Cheap But Not Free: Commentary on Pfau 'Cheap Repairs'*

Adam Albright
Massachusetts Institute of Technology

Pfau's (2007) study tests several important predictions of a modular, non-interactive framework such as Distributed Morphology (DM). In this commentary, I argue that one of these, gender accommodation, is not actually predicted by any intrinsic feature of DM. However, the effect is at least compatible with a certain version of it, and in fact, usefully informs our understanding of how lexical gender is accessed and copied. I also discuss a variety of other interesting but as yet untested predictions for speech errors.

1. Introduction

Morphologically rich languages are a minefield of possible errors. In even a moderately complex language like German, every sentence involves some amount of gender, person, number, and case agreement. One might expect that errors concerning one aspect of a sentence (for example, the lexical items involved) could have far-reaching consequences, leading to a wide variety of ill-formed structures. What is so striking about the cases discussed by Pfau (2007) is just how minimal the morphological impact of an error can be. Frequently, a single lexical exchange or substitution does not lead to a cascade of broken agreement relations, but is instead "repaired" by a series of adjustments to accommodate the new structure. This is seen quite clearly in Pfau's example (31), repeated here in (1)[1]:

(1) Intended: Da war der Bann ge-broch-en
then was the-MASC.SG spell(MASC) PART-break-PART
'Then the spell was broken'

Error:	Da	war	der	***Broch***	ge-***bann***-en	(swap)
Repair:	Da	war	der	Bruch	ge-bann-en	(vowel in √BRECH)
Repair:	Da	war	der	Bruch	ge-bann-t	(past participle suffix)

* I owe special thanks to Roland Pfau for a stimulating presentation and extended discussion of related issues, and also to Heidi Harley and Alec Marantz for helpful discussion and pointers; all subsequent errors and confusions are my own. Thanks also to participants in the workshop, and particularly to the organizers, Carson Schütze and Vic Ferreira.

[1] I follow Pfau's typographic conventions for slip examples, described in his footnote 5.

MIT Working Papers in Linguistics 53, 35–51
The State of the Art in Speech Error Research: Proceedings of the LSA Institute Workshop
Carson T. Schütze and Victor S. Ferreira (Eds.)

When the consequences of an error are seen as a sequence of repairs, the necessary accommodations can be quite elaborate. The important observation, however, is that although the degree of accommodation in (1) is remarkable, it is no more remarkable than the series of agreement operations needed to produce a typical (non-errorful) German sentence. Ideally, such accommodations should be analyzed not as the result of some special clean-up operation that is seen only in errors, but rather, as the normal workings of the morphological system given a slightly incorrect input (Fromkin 1971). As Pfau shows, a grammatical framework such as Distributed Morphology (DM: Halle and Marantz 1993) can be an important tool for pursuing such an analysis, since it provides a set of specific hypotheses about which operations occur "downstream" of other operations (and thus are predicted to show accommodation effects). My goals in this commentary are the following: first, I will discuss in some detail Pfau's analysis of gender accommodations. I will argue that such accommodations are not predicted as straightforwardly as claimed by the DM architecture, but that they nevertheless usefully inform our understanding of how grammatical gender is entered into the derivation. The conclusion will be that although gender accommodation is not incompatible with the general architecture, it does provide evidence against a very strongly modular version of the theory. I will then turn to some other predictions of the DM framework concerning the predicted (lack of) influence of phonological form on earlier stages of the derivation. Finally, I will discuss a handful of additional predictions that the theory makes, but which remain thus far untested. The issues that emerge here do not undermine Pfau's basic argument, but merely serve to point out areas where further investigation is needed.

2. The treatment of gender in Distributed Morphology

Pfau's corpus analysis reveals two distinct effects of gender. First, when a wrong word is selected, it frequently has the same canonical gender as the intended target (the "identical gender effect"). This gender-matching effect in naturally occurring errors has been noted previously in the literature (Berg 1992; Marx 1999; Arnaud 1999; Pfau 2000), and is observed no matter whether the intrusive word is semantically or phonologically similar to the target word. The top graph in Figure 1 shows the prevalence of gender-preserving substitutions, using data from Pfau (2000): the "identity" line is significantly higher than the "mismatch" line. The tendency for substitutions to match the gender of the intended word shows that gender information is available even before the target word has been selected conclusively.[2] One might say, in the terms of Dell (1986), that activating the target word co-activates other words sharing not only similar semantics and phonological form, but also the same morphological gender.

[2] This conclusion is also supported by the fact that speakers in a tip-of-the-tongue state are often able to identify the gender of the intended word, even if they cannot recall the details of its phonological form (e.g., Miozzo and Caramazza 1997).

