J. Linguistics **43** (2007), 683–694. © 2007 Cambridge University Press
doi:10.1017/S002222670700480X Printed in the United Kingdom

NOTES AND DISCUSSION

Inherent variability and Minimalism: Comments on Adger's 'Combinatorial variability'

RICHARD HUDSON

(Received 12 February 2007; revised 30 April 2007)

Adger (2006) claims that the Minimalist Program provides a suitable theoretical framework for analysing at least one example of inherent variability: the variation between *was* and *were* after *you* and *we* in the Scottish town of Buckie. Drawing on the feature analysis of pronouns and the assumption that lexical items normally have equal probabilities, his analysis provides two 'routes' to *we/you was*, but only one to *we/you were*, thereby explaining why the former is on average twice as common as the latter. This comment points out four serious flaws in his argument: it ignores important interactions among sex, age and subject pronoun; hardly any social groups actually show the predicted average 2:1 ratio; there is no general tendency for lexical items to have equal probability of being used; the effects of the subject may be better stated in terms of the lexemes *you* and *we* rather than as semantic features. The conclusion is that inherent variability supports a usage-based theory rather than Minimalism.

1. Introduction

Are the findings of variationist sociolinguistics compatible with standard versions of Minimalism, as claimed in this journal by Adger (Adger 2006)? In 1986 one could write that 'none of these mainstream theories [including generative grammar] pays any attention whatsoever to what sociolinguists have been discovering' (Hudson 1986: 1053) and in 1995 that 'this work [on sociolinguistics] has had no influence at all on the most popular theories of syntax' (Hudson 1995: 1514). Indeed, some generative grammarians denied that variation could in principle be relevant to a theory of competence (N. Smith 1989: 178–186) and as recently as 1998 it was possible to write, in relation to variationist sociolinguistics and Chomskyan generative theory: 'There have been few real attempts to marry these seemingly divergent positions' (Wilson & Henry 2007). If this marriage is to take place, it is most likely to be arranged by experts in generative grammar, so we must welcome the recent flurry of attempts (Cornips & Corrigan 2000, 2005; Adger & Smith 2005; Henry 2005; Adger 2006; Parrott 2007; van Gelderen 2007).

Adger's paper is a particularly useful example of what Minimalism can offer because (unlike most other attempts) it at least tries to explain the

https://doi.org/10.1017/S002222670700480X Published online by Cambridge University Press

observed frequencies as well as the structural variation. To summarise Adger's analysis, it concerns the variation between *was* and *were* after *we* and *you* in the small Scottish town of Buckie (Smith 2000, Adger & Smith 2005). Adger reports that the ratio of *was* to *were* is 2:1 after *we* and singular *you*, and explains this ratio as a consequence of the internal organisation of Buckie grammar. For the past tense of BE with *we* or *you* (singular) as subject, this grammar provides three lexical items with different but overlapping feature structures. Two of these lexical items are pronounced *was*; only one is pronounced *were*. On the assumption that each item has an equal chance of being selected, this explains the observed 2:1 ratio in favour of *was* without, crucially, invoking probabilities inside the grammar. Moreover, Adger reports that Buckie is so small and isolated that speakers all show about the same tendencies in their speech. Consequently, these variable data can be explained, he claims, without including any social distinctions in the grammar.

If Adger is right, then at least one example of Labov's 'inherent variability' (Labov 1969), which he calls 'intra-personal morphosyntactic variation', can be explained in a Minimalist grammar. But is even this modest conclusion justified? Section 2 will point out four weaknesses of Adger's analysis, but the discussion will show that inherent variability can only be accommodated in a theory of language structure which can also accommodate usage information about typical speakers and about frequencies – a point that I have been arguing for some time (Hudson 1980: 188–190; 1986; 1995; 1996: 243–257; 1997; 2007a; 2007b: 246–248). If I am right, then Minimalism needs at least some fairly serious revision; but I have to leave it to the experts to decide whether the necessary changes are possible. I return in section 3 to a very brief outline of a radical alternative to Minimalism which is certainly compatible with inherent variability.

