phase RCM's sublexical processing had improved, and her rate of semantic errors was far lower.

Hillis et al. (1999) interpreted this pattern of results as evidence for summation of lexical and sublexical information in lexical selection. According to their account, dictation of a stimulus word (e.g., "knife") activated a lexical-semantic representation, leading in turn to activation of the orthographic lexemes for multiple semantically-related words (e.g., *knife, fork, spoon, plate*). Because RCM's orthographic output lexicon was damaged, a nontarget lexeme (e.g., *fork*) was sometimes more strongly activated than the target lexeme (*knife*). In the first testing phase no additional constraints were brought to bear on lexical selection, and as a consequence semantic errors were frequent (e.g., *knife* → *fork*). However, by the time of the second testing phase RCM's orthography–phonology conversion process was functioning to some extent. When a word was dictated, this sublexical process generated a potential spelling (e.g., *nife*), and this spelling representation somehow influenced the activation of lexeme representations in the orthographic output lexicon. Even when the sublexically generated spelling was not entirely correct (as in the case of *nife*), it usually favoured the target within the pool of lexemes activated from semantic input (e.g., *knife, fork, spoon, plate*), thereby reducing the incidence of semantic errors.

Although these results and others have been presented in support of the summation hypothesis, specific mechanisms for summation have not usually been discussed. We suggest that feedback provides such a mechanism (see also Folk et al., 2002). In the case of spelling, grapheme-to-lexeme feedback constitutes a plausible mechanism for summation effects: Sublexical processes activate grapheme representations, which feed activation back to the orthographic lexemes, in this way influencing lexical selection in favour of the target word. Similarly, phoneme-to-lexeme feedback could conceivably account for summation effects in reading.

Manuscript received 24 November 2003
Revised manuscript received 14 January 2005
Revised manuscript accepted 25 January 2005

REFERENCES

Baars, B. J., Motley, J. T., & MacKay, D. (1975). Output editing for lexical status from artificially elicited slips of the tongue. *Journal of Verbal Learning and Verbal Behavior*, *14*, 382–391.

Best, W. (1996). When racquets are baskets but baskets are biscuits, where do the words come from? A single case study of formal paraphasic errors in aphasia. *Cognitive Neuropsychology*, *13*, 443–480.

Caramazza, A., & Miceli, G. (1990). The structure of graphemic representations. *Cognition*, *37*, 243–297.

Caramazza, A., Miceli, G., Villa, G., & Romani, C. (1987). The role of the graphemic buffer in spelling:

Evidence from a case of acquired dysgraphia. *Cognition*, *26*, 59–85.

Cohen, L., & Dehaene, S. (1998). Competition between past and present: Assessment and interpretation of verbal perseverations. *Brain*, *121*, 1641–1659.

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

Dell, G. S., & Gordon, J. K. (2003). Neighbors in the lexicon: Friends or foes? In N. O. Schiller & A. S. Meyer (Eds.), *Phonetics and phonology in language comprehension and production: Differences and similarities*. New York: Mouton de Gruyter.

Dell, G. S., & Reich, P. A. (1980). Toward a unified model of slips of the tongue. In V. A. Fromkin (Ed.), *Errors in linguistic performance: Slips of the tongue, ear, pen, and hand* (pp. 473–486). New York: Academic Press.

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

Dell, G. S., Schwartz, M. F., Martin, N., Saffran, E. M., & Gagnon, D. A. (1997). Lexical access in aphasic and nonaphasic speakers. *Psychological Review*, *104*, 801–838.

Dunn, L. M., & Dunn, L. M. (1981). *Peabody Picture Vocabulary Test-Revised.* Circle Pines, MN: American Guidance Service.

Folk, J., Rapp, B., & Goldrick, M. (2002). Interaction of lexical and sublexical information in spelling: What's the point? *Cognitive Neuropsychology*, *19*, 653–671.

Garrett, M. F. (1976). Syntactic processes in sentence production. In R. J. Wales & E. Walker (Eds.), *New approaches to language mechanisms* (pp. 231–255). Amsterdam: North Holland.