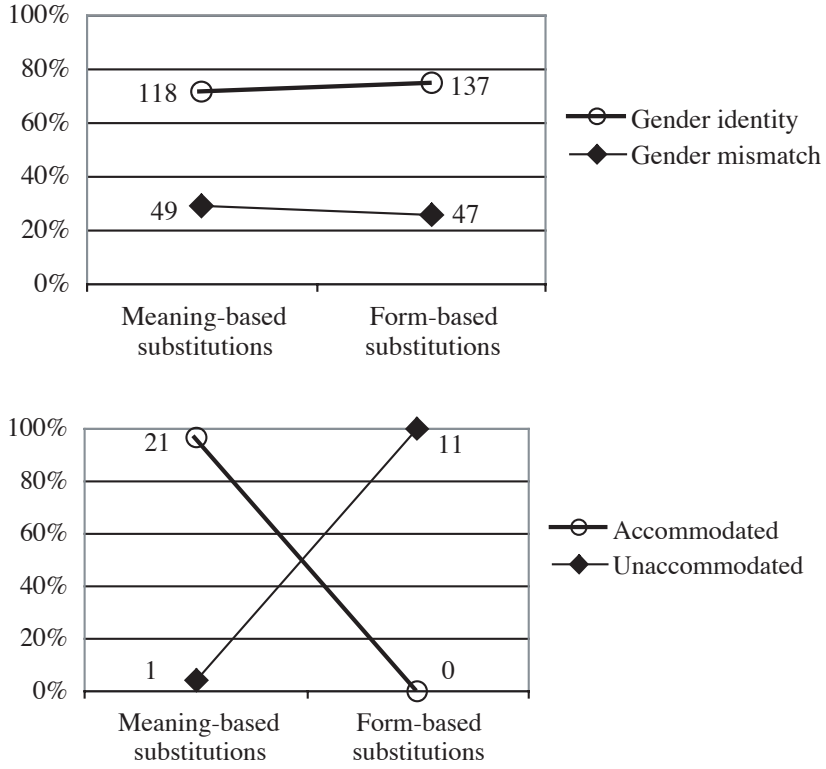

Figure 1
Rates of gender matching (top) and accommodation (bottom) (Pfau 2000)

A more striking finding concerns those instances in which substitutions do *not* match in gender: Pfau observes that in his corpus, when the intrusive word bears a semantic relation to the intended target, the surrounding determiner phrase is almost always adjusted to accommodate the gender of the new noun (21/22 cases), but when the intrusive word is a phonologically-motivated substitution, no accommodation is seen (0/11 cases) (Figure 1 bottom = Pfau's Table 1). This suggests that not only is the gender of the target word available before its phonological content has been fully retrieved, but also that the copying operations that lead to agreement take place some time between the initial (semantically-based) selection and the final phonological selection. In order to capture this effect, we need a model in which gender features are entered into the derivation early (before morphological feature copy operations take place), but phonological features are introduced later, after agreement has already taken place.

As Pfau points out, this order of operations is broadly consistent with the architecture of Distributed Morphology, in which syntax and morphology operate on abstract roots and morphosyntactic features, while phonological informa-

tion is introduced only later in the derivation. He accomplishes the early insertion of gender features by hypothesizing that abstract roots are already specified for gender when they are entered into the syntax. This move appears to be motivated empirically, both by the fact that we see accommodation in speech errors, and also, more fundamentally, by the fact that there is agreement for morphological gender in the first place. It is also not unprecedented; for example, Müller (2004) assumes that gender (and even inflectional class) features of lexical items are provided to the syntax. As far as I can tell, although early access to gender features is taken by many analysts as a background assumption, very few arguments have been put forward for why it could not be otherwise. Thus it is worth pausing to consider how the speech error data contributes to the argument for early insertion of gender features, how this fits into the general architecture of DM, and what further predictions are made.

2.1 What does syntax know about lexical gender?

As mentioned above, a fundamental tenet of Distributed Morphology is that syntax operates on roots and bundles of morphosyntactic features, rather than on fully specified lexical items. This immediately raises the question of what kinds of features are actually available from the outset, in the input to syntax. This is clearly an empirical question, but a conceptually appealing null hypothesis is that syntax cares only about a limited set of universal syntactic and semantic features (person, number, animacy, etc.), and that lexically arbitrary and language-particular morphological and phonological features (such as the phonemes, the inflectional class, etc.) are available only after the specific vocabulary items have been selected and inserted (*feature disjointness*; Marantz 1995; Embick 2000, p. 188).

Gender agreement appears to present a challenge to this strong hypothesis, since on the face of it, morphological gender is exactly the type of language-particular and lexically arbitrary feature that would be excluded from the input to syntax. This naturally leads one to wonder whether it is possible to get around the need for early access to gender features, and preserve the hypothesis that lexically arbitrary information is inserted late. One possibility would be to have the syntax and morphology "prepare for all eventualities," freely generating DPs with all different gender features attached to them, and only pronouncing the ones that turn out to be compatible with the intended lexical items. Concretely, consider the German sentence in (2):

(2)	Das	Mädchen	bricht	einen	Stock.
	the-NEUT.SG	girl	break-3.SG	a-MASC.SG	stick(MASC)
	'The girl breaks the stick.'				

To generate this sentence without knowing the genders ahead of time, the syntax would need to generate 9 ($=3^2$) parallel structures, including masculine, feminine, and neuter versions of each root:

(3)	i.	[NOM][+def][+masc]	√GIRL	√BREAK	[ACC][−def][+masc]	√STICK
	ii.	[NOM][+def][+fem]	√GIRL	√BREAK	[ACC][−def][+masc]	√STICK
	iii.	[NOM][+def][+neut]	√GIRL	√BREAK	[ACC][−def][+masc]	√STICK
	iv.	[NOM][+def][+masc]	√GIRL	√BREAK	[ACC][−def][+fem]	√STICK
	v.	[NOM][+def][+fem]	√GIRL	√BREAK	[ACC][−def][+fem]	√STICK
	vi.	[NOM][+def][+neut]	√GIRL	√BREAK	[ACC][−def][+fem]	√STICK
	vii.	[NOM][+def][+masc]	√GIRL	√BREAK	[ACC][−def][+neut]	√STICK
	viii.	[NOM][+def][+fem]	√GIRL	√BREAK	[ACC][−def][+neut]	√STICK
	ix.	[NOM][+def][+neut]	√GIRL	√BREAK	[ACC][−def][+neut]	√STICK

The intended vocabulary items √GIRL ↔ $[\text{meːdçən}]_{[+\text{neut}]}$ and √STICK ↔ $[\text{ʃtɔk}]_{[+\text{masc}]}$ can be inserted only into a structure that provides features for a neuter subject and masculine object (3.iii), so this is the one that is spelled out successfully. Although such a solution seems inelegant, it works mechanically, and has the virtue of allowing us to maintain Late Insertion for gender features.