2. Flaws in the argument

Adger's analysis of *was*/*were* variation in Buckie fails for a number of reasons. Some of these are 'statistical' because they concern his interpretation of the observed figures, while others involve 'structural' issues in the grammar that he proposes.

2.1 *Statistics: age and sex have a significant influence*

Adger claims (p. 525) that

> *was*/*were* variability in Buckie is not affected by extra-linguistic factors in any clearly systematic way … each generation of speakers has a very similar statistical pattern for the use of *was*/*were* in the various person/number combinations, although it is true that the older generation have more *was* in general. As regards the gender of the language users, Smith

https://doi.org/10.1017/S002222670700480X Published online by Cambridge University Press

	You (singular)		We	
Speaker age	*You was*	*You were*	*We was*	*We were*
Old	45	5	113	36
Middle	23	12	32	41
Young	43	33	101	45
All	111 (=69%)	50	246 (=67%)	122

Table 1
Was or *were* after *you* and *we* for three age groups in Buckie

shows that it is only middle-aged females who have a markedly more standard pattern (although even these speakers show the same basic pattern of categoricality versus variability, and in fact show a broadly similar pattern of frequency distribution).

This generalisation allows him to ignore age and sex differences in a grammar for the whole community, thus avoiding one of the main challenges of inherent variability.

In fact, the generalisation is far from true. The raw data for *was*/*were* with *you* and *we* are shown in table 1 (from Smith 2000: 61). The average in the bottom line does indeed support Adger's generalisation that on average speakers use about two *was*s for every *were*; but for both subject types, the age differences are very highly significant ($p < 0.005$ in both cases), and the effect of the subject type is significant ($p < 0.05$) for both of the older age groups and almost significant ($p = 0.062$) for the young speakers. The figures are presented more accessibly in figure 1, which shows for each pronoun the number of *was* tokens as a percentage of the total of *was* plus *were*.

The most striking feature of figure 2 is the clear effect of age when the subject is *you*, with a general decline in the use of *was*. However, *we* produces a very different age effect, with a reverse of the decline between middle-aged and young speakers. As Adger notes, this is mainly due to the female middle-aged speakers shunning *we was*, in contrast with young males who increasingly revel in it (Smith 2000: 64); this complex pattern is shown in figure 2. Without real-time data it is not possible to know what changes have actually taken place to produce this effect, but one possibility is that middle-aged women have changed (considerably) within their own lifetime while young men have changed (slightly) in comparison with older men; and equally speculatively, it is possible that these changes were motivated by a desire on the part of these groups to distance themselves from one another – the more often young men said *we was*, the more often middle-aged women, who disapproved of this usage, said *we were*.

https://doi.org/10.1017/S002222670700480X Published online by Cambridge University Press

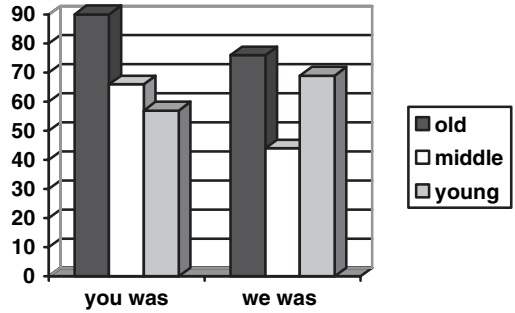


Figure 1
Was as a percentage of *was*/*were* after *you* and *we* for three age groups in Buckie

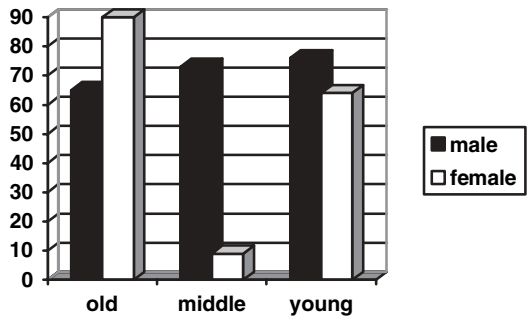


Figure 2
We was as a percentage of *we was*/*were* for three age groups and two sexes in Buckie

These figures clearly contradict Adger's claim that age and sex have so little influence on the *was*/*were* variable that they can be ignored, and more generally that the *was*/*were* variable has no 'social meaning' (p. 526, fn. 9) – i.e. that it carries no information about the speaker. But even if there are age and sex differences, it does not follow immediately that these differences must in some way be part of the grammar. Maybe it is possible to explain them in some other way. Let us consider the two most plausible alternative candidates: distinct grammars, and prescriptive knowledge (both suggested to me by David Adger, p.c.).