Goldrick, M., & Rapp, B. (2002). A restricted interaction account (RIA) of spoken word production: The best of both worlds. *Aphasiology*, *16*, 20–55.

Goodglass, H., & Kaplan, E. (1983). *The assessment of aphasia and related disorders.* Philadelphia: Lea & Febiger.

Goodman, R. A., & Caramazza, A. (1985). *The Johns Hopkins University Dysgraphia Battery.* Baltimore, MD: Johns Hopkins University.

Graham, N. L., Patterson, K., & Hodges, J. R. (1997). Progressive dysgraphia: Co-occurrence of central and peripheral impairments. *Cognitive Neuropsychology*, *14*, 975–1005.

Graham, N. L., Patterson, K., & Hodges, J. R. (2000). The impact of semantic memory impairment on spelling: Evidence from semantic dementia. *Neuropsychologia*, *38*, 143–163.

Harley, T. A. (1984). A critique of top-down independent levels models of speech production: Evidence from non-plan-internal speech errors. *Cognitive Science*, *8*, 191–219.

Hillis, A. E., & Caramazza, A. (1991a). Category-specific naming and comprehension impairment: A double dissociation. *Brain*, *114*, 2081–2094.

Hillis, A. E., & Caramazza, A. (1991b). Mechanisms for accessing lexical representations for output: Evidence from a category-specific semantic deficit. *Brain and Language*, *40*, 106–144.

Hillis, A. E., & Caramazza, A. (1995). Converging evidence for the interaction of semantic and sublexical phonological information in accessing lexical representations for spoken output. *Cognitive Neuropsychology*, *12*, 187–227.

Hillis, A. E., Rapp, B. C., & Caramazza, A. (1999). When a rose is a rose in speech but a tulip in writing. *Cortex*, *35*, 337–356.

Houghton, G., & Zorzi, M. (2003). Normal and impaired spelling in a connectionist dual-route architecture. *Cognitive Neuropsychology*, *20*, 115–162.

Levelt, W. J. M., Roelofs, A., & Meyer, A. S. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences*, *22*, 1–75.

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

Martin, N., Gagnon, D. A., Schwartz, M. F., Dell, G. S., & Saffran, E. M. (1996). Phonological facilitation of semantic errors in normal and aphasic speakers. *Language and Cognitive Processes*, *11*, 257–282.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review*, *86*, 287–330.

McCloskey, M., Badecker, W., Goodman-Schulman, R., & Aliminosa, D. (1994). The structure of graphemic representations in spelling: Evidence from a case of acquired dysgraphia. *Cognitive Neuropsychology*, *11*, 341–392.

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

Miceli, G., Capasso, R., & Caramazza, A. (1994). The interaction of lexical and sublexical processes in reading, writing, and repetition. *Neuropsychologia*, *32*, 317–333.

Patterson, K. (1986). Lexical but nonsemantic spelling? *Cognitive Neuropsychology*, *18*, 729–748.

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

Rapp, B., & Goldrick, M. (2000). Discreteness and interactivity in spoken word production. *Psychological Review*, *107*, 460–499.

Rapp, B., & Goldrick, M. (2003). Feedback by any other name is still interactivity: A reply to Roelofs' comment on Rapp and Goldrick (2000). *Psychological Review.*

Roeltgen, D. P., Rothi, L. G., & Heilman, K. M. (1986). Linguistic semantic agraphia: A dissociation of the lexical spelling system from semantics. *Brain and Language*, *27*, 257–280.

Romani, C., Olson, A., Ward, J., & Ercolani, M. G. (2002). Formal lexical paragraphias in a single case study: How "masterpiece" can become "mister-pieman" and "curiosity" "suretoy". *Brain and Language*, *83*, 300–334.

Stemberger, J. P. (1985). An interactive activation model of language production. In A. W. Ellis (Ed.), *Progress in the psychology of language* (pp. 143–186). Hove, UK: Lawrence Erlbaum Associates Ltd.

Tainturier, M.-J., & Rapp, B. (2001). The spelling process. In B. Rapp (Ed.), *The handbook of cognitive neuropsychology: What deficits reveal about the human mind* (pp. 263–289). Philadelphia: Psychology Press.