In fact, Embick (2000) contemplates a parallel maneuver to analyze Latin deponent verbs, which have the lexically arbitrary property of requiring passive morphology but otherwise require active syntax. In Latin, the solution requires generating structures in which the verb root is given a passive feature, but occurs in what is otherwise an active structure (no passive feature under *v*). As with the example in (3), these "vacuously passive" structures are used just in case Vocabulary Insertion wishes to insert a deponent verb; additional assumptions are needed to ensure that they cannot be spelled out using active verbs marked with gratuitous passive morphology. Embick points out that for deponent verbs, this approach involves not just an inefficiency, but a fundamental change in how the feature [passive] is conceived, since the syntax must be able to arbitrarily insert [passive] features under verb roots (rather than under *v*, where it normally occurs just in case it is licensed by a genuinely passive structure). The situation may be simpler for morphological gender, however, since at least in many languages there is no syntactic difference between gender that has been assigned for semantic reasons ("natural gender") or for arbitrary lexical reasons.[3] As a result, even if we allow the syntax to arbitrarily add gender features to noun roots in anticipation of inserting vocabulary items of various genders, we do not need to perform any complicated maneuvers to distinguish syntactically between gender features that are licensed by semantics on the one hand, and those that will need to be licensed arbitrarily by a particular vocabulary item on the other. Thus, as far as I can tell, for analyzing normal German productions, the main cost of not providing lexical gender to the syntax ahead of time is the ineffi-

[3] Sauerland (2005) discusses a possible exception in German, involving pronouns that can agree with the antecedent's grammatical gender, or the referent's natural gender:

(i)	Kein	Mädchen	glaubt,	dass	es/sie	überfordert	wird.
	no	girl(NEUT)	believes	that	it(NEUT)/she(FEM)	overchallenged	is

Deriving the non-agreeing structure in which [+fem] is licensed in the syntax/semantics could indeed require a maneuver parallel to Embick's Latin example, in which [pass] under *v* is licensed syntactically in truly passive structures. See also Carstens (2000) and Pesetsky and Torrego (2007) for discussion of additional issues in the syntactic analysis of gender features.

ciency of constantly needing to provide parallel derivations in anticipation of various possible vocabulary items.

Importantly, although this analysis works (admittedly rather awkwardly) for the surface facts of German gender agreement, it makes wrong predictions about errors and accommodation. If it were true that syntax blindly generated parallel candidate structures for different genders, then in principle, agreement frames for all different genders would be available at the time of Vocabulary Insertion. This implies that if something went wrong and the wrong vocabulary item was inserted, then accommodation should always be possible since the syntax has provided a whole set of agreeing frames as candidate structures.

(4) Input to syntax:

[+def] √GIRL √BREAK [–def] √STICK

Output of syntax:

i. [NOM][+def][+neut] √GIRL √BREAK [ACC][–def][+masc] √STICK
ii. [NOM][+def][+neut] √GIRL √BREAK [ACC][–def][+fem] √STICK
(etc., viz. (3))

(Erroneous) vocabulary insertion

√GIRL ↔ [meːdçən]$_{[+neut]}$
~~√STICK ↔ [ʃtɔk]$_{[+masc]}$~~ ⇐ √SOCK ↔ [zɔkə]$_{[+fem]}$

Output, with "accommodation": Das Mädchen bricht eine ***Socke***

In this hypothetical example, the vocabulary item *Socke* 'sock' has accidentally been inserted instead of the intended *Stock* 'stick', due to their phonological similarity. This error could have one of two possible consequences. Suppose first that the gender of the intrusive item is retrieved along with its phonological form; in this case, the intrusive item should be inserted into the appropriate (accommodating) structure ((4.ii)): *Das Mädchen bricht eine*$_{[+fem]}$ ***Socke***$_{[+fem]}$. If, on the other hand, the gender of √SOCK is not retrieved, then no mechanism would be available to assign a gender to the object, and there would be no basis for choosing among the candidate structures; perhaps one would be chosen at random, or perhaps the derivation would simply crash. Thus, we predict that phonological substitutions should either receive accommodation, or get some randomly selected erroneous gender. Crucially, we would not expect phonological substitutions to maintain the gender of the intended word, since the syntax doesn't know the gender of the intended item, and thus has no commitment to providing a structure that agrees with it.

Recall that this prediction is disconfirmed by Pfau's corpus counts, which show that phonological substitutions tend *not* to be accommodated (Figure 1), and that they instead preserve the gender of the intended word: *ein*$_{[+masc]}$ *neuer*$_{[+masc]}$ ***Luft***$_{[+neut]}$, *äh, Duft*$_{[+masc]}$ 'a new air, uh, fragrance' (Pfau's example (17b)). To capture this effect, the actual gender of the intended word must already be known at the time when features are copied to create agreement. Thus, the data from gender accommodations provides evidence against the strong version of Late Insertion/Feature Disjointness in which lexically arbitrary gender features are inserted together with phonological information at the end of the

derivation, and supports the claim that gender features are available relatively early in the derivation. The speech error data is much stronger and more compelling, in fact, than the surface data of normal German productions. This strikes me as a significant result; the error data goes beyond simply confirming conclusions that had already been forced by other facts, but actually help us decide between logically possible theories of "normal" productions.