If middle-aged women had different grammars from the other speakers, this might indeed explain why their output frequencies are different. But what kind of grammatical difference could explain an output ratio of about 1:9 in favour of *were*? If the 2:1 ratio of other speakers reflects the number of distinct lexical items pronounced *was*, are we to believe that middle-aged women have nine different lexical items pronounced *were*? Even less plausibly, if they still have the same two *was*s as other speakers, we would

https://doi.org/10.1017/S002222670700480X Published online by Cambridge University Press

need no fewer than 18 distinct *were*s to explain the 1:9 ratio. This avenue looks unpromising, though it cannot be ruled out altogether.

Prescriptive knowledge is a more attractive explanation as there is some evidence that speakers are sufficiently aware of the difference between *we was* and *you was* to comment on it. One middle-aged woman who never used *we was* but used *you was* a great deal showed considerable self-awareness: 'Well, maybe I would say *you was* but I would never say *we was*. I ken that's just wrong …' (Jennifer Smith, p.c.). If this were true of all middle-aged women in Buckie, then it might well explain why they avoid *we was* even though it is allowed by their grammar on just the same basis as *you was*. Admittedly this answer would create another question: Why is it only middle-aged women who feel this way about *we was*? But this question could reasonably be left for sociologists and educationalists to worry about; the linguistic analysis, with its grammar insulated against social influences, would be saved.

The trouble with this argument is that it is no stronger than the arguments for excluding prescriptive knowledge from grammar. It is easy to exclude certain kinds of information from grammar by fiat, and one such exclusion could be applied to all 'social knowledge' (that *tummy* is for children, that *attempt* is formal, that *we was* is wrong, and so on). In this way it would be possible to rescue the boundary round grammar by distinguishing 'the grammar proper' from other kinds of knowledge such as the 'user's manual' (Culy 1996, Zwicky 1999); but it is easier to imagine such modules than to justify them either linguistically or psychologically.

One possible defence would be to argue that the objects in these two kinds of knowledge are in fact different (David Adger, p.c.): that grammar is about lexical entries such as '[*u*author:+] *was*' (p. 521), whereas prescriptive knowledge is about word combinations such as *we was*, which are not part of the grammar as such even if they are generated by it. But this defence fails for two reasons. First, prescriptive knowledge is not reserved for multiword phrases, but can be about single words (e.g. *ain't*). And second, the two-word description of *we was* implies an analysis in which *was* has the features [*u*author:+, *u*singular:−], which is sufficient to identify the subject as *we*, so the prescriptive knowledge may well apply to this single word; indeed, it is hard to see how '[*u*author:+] *was*' could be learned except by induction from a collection of more specific examples such as *we was* and *you was*. In short, prescriptive knowledge and grammatical knowledge may apply to the same objects.

Another kind of defence of a fundamental distinction between grammatical and prescriptive knowledge would be to demonstrate differences in how they are acquired or used; for example, if prescriptive knowledge was only learned at school and only used in school-like situations then the two might be assigned to different parts of our mind. But this defence faces two problems. One is that it leaves unanswered the question why the

https://doi.org/10.1017/S002222670700480X Published online by Cambridge University Press

school-teachers of Buckie picked out *we was* (but not *you was*) for special condemnation; if this was because they were middle-aged women, then the school teaching would be the result of the general trend, and not its cause. The other problem is that the ban on *we was* does not just apply in special situations; indeed, the woman who condemned it as 'just wrong' never used it even in casual family situations (Jennifer Smith, p.c.). A piece of knowledge which applies regardless of situation is hard to distinguish from an ordinary part of competence.

In conclusion, then, the Buckie data show that the choice between *was* and *were* is much less uniform than Adger assumes. Performance is heavily influenced not only by the subject (*you* or *we*) but also by the speaker's age and sex in a complex interactive pattern – a classic example of the micro-links between fine-grained linguistic structure and social structure which make inherent variability so fascinating and challenging for theoretical linguistics. These performance differences do not appear to be due to external influences such as prescriptive attitudes, so they must in some sense be part of speakers' competence.