I note in passing that the accommodation data shows that not just roots, but also derivational affixes must be specified for gender early in the derivation. The evidence for this comes from the fact that when two roots are exchanged and receive different nominal suffixes from the target sentence, the surrounding context reflects the gender of the new suffix, as the following example from Pfau (2000) shows:

(5)	Er	war	nur	darauf aus …		
	he	was	only	interested in		
	Intended:		sein	Bedürf-nis	zu	be-friedig-en
			his.NEUT	need(NEUT)	to	satisfy
	Error:		seine	***Befriedig***-ung	zu	be-***dürf***-en
			his.FEM	satisfaction(FEM)	to	need

Although the accommodation of the feminine gender introduced by *-ung* looks like a trivial variant of the more general case of gender accommodation, in fact it requires a different kind of look-ahead from a DM perspective. Here, a nominalization of √SATISFY is constructed, but the information that √SATISFY selects the morpheme /-ʊŋ/, and that /-ʊŋ/ is feminine, must all be determined before gender agreement takes place. Pfau recognizes this issue, and postulates that morpheme insertion must precede the feature-copying operation (section 5).[4] As with the insertion of gender on roots, this order of operations involves earlier integration of lexically arbitrary information than a strong version of Late Insertion would predict. However, once again, the speech error data provides stronger evidence in support of this complication than can be found in normal productions. The upshot of this section, then, is that Pfau's findings concerning the asymmetry in gender accommodation go beyond simply confirming a prediction of Distributed Morphology; they actually disconfirm a strong version of the theory, and inform our understanding of gender agreement.

2.2 What does lexical gender know about phonology?

The conclusion that lexically arbitrary gender information figures early in the derivation also has interesting implications for the relation between lexical gender and phonology. So far, we have focused on scenarios in which the gender feature is retrieved correctly, and the wrong root or wrong phonological string is retrieved. But what if it is the gender of a root that is incorrectly retrieved? Under the hypothesis that gender information is accessed after the semantic selec-

[4] Kornfeld (2004) makes a proposal that is similar in spirit, though mechanically different.

tion of roots but before the insertion of their phonological form, a strong prediction emerges: we expect possible gender errors due to activation of semantically related roots, but not based on activation of phonologically similar roots. That is, if a word is produced with the incorrect gender, this may be the result of the influence of semantically related words, but should not be influenced by the gender of phonologically similar words. As far as I know, this prediction has not been tested directly, but there are some hints that phonological form may indeed influence the processing of grammatical gender. In this section, I consider two such types of evidence.

2.2.1 Phonologically motivated gender assignment

The first source of data concerns the existence of "gender gangs" of phonologically similar words. It has often been noted that although it is impossible to predict the gender of a German noun with 100% accuracy, there are statistically significant tendencies that can make it at least somewhat possible to guess the gender of a noun. Zubin and Köpcke (1981, et seq.) demonstrate the existence of numerous predictors of noun gender in German, including both semantic generalizations (nouns denoting alcoholic beverages other than beer tend to be masculine) and phonological generalizations (nouns ending in (non-strident) fricative-stop clusters tend to be feminine). Both semantic and phonological correlates of grammatical gender are observed in other languages, as well (e.g., Serbian: Mirković, MacDonald and Seidenberg 2005). What is important for the present discussion is that these correlations are not just static facts about the language; speakers appear to have active knowledge of them, and they are reflected both in historical change (Zubin and Köpcke 1981) and also in behavior on novel words (Köpcke and Zubin 1983; Levine 1999; Schwichtenberg and Schiller 2004). This ability is often taken to indicate that speakers have something like a grammar of gender assignment, which uses semantic, morphological, and phonological properties of words to predict their lexical gender (Köpcke 1982; Corbett 1991).[5]

In the DM architecture that Pfau assumes, a grammar of gender assignment would have to operate under very severe restrictions. In particular, since the gender of a root or morpheme must be determined before its phonological features have been inserted, then one would not expect to be able to assign gender on the basis of phonological form. Empirically, this does not seem to be quite right. Although there indeed appears to be no language with gender assignment based exclusively on phonological form (Corbett 1991, p. 34), phonology does figure prominently in many gender systems—as seen, for example, in the wug tests of Köpcke and Zubin (1983) and others. This presents something of a paradox: the non-accommodation of form-based substitutions in speech errors seems to show that gender is accessed before phonological form, but the fact that gender assignment can make reference to phonological features seems to require the opposite. One possible way to reconcile these conflicting demands would be to suppose that form-based gender assignment does not reside within

[5] Naturally, this grammar cannot be 100% accurate for existing words, but presumably it contains a set of rules that make use of statistically significant correlations to assign the "regular" gender patterns, and the remaining cases rely purely on lexical knowledge.

the active production grammar, but rather, is part of a learning strategy for filling in gender features of unknown words. If this is on the right track, then we predict that speakers should never make production errors that involve phonologically-motivated "gender regularization" of known words—that is, accidentally using the gender that would be expected, given the phonological form of the word (e.g., **die Duft* 'the-FEM fragrance', where *Duft* is masculine, but expected to be feminine because of the final [ft] cluster). This implies that historical changes that bring individual words into line with their phonologically expected gender (e.g., German *die Frucht* 'fruit' (feminine) from older *der Frucht* (masculine), based likewise on the final fricative + stop sequence) would have to be attributed solely to incomplete or imperfect learning, and could not have received any help from regularizing speech errors in the input to subsequent generations. The prediction that gender errors outside of substitution contexts should be influenced only by semantics and not phonology could be quite difficult to test, given the overall rarity of such errors. However, to the extent that they exist at all, a systematic investigation of wrong-gender errors could be informative in determining to what extent semantic or phonological generalizations play an on-line role in filling in the gender features of nouns.