2.2 *Statistics: few speaker groups show the predicted 2:1 ratio*

The second statistical weakness of Adger's analysis follows directly from the first. His claim that *was* is twice as frequent as *were* after both *you* and *we* is based on a gross averaging of the data. This average conceals the much more complicated picture in figure 1 and figure 2, where very few speaker groups show the 2:1 ratio which his analysis predicts, and where the patterns for *we* and *you* are rather different. Of course, it is possible that even these group averages conceal significant differences among individuals but individual figures are not available, and in any case, such differences would merely deepen the problem.

If the speakers all have the same grammar, and the grammar predicts a 2:1 ratio, then all the cases where the observed ratio is significantly different from 2:1 must either refute the analysis, or receive some kind of special explanation in terms of 'extra-linguistic' influences (such as the possible effects of prescriptivism that we considered briefly above). We cannot rule out the possibility of an extra-linguistic explanation, but it seems most unlikely that such an explanation is possible for the purely linguistic difference between the two subjects, *we* and *you*.

But if, on the other hand, the speakers have different grammars, what kinds of difference are possible? One possibility is that they have structural differences which predict the observed ratios, but, as I pointed out earlier, these are much less simple than the 2:1 ratio of Adger's explanation; for example, the middle-aged women's 90% (figure 2) implies nine lexical entries for *was* compared with only one for *were*. While worth pursuing, this analysis seems implausible. The other possibility is that the grammars differ

https://doi.org/10.1017/S002222670700480X Published online by Cambridge University Press

quantitatively – i.e. they contain the same elements, in the same structural relations, but with some kind of quantitative difference which affects their availability. Adger explicitly rejects this kind of analysis:

> In my model, the input probabilities for the sets of lexical items that constitute variants are all equal, but I agree with Bender that certain social and psychological factors may very well alter these input probabilities. I disagree that there's any need to build these social factors into the linguistic information in lexical items, however. (p. 527)

Nevertheless, it is hard to see a viable alternative to an analysis such as Bender's. Section 3 develops this point further.

2.3 *Structure: biased usage does not require inflection*

Adger's explanation for the variation between *was* and *were* relies on the assumption that every lexical item has an equal chance of being used (subject to external disturbances which lie outside the grammar). The case of *was/were* looks like an exception because *was* is used more often than *were*, but is actually regular because there happen to be two distinct lexical items that happen to be both pronounced *was* and that are both compatible with the inflectional features required with *you* and *we*. In other words, the only reason why usage is biased in favour of *was* is because it is part of a rather special inflectional paradigm.

The explanation is ingenious, but no evidence is offered for the underlying assumption that every lexical item has an equal chance of being used, which predicts that whenever two items share the same meaning, they should each have about 50% of the total usage. Common experience suggests that this is not so; for example, pairs of synonyms like *try* and *attempt* (as in *try/attempt to open the door*) offer speakers a lexical choice, but stylistic differences strongly favour *try* in ordinary casual conversation. Research evidence supports this conclusion. Take, for example, the pairs of compound pronouns such as *someone* and *somebody*, *no one* and *nobody*, and so on, which appear to be exactly synonymous. This choice tends to be resolved in favour of *-body* in casual speech and *-one* in formal writing; so *-body* words account for between 65% and 82% of the total in spoken corpora, compared with only 9%–48% (depending on formality) in written corpora (Jane Edwards, message 5.1196 on the Linguist List, 1994; Carter & McCarthy 2006: 14). In other words, where there is a free choice between two lexical items, the choice is likely to be biased in favour of one alternative, and to be strongly influenced by contextual factors such as formality and medium. In these cases there is no way to explain the bias in terms of inflectional features. (Nor, incidentally, can it be due to differences in the frequencies of the meanings concerned, contra to Newmeyer 2003, 2006; after all, the meaning of *somebody* is precisely the same as that of *someone*.)

https://doi.org/10.1017/S002222670700480X Published online by Cambridge University Press

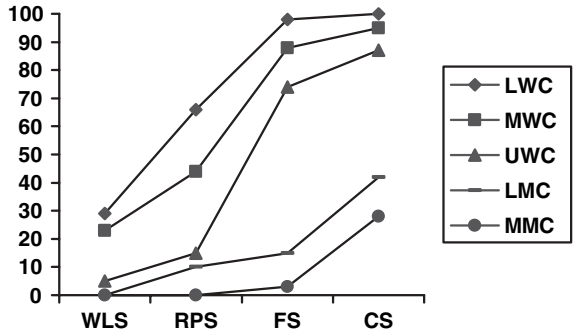


Figure 3
Alveolars as percent of all (*ing*) in Norwich (Trudgill 1974: 92)

This conclusion is typical of findings in quantitative sociolinguistics, where the data normally show context-sensitive bias in favour of one of two synonymous alternatives, and examples can be found in any introduction to variationist sociolinguistics (Hudson 1996, Coupland & Jaworski 1997, Trudgill 2000, Chambers, Trudgill & Schilling-Estes 2001, Wardhaugh 2005). To take one further example at random, the pronunciation of the suffix *-ing* varies between velar (*ing*) and alveolar (*in'*) in Norwich (as in every other dialect of English), with apparently continuous variation depending on social class (from Lower, Middle and Upper Working Class to Lower and Middle Middle Class) and style (from Word List Style, through Reading Passage and Formal Style, to Casual Style). Figure 3 shows a typical statistical pattern for variation, with a continuous range of percentages for the *in'* pronunciation from 0 to 100. Admittedly, each of these percentages aggregates the results for a number of speakers, but when individual performances are reported in other studies they tend to be broadly in line with the group scores; and in any case, the more individual variation there is, the more continuous the variation. This continuous variation, with contextual influences, is not at all what Adger's theory predicts.

Adger's theory therefore rests on the unsupported assertion that in general lexical choices are random: 'I have assumed that there is a random choice of lexical items (that is, that there is an equal probability that any of the three lexical items is chosen)' (p. 511). It is true that he accepts that usage itself can produce a bias, but for him this is an occasional perturbation of the basic pattern; thus the passage just quoted continues as follows:

> However, this is almost certainly not true in all cases. Choice of a lexical item by a speaker in any particular utterance is potentially influenced by social and/or psychological factors, so that a particular lexical item may

https://doi.org/10.1017/S002222670700480X Published online by Cambridge University Press

> have a higher probability of being chosen in a particular utterance (for example, if that lexical item has been recently accessed, it may be easier to access again; or if a lexical item is simply more frequent overall, it may be easier to access). (ibid.)

Section 3 sketches a theory in which usage (recency and frequency) is the only source of influence on lexical choices. In this theory, the Buckie pattern for *was*/*were* alternation is a self-perpetuating pattern of behaviour which needs no explanation in the grammar's structure.

2.4 *Structure: subject restrictions are lexical, not semantic*

My final criticism is a rather technical point about Adger's use of features for subject selection, using features such as '*u*singular', where *u* stands for 'uninterpretable'. For example, when the subject is singular *you*, the verb's features include: [*u*singular:+, *u*participant:+, *u*author:–] which are paired with the corresponding interpretable features [singular:+, participant:+, author:−] on the subject pronoun. Crucially, 'interpretable' features are interpretable in terms of meaning (p. 510), and not in terms of phonology or mere lexical identity; so the features identify the subject in terms of its meaning (e.g. *we* is a word that refers to one person who is a participant and who is not the 'author', i.e. the speaker) rather than identifying it directly as the pronoun *you* or *we*. This may seem harmless, but in fact it is potentially problematic because *you* is ambiguous, and has a 'generic' interpretation as in (1) (which might be addressed to a man) or (2), where the addressee could not possibly be the referent of *you*.

(1) When you're having a baby you have to push hard.
(2) In the old days, if they caught you stealing you were put in the stocks.

No information seems to be available about how such examples are treated in Buckie, but Adger's analysis predicts a completely different pattern of alternation between *was* and *were* because the verb *were* in (1)–(2) must be a different lexical item from the one in (say) *You were late last night*.