2.2.2 The interaction of semantic and phonological similarity

An even more fundamental prediction of the DM architecture is that phonological information should not be available to affect the semantic or syntactic selection of roots and morphemes, since phonological features are inserted late, during Vocabulary Insertion. In this section, I briefly mention two ways in which speech error data potentially show early phonological influences.

One plausible interaction between phonological form and morpheme selection can be observed in Pfau's example (30c), repeated in (6):

(6) Intended: schreibt man das mit Bindestrich?
write-3SG one that with connect-line(MASC)
'Is that written with a hyphen?'

Error: mit Binde***schrift***
with connect-writing(FEM)

Pfau asserts, quite reasonably, that the error is caused by accidentally inserting the root √SCHREIB 'write' in both the verb and object positions, creating the need for a nominal form of √SCHREIB. Furthermore, it is claimed that among the various possible nominalizations of √SCHREIB (*Schreibung* 'spelling', *Schreiber* 'writer', *Schrift* 'script'), *Schrift* is the best syntactic/semantic match to the intended stative DP structure that would have been appropriate for *Strich*. Pfau also admits, however, that the choice may also be influenced by the phonological similarity of *Strich* [ʃtrɪç] and *Schrift* [ʃrɪft]: both are monosyllabic, they have similar syllable structure, the same vowel, and so on. In other words, although the substitution of √SCHREIB for √STRICH must occur early in the derivation, it looks like the spell-out of √SCHREIB is being influenced by phonological similarity to the phonological form of √STRICH.

The DM architecture does not provide any obvious way for the phonological form of intended *Strich* to play a role, since even if √STRICH is activated as part of the competition for semantic selection of roots, its phonological form would not have been consulted yet. The phonological form of *Strich* can play a role only at the time of vocabulary insertion; yet, by hypothesis, √STRICH has been accidentally omitted from the syntactic structure during the initial root insertion, and there is no reason for it to be considered for vocabulary insertion. This means that there can be no real influence of phonological form on the choice of roots or morphemes; if the intended word happened to sound like *Schrift* but had agentive or eventive semantics, there would be no mechanism that would prefer *Schrift* just because of its phonological resemblance to intended *Strich*. The prediction of the theory, therefore, is that substitutions that occur early in the derivation should not be encouraged additionally by phonological similarity.

In point of fact, various authors have claimed that semantic substitutions *are* facilitated by phonological similarity (e.g., Dell 1986; Cutting and Ferreira 1999). For example, T. Harley and MacAndrew (2001) found that among errors involving semantic substitutions, there was a greater than chance likelihood that the words involved also share the same initial phoneme. They take as evidence for a more interactive model, in which phonological information is available already at the point of semantic selection. This result must be interpreted with care, however, because just as semantic substitutions may involve accidental phonologically similarity a certain proportion of the time, phonological substitutions may also occasionally involve accidental semantic similarity (see also Ferreira and Griffin 2003, and references cited therein, for discussion). In order to evaluate whether semantic and phonological similarity truly interact, we must use a baseline that takes into account not only the expected rate of chance phonological similarity, but also chance semantic similarity. This is certainly more difficult to estimate, and results so far from more carefully controlled comparisons seem mixed. Phonological similarity, at least in the extreme case of homophony, does appear to increase the number of semantically based substitutions (Ferreira and Griffin 2003), but the interaction may be limited, both in type and degree (Griffin 2002). The extent of interaction between form and meaning in lexical access is a matter of much on-going research and debate, and a full discussion of the issues is beyond the scope of this commentary. It should be clear, however, that results in this area bear on whether the model that Pfau assumes to explain the difference in gender accommodation between semantic and phonological substitutions can also account for the broader range of substitution data.

A second relevant finding comes from a study by Vigliocco et al. (2004), who found that the phonological form of the agreeing determiner appears to influence the rate of gender matching in semantic substitutions. Vigliocco et al. used an induced speech errors technique to elicit noun substitutions in a speeded picture naming task. They found, contrary to what one would expect given the identical gender effect on the top graph of Figure 1, that when naming items under time pressure ("cat," "desk," "zebra," …) German-speaking subjects were *not* significantly more likely to substitute words that shared the same gender as the target word. However, when the task was changed to require the definite article ($der_{[+masc]}$, $die_{[+fem]}$, or $das_{[+neut]}$) along with the noun, substitutions of same-

gender items did become more likely that different-gender items. Of greatest relevance to the current discussion, in a third task which required the indefinite article (*ein*$_{[+masc]}$, *eine*$_{[+fem]}$, *ein*$_{[+neut]}$), a tendency was found to avoid substituting feminine nouns for masculine or neuter nouns, but there was no reluctance to substitute masculine nouns for neuter nouns (or vice versa). In other words, in these studies, gender preservation is found only when a change in gender would also require a change in the determiner. Thus, it appears that the eventual surface form of the determiner is able to influence the likelihood of a given semantic substitution. This finding raises the possibility that the "identical gender effect" is actually an identical determiner effect (or, more generally, an identical frame effect).[6]

In fact, a similar effect may possibly be lurking in Pfau's own data. As it turns out, 63 out of 96 (= 66%) examples of substitutions that violate gender matching actually satisfy determiner-matching, either by having no determiner or by having an ambiguous determiner (Pfau's example (16)). Unfortunately, we cannot evaluate the relation between Vigliocco et al.'s finding and these corpus results without first knowing a baseline rate of how often nouns occur in unambiguously gender-marked contexts in German. However, this number implies at least the possibility that even semantic substitutions are sensitive to the surface morphological/phonological structure of the agreement context. The question, of course, is how this could happen if this information is available only later in the derivation.