The alternative to Adger's analysis is one which identifies the subject in terms of lexical items rather than in terms of semantic features. In this kind of analysis, the two kinds of *you* could be unified under a single polysemous lexical item, and the prediction for *was*/*were* would be that the two uses of *you* would show the same alternation pattern. Of course, ultimately this is a matter of fact, and Adger's analysis may turn out to be vindicated by the facts, but the general approach that he offers predicts that contextual influences on variation can never be defined in purely lexical terms. This seems unlikely to be true.

https://doi.org/10.1017/S002222670700480X Published online by Cambridge University Press

3. Modularity or usage?

Adger summarises his argument in this way:

> If the approach I defend here is tenable, then we have a clear rapprochement between transformational generative grammar and variationist sociolinguistics. The grammar produces variants in a way that predicts particular probability distributions, but those probabilities can be perturbed at the point of use by factors such as ease of lexical access, recency effects, metalinguistic or social judgements on the form, etc. (p. 506)

I have argued that this approach is not in fact tenable, so it does not represent a 'rapprochement between transformational generative grammar and variationist sociolinguistics'.

Indeed, I believe that such a rapprochement is inherently impossible in any theory which assumes that language is a 'purely linguistic' module free of all social and usage information. Adger rightly sees this as a major issue when he rejects Bender's analysis (Bender 2000) in terms of 'social meaning' in the grammar:

> This is an intriguing position, but one which I wholly reject, mainly for broader reasons of modularity, dissociation of linguistic and social skills, etc. (p. 525)

What is at stake is the fundamental claim that language is a distinct mental module which cannot include 'non-linguistic' information. The findings of variationist sociolinguistics offer a particularly strong prima facie challenge to this claim, so the debate about the *was/were* alternation in Buckie takes us right to the heart of linguistic theory.

The modularity claim is hard to evaluate in detail without a clear specification of what it entails. What, precisely, is it that is claimed to be modular in language: its structure, its storage, its acquisition or its processing? Each of these properties of language entails a different theory of modularity, and any one of them would explain some of the dissociations that Adger mentions, but they clearly have very different implications for sociolinguistics. For example, if the modularity only applied to storage or processing, the structure of language could be intimately connected with social (and other) knowledge, so that Bender's social meaning could be part of a lexical entry, and indeed the grammar might even include frequency and recency effects. But if modularity applies to language structure, as Adger seems to assume, this will not be possible. In that case, of course, there will be serious questions about how properties of general memory such as the effects of frequency can apply to elements of language structure such as individual lexical items. The debate is clearly important but progress is unlikely without more clarity.

https://doi.org/10.1017/S002222670700480X Published online by Cambridge University Press

What, then, do I recommend instead? The data from inherent variability converge with a great deal of other evidence in support of a 'usage-based' account of language learning (Barlow & Kemmer 2000, Hudson 2007b: 52–59) in which syntactic patterns (among others) are learned inductively on the basis of experience, with a great deal of very specific information stored in memory about patterns such as subject–verb pairs (Goldberg 2006). This is exactly what the Buckie data show: that memorised items include *we was*, *you was*, *we were* and *you were*. Moreover, in this usage-based account, our memories of tokens may include their contextual specifics such as who uttered them and when (Bender's 'social meaning'), and will certainly reflect their frequencies in usage – the two crucial elements in any account of inherent variability.

Fortunately, however, we do not face a choice between explaining dissociations and explaining socially determined variation. Even a usage-based theory allows different kinds of knowledge to be dissociated because the total network of knowledge is by no means undifferentiated; but these dissociations arise largely from the structure of experience rather than from imposed 'modules' of the mind. And most importantly for the project of integrating the data of sociolinguistics into linguistic theory, it allows the elements of I-language to relate freely to those of what we might call 'I-society'.

REFERENCES

Adger, David. 2006. Combinatorial variability. *Journal of Linguistics* 42, 503–530.
Adger, David & Jennifer Smith. 2005. Variation and the Minimalist Programme. In Cornips & Corrigan (eds.), 149–178.
Barlow, Michael & Suzanne Kemmer. 2000. *Usage based models of language*. Stanford: CSLI.
Bender, Emily. 2000. *Syntactic variation and linguistic competence: The case of AAVE copula absence*. Ph.D. dissertation, Stanford University.
Carter, Ronald & Michael McCarthy. 2006. *Cambridge grammar of English: A comprehensive guide. Spoken and written English grammar and usage*. Cambridge: Cambridge University Press.
Chambers, Jack, Peter Trudgill & Natalie Schilling-Estes. 2001. *The handbook of language variation and change*. Oxford: Blackwell.
Cornips, Leonie & Karen Corrigan. 2000. Convergence and divergence in grammar. In Peter Auer, Frans Hinskens & Paul Kerswill (eds.), *Dialect change: Convergence and divergence in European languages*, 96–134. Cambridge: Cambridge University Press.
Cornips, Leonie & Karen Corrigan. 2005. *Syntax and variation: Reconciling the biological and the social*. Amsterdam: John Benjamins.
Coupland, Nikolas & Adam Jaworski. 1997. *Sociolinguistics. A reader and coursebook*. London: Macmillan.
Culy, Christopher. 1996. Null objects in English recipes. *Language Variation and Change* 8, 91–124.
Goldberg, Adele. 2006. *Constructions at work: The nature of generalization in language*. Oxford: Oxford University Press.
Henry, Alison. 2005. Idiolectal variation and syntactic theory. In Cornips & Corrigan (eds.), 109–122.
Hudson, Richard. 1980. *Sociolinguistics*, 1st edn. Cambridge: Cambridge University Press.
Hudson, Richard. 1986. Sociolinguistics and the theory of grammar. *Linguistics* 24, 1053–1078.

https://doi.org/10.1017/S002222670700480X Published online by Cambridge University Press

Hudson, Richard. 1995. Syntax and sociolinguistics. In Joachim Jacobs, Arnim von Stechow, Wolfgang Sternefeld & Theo Venneman (eds.), *Syntax: An international handbook*, vol. 2, 1514–1528. Berlin: Walter de Gruyter.
Hudson, Richard. 1996. *Sociolinguistics*, 2nd edn. Cambridge: Cambridge University Press.
Hudson, Richard. 1997. Inherent variability and linguistic theory. *Cognitive Linguistics* 8, 73–108.
Hudson, Richard. 2007a. English dialect syntax in Word Grammar. *English Language and Linguistics* 11, 383–405.
Hudson, Richard. 2007b. *Language Networks: The new Word Grammar*. Oxford: Oxford University Press.
Labov, William. 1969. Contraction, deletion and inherent variability of the English copula. *Language* 45, 715–762.
Newmeyer, Frederick J. 2003. Grammar is grammar and usage is usage. *Language* 79, 682–707.
Newmeyer, Frederick J. 2006. On Gahl and Garnsey on grammar and usage. *Language* 82, 399–404.
Parrott, Jeffrey. 2007. *Distributed morphological mechanisms of Labovian variation in morphosyntax*. Ph.D. dissertation, Georgetown University.
Smith, Jennifer. 2000. *Synchrony and diachrony in the evolution of English: Evidence from Scotland*. Ph.D. dissertation, University of York.
Smith, Neil. 1989. *The twitter machine: Reflections on language*. Oxford: Blackwell.
Trudgill, Peter. 1974. *The social differentiation of English in Norwich*. Cambridge: Cambridge University Press.
Trudgill, Peter. 2000. *Sociolinguistics: An introduction to language and society*, 4th edn. London: Penguin.
van Gelderen, Elly. 2007. *Principles and parameters in change*. In Cornips & Corrigan (eds.), 179–198.
Wardhaugh, Ronald. 2005. *An introduction to sociolinguistics*, 5th edn. Oxford: Blackwell.
Wilson, John & Alison Henry. 2007. Parameter setting within a socially realistic linguistics. *Language in Society* 27, 1–21.
Zwicky, Arnold. 1999. The grammar and the user's manual. Unpublished. http://www-csli.stanford.edu/~zwicky/forum.pdf (27 April 2007).

Author's address: Richard Hudson, dick@ling.ucl.ac.uk

https://doi.org/10.1017/S002222670700480X Published online by Cambridge University Press