Although I cannot offer a resolution here, when these facts are viewed from a DM perspective, they do suggest a possible hypothesis to pursue. The principle of Late Insertion does not lead one to expect purely phonological identity to influence lexical selection, but it is possible to imagine scenarios in which substitutions would emerge more readily if the masculine and neuter indefinite determiners (both *ein*) are morphologically identical—that is, are the very same vocabulary item, spelling out featurally identical contexts. For example, suppose that the morphological structures of *ein Mann* 'a-MASC man(MASC)' and *ein Haus* 'a-NEUT house(NEUT)' are identical: 'a$_{[-fem]}$ √X' or perhaps even 'a$_{[gender\ unspecified]}$ √X' (depending on the details of how gender features are marked). This underspecified structure would allow insertion of nouns of all compatible genders. If this is right, then the principles of underspecification and impoverishment (Halle 1997; H. Harley 2004) could also have an important role to play determining the shape of possible substitutions and accommodations. It also suggests an empirical test—namely, that frame matching effects should be found only in cases of true (principled) syncretism, and not in cases of accidentally homophonous determiners. Although distinguishing between these two cases is notoriously difficult problem, one could imagine a contrast between, say, the context *ein X* 'a X', in which *ein* is either masculine or neuter nominative, and *den X* 'the X', in which *den* could be either masculine accusative singular or the dative plural of any gender. Assuming for the sake of illustration that the former is a principled neutralization and the latter is an accidental ho-

[6] The definition of what constitutes a "frame" (determiners, agreement suffixes on the noun itself, on nearby adjectives, etc.) and how this plays a role has been a matter of considerable debates; see Schriefers, Jescheniak and Hantsch 2005 for discussion and references.

mophony, this leads to the hypothesis that masc./neut. substitutions should be possible with *ein*, but masc.sg./neut.pl. substitutions should not be possible with *den*.

To summarize, in this section we have seen two potential challenges to the idea that phonological information is unavailable during the semantic and syntactic planning stages: phonological facilitation for semantic substitutions, and facilitation by phonologically identical agreement frames. In both cases, further inquiry is needed to determine to what extent the data is compatible with a modular, unidirectional model such as Distributed Morphology, but the theory has the virtue of providing a number of clear and interesting predictions.

3. Is free too cheap?

In the previous section, I discussed several ways in which speech error studies such as Pfau's might help inform our understanding of the mechanics of how information about words is accessed, how gender agreement is carried out, and so on. A recurring theme has been that a formal, modular framework like DM makes usefully strong predictions about possible and impossible errors, and error data can be valuable in refining hypotheses about how particular parts of the system work. The other major issue I would like to discuss concerns the more basic claim that given the right theory of morphological derivation, no additional surface repair or "clean-up" operations are needed to turn unintended inputs into well-formed surface structures. The suggestion I would like to make is that this hypothesis may actually make repairs *too* easy, and that there may exist additional constraints that limit errors from making use of the full computational power of the morphological system.

Pfau's central thesis is that the normal morphological system is available to fix erroneous inputs, since at least some of them occur "upstream" of the computational system. Consider, for example, the following error from Fromkin (1973): *I think it's* ***care****ful to measure with* ***reason*** instead of intended … *reasonable to measure with care*. Here, the analysis is that a swap has taken place at the abstract root level, where the root √REASON has moved out of its degree phrase and switched places with √CARE, the complement of a preposition. The claim is that whereas the morphology would have taken the intended √REASON, [+deg] structure and spelled it out as *reasonable*, it takes the new √CARE, [+deg] structure and spells it out as *careful*. In this case, the switch from *-able* to *-ful* follows from the fact that the morphological rules for creating degree adjectives in English happen to have a special case that specifies that the root √REASON should receive *-able*, and a more general rule applying the productive affix *-ful* in other cases (including after √CARE). But what if the swap had involved a root that does not happen to have an existing degree adjective? For example, what if the intended utterance had been *I think it's reasonable to measure with gusto*? Here, the syntactic structure would involve a degree phrase √GUSTO, [+deg]. The general, productive affix for creating degree adjectives is *-ful*, and although *gustoful* is not an existing word for most speakers, it does

seem to be the preferred option for nonce adjectives based on *gusto*.[7] Ordinarily, morphology is (relatively) free to create new derivations as necessary, to fill the syntactic needs of the sentence. If repairs in speech errors are simply the ordinary workings of the morphological system, then this leads to the prediction that errors, too, should be free to productively create neologisms.

There is at least suggestive evidence from self-corrections that the system is not quite so simple, and that errors generally tend not to lead to productively formed neologisms. The following example (Pfau's (20c)) seems to be quite typical:

(7)	dass	ein	***Tänz***er,	äh,	dass	ein	Linguist	so	wild	tanzt…
	that	a	dancer,	um,	that	a	linguist	so	wildly	dance.3SG

In this example, the roots √LINGU(IST) and √TANZ have switched places, so √TANZ finds itself in an agentive noun position, which is spelled out with *-er* (productively, but also already lexicalized for this root). The root √LINGU(IST), by the same token, now finds itself in the position of a verb. Even if no such verb actually exists, in principle, it should be possible to add productive verbalizing morphology and spell out the structure (*linguistiziert*? *linguistiert*?). I conjecture that it is precisely the lack of a suitable existing verb that led the speaker to an immediate self-correction. This suggests that a strong lexical filter on possible affixal accommodations: morphological structures that do not correspond to pre-existing words tend to be recognized and blocked.

If this is right, it seems to imply a possible role for surface filters, after all. Certainly, the examples that Pfau discusses do show, quite clearly, that at least some errors occur very early in the computational process—early enough, in fact, to allow selection of different affixes in combination with different roots (Fromkin's *-able* vs. *-ful* example). This does not necessarily mean, however, that the system proceeds to operate blindly, taking a slightly incorrect input and running with it. In fact, a surface filter on pre-existing combinations seems to lead to a very common and powerful repair, of stopping and correcting the utterance. It is worth noting that the filter that attempts to block neologisms is almost certainly not specific to detecting errors. Even when the neologism is actually produced, speakers often seem to be conscious of the fact that they are saying a novel combination, and they frequently produce it with special prosody or comment on the fact that they are creating a new word. I should reiterate that I am not trying to argue for additional repair mechanisms that are unique to error situations or duplicate the workings of the morphological system. I am merely pointing out that errors seem to refrain from employing the full computational power of the system, and that some type of surface filter seems to be involved in blocking novel formations from getting spelled out.

[7] E.g., "What better a vehicle to explore the uncharted than the gustoful written word?" [http://plicity.org/?/unknown]

4. Errors in other types of operations?

All of the errors that Pfau discusses involve problems accessing lists (roots, or vocabulary items), or errors placing the retrieved items into the syntactic structure. It is useful to consider briefly what other types of errors a DM-style approach predicts, if something goes wrong elsewhere in the system. As it turns out, several other possible error types find plausible real-world examples to support them.

One place where errors seem likely to occur is at the point when morphemes are inserted into the morphological structure. One might, for example, fail to notice fast enough that a particular morpheme is specified to occur after a given root—for example, that √TAG 'day' is among the roots that the plural suffix [ə] is specified to attach to. In this case, incomplete retrieval of the specified contexts of a given rule could allow another, more general rule to accidentally apply. This yields what seems to be a correct prediction, that affixation errors should favor more general, default affixes. One source of evidence for this comes from a constant low-level rate of overregularization among not just children, but also adults. Fromkin (1971) cites the example … *and so, in conclude-**ment***, which could be interpreted as the productive suffix *-ment* stepping in when the suffix *-(t)ion* has failed to be inserted. Similarly, adults make occasional overregularization errors on (canonically) irregular past tense forms (Stemberger 1985). The rate of such errors in existing corpora is quite low, but it seems quite likely that their true rate of occurrence may actually be higher, either because they are hard to notice, or because it is difficult to be certain that it is an error from the point of view of the speaker. Rarely can we be confident that the speaker actually knows the lexicalized irregular form (i.e., *conclusion*), which makes it easy to underestimate the pervasiveness of such errors. Another reason to believe that such errors may actually occur with some frequency is that historical change tends to reflect regularization to a default form. If adult errors help to drive such changes, as if often thought to be the case, then we would expect to find non-negligible numbers of such errors. A desirable goal for future research, therefore, would be more careful documentation of such regularization errors, which are predicted by the theory, but which seem to be underrepresented in current error corpora.

DM also allows for lexically restricted phonological readjustment rules. This means that if the set of lexical items that are specified to undergo a particular readjustment is not retrieved completely, the process may underapply, exactly parallel to the case just discussed in which a lexically restricted morpheme fails to be inserted. Concretely, this predicts errors like **der Brech* instead of *der Bruch* 'the break', in which the [ɛ] → [ʊ] rule has failed to apply. These errors also seem to occur; they are seen, for example, in failures of Trisyllabic Shortening in English, creating errors such as *obsc*[iː]*nity* instead of *obsc*[ɛ]*nity*. I have occasionally observed such errors,[8] and they are also mirrored by the general trend of historical change (for example *obesity* has shifted from older *ob*[ɛ]*sity* to modern *ob*[iː]*sity*). Here too, such errors are not particu-

[8] For example, *obsc*[iː]*nity* occurred on Apr 12, 2004 on the NPR talk show *Brain Brew*. In this particular case, the error was immediately self-corrected, so we can be certain that it was not the usual pronunciation by this speaker.

larly well attested in existing error corpora. I suspect that this is either because they seem "practically correct" and thus pass by unnoticed, or else listeners may assume that the speaker has a slightly different lexical distribution of the process in question (i.e., it is easy to believe that a particular speaker always says *obsc*[iː]*nity*). It is possible that larger-scale corpora, in which errors are identified with the help of an automated speech recognition system, may help to arrive at a more accurate estimate of the rate of occurrence of such errors, and a better catalog of their precise content.

5. Conclusion

The goal of this commentary has been to highlight some ways in which a formal framework like DM, which involves a number of distinct modules and operations, can be a valuable tool in making predictions about possible error types, and conversely, how systematic studies of errors can uncover details of the workings of the system that are difficult to see from canonical productions. Pfau's study is an excellent demonstration of how results can be achieved in both directions. For the most part, the issues raised here do not challenge the basic claims, but merely point to predictions that await systematic attention. In many cases, it appears that existing corpora are inadequate to test these predictions, and further empirical work is needed, employing a combination of larger (automatically labeled) corpora, and experimental techniques.

References

Arnaud, Pierre (1999). Target-error resemblance in French word-substitution speech errors and the mental lexicon. *Applied Psycholinguistics* 20, 269–287.

Berg, Thomas (1992). Prelexical and postlexical features in language production. *Applied Psycholinguistics* 13, 199–235.

Carstens, Vicki (2000). Concord in Minimalist Theory. *Linguistic Inquiry* 31, 319–355.

Corbett, Greville (1991). *Gender*. Cambridge: Cambridge University Press.

Cutting, J. Cooper and Victor Ferreira (1999). Semantic and phonological information flow in the production lexicon. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 25, 318–344.

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

Embick, David (2000). Features, syntax, and categories in the Latin Perfect. *Linguistic Inquiry* 31, 185–230.

Ferreira, Victor and Zenzi Griffin (2003). Phonological influences on lexical (mis)selection. *Psychological Science* 14, 86–90.

Griffin, Zenzi (2002). Recency effects for meaning and form in word selection. *Brain and Language* 80, 465–487.

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

Fromkin, Victoria (1973). Introduction. In *Speech errors as linguistic evidence*, ed. V. Fromkin, 11–45. The Hague: Mouton.

Halle, Morris (1997). Distributed Morphology: Impoverishment and Fission. *MITWPL* 30, 425–439.

Halle, Morris and Alec Marantz (1993). Distributed Morphology and the pieces of inflection. In *The view from Building 20: Essays in linguistics in honor of Sylvain Bromberger*, ed. Kenneth Hale and Samuel Jay Keyser, 111–176. Cambridge, MA: MIT Press.

Harley, Heidi (2004). The importance of Impoverishment. Ms., University of Arizona.

Harley, Trevor and Siobhan MacAndrew (2001). Constraints upon word substitution speech errors. *Journal of Psycholinguistic Research* 30, 395–418.

Kornfeld, Laura (2004). Derivational processes and lexical categories. Ms., Universidad de Buenos Aires/CONICET.

Köpcke, Klaus-Michael (1982). *Untersuchung zum Genussystem der deutschen Gegenwartssprache*. Tübingen: Niemeyer.

Köpcke, Klaus-Michael and David Zubin (1983). Die kognitive Organisation der Genuszuweisung zu den einsilbigen Nomen der deutschen Gegenwartssprache. *Zeitschrift für germanistische Linguistik* 11, 166–182.

Levine, Glenn (1999). Gender assignment to German nonsense nouns: What does the native speaker know that the non-native speaker doesn't? *Proceedings of the 23rd Annual Boston University Conference on Language Development, Vol. 2*, ed. Annabel Greenhill, Heather Littlefield and Cheryl Tano, 397–406. Somerville, MA: Cascadilla Press.

Marantz, Alec (1995). A late note on Late Insertion. In *Explorations in Generative Grammar*, ed. Young-Sun Kim et. al., 396–413. Seoul: Hankuk Publishing Co.

Marx, Edeltraud (1999). Gender processing in speech production: Evidence from German speech errors. *Journal of Psycholinguistic Research* 28, 601–621.

Miozzo, Michele and Alfonso Caramazza (1997). Retrieval of lexical-syntactic features in tip-of-the tongue states. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 23, 1410–1423.

Mirković, Jelena, Maryellen C. MacDonald and Mark S. Seidenberg (2005). Where does gender come from? Evidence from a complex inflectional system. *Language and Cognitive Processes* 20, 139–167.

Müller, Gereon (2004). Syncretism and iconicity in Icelandic noun declensions: A Distributed Morphology approach. *Yearbook of Morphology 2004*, 229–271.

Pesetsky, David and Esther Torrego (2007). The syntax of valuation and the interpretability of features. In *Phrasal and clausal architecture: Syntactic derivation and interpretation—In honor of Joseph E. Emonds*, ed. Simin Karimi, Vida Samiian and Wendy K. Wilkins, 262–294. Amsterdam: Benjamins.

Pfau, Roland (2000). Features and categories in language production. Doctoral dissertation, Johann Wolfgang Goethe University, Frankfurt/Main.

Pfau, Roland (2007). Cheap repairs: A Distributed Morphology toolkit for sentence construction. This volume.

Sauerland, Uli (2005). On the semantic markedness ofφ φ-features. Ms., Zentrum für allgemeine Sprachwissenschaft, Berlin.

Schriefers, Herbert, Jörg Jescheniak and Ansgar Hantsch (2005). Selection of gender-marked morphemes in speech production. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 31, 159–168.

Schwichtenberg, Beate and Niels O. Schiller (2004). Semantic gender assignment regularities in German. *Brain and Language* 90, 326–337.

Stemberger, Joseph P. (1985). *The lexicon in a model of language production*. New York: Garland Publishing.

Vigliocco, Gabriella, David P. Vinson, Peter Indefrey, Willem J.M. Levelt and Frauke Hellwig (2004). The role of grammatical gender and semantics in German word production. *Journal of Experimental Psychology, Learning, Memory, and Cognition* 30, 483–497.

Zubin, David and Klaus-Michael Köpcke (1981). Gender: A less than arbitrary grammatical category. In *CLS* 17, ed. R. Hendrick, C. Masek and M. Miller, 439–449. Chicago: Chicago Linguistic Society.

MIT Linguistics, 32-D808
77 Massachusetts Ave.
Cambridge, MA 02139
USA

albright@mit.edu
http://www.mit.edu/~albright