Does contextual strength modulate the subordinate bias effect? A reply to Kellas and Vu

KATHERINE S. BINDER

Mount Holyoke College, South Hadley, Massachusetts
and

KEITH RAYNER

University of Massachusetts, Amherst, Massachusetts

In their reply to Binder and Rayner (1998), Kellas and Vu (1999) raised questions about the criteria we used to exclude items from the Kellas, Martin, Yehling, Herman, and Vu (1995) stimulus set. In this reply, we further document these criteria and also address the issue of local versus published norms. We continue to believe that the stimulus set used by Kellas et al. (1995) was problematic. We also address the issue of strength of context, a concept used in earlier research that dealt with the subordinate bias effect. We argue that the contexts used by Kellas et al. (1995) were no stronger than the contexts previously used that established this effect. Therefore, we continue to think that our finding that context does not eliminate the subordinate bias effect is valid.

Over the past few decades, there have been many investigations into the processes surrounding the resolution of lexically ambiguous words (for a review, see Simpson, 1994). Investigators have concluded that these processes are greatly influenced by two factors: relative meaning frequency and contextually biasing information (Binder & Morris, 1995; Dopkins, Morris, & Rayner, 1992; Duffy, Morris, & Rayner, 1988; Kellas, Paul, Martin, & Simpson, 1991; Neill, 1989; Neill, Hilliard, & Cooper, 1988; Paul, Kellas, Martin, & Clark, 1992; Rayner, Pacht, & Duffy, 1994; Simpson, 1984, 1994). Two models, which are quite similar, have been proposed to account for the data that led researchers to this conclusion. These models are the context-sensitive model (Kellas et al., 1991) and the reordered access model (Duffy et al., 1988). According to both models, relative meaning frequency and contextual bias are important variables in the resolution of lexical ambiguity. In our laboratory, we have recorded the eye movements of readers while they were reading sentences that contained lexically ambiguous words. Two important findings have emerged. First, when readers encounter a balanced ambiguous word (a word with two equally frequent interpretations) in a neutral context, they look at that word longer

than at a control word that is matched on length and frequency. Second, when readers encounter a biased ambiguous word (a word with one highly dominant interpretation) in a context that biases a reader toward its subordinate interpretation, readers look longer at that word than at a matched control word. The latter finding has been termed the *subordinate bias effect* (Pacht & Rayner, 1993; Rayner et al., 1994).

In contrast to our results, Kellas, Martin, Yehling, Herman, and Vu (1995) found that a strong biasing context led to the elimination of the subordinate bias effect. The discrepancy between the findings had generally been attributed to task differences: Kellas et al. (1995) used naming, self-paced reading, and the Stroop task; we relied on eye movement measures. In order to substantiate the claim that the discrepancy was the result of task differences, we obtained the set of materials used by Kellas et al. and conducted two experiments to pinpoint the discrepancy (Binder & Rayner, 1998). We examined both eye movements and self-paced reading times, and we obtained the subordinate bias effect whether biased ambiguous words were embedded *either* in strong or in weak contexts. Therefore, we (Binder & Rayner, 1998) argued that strongly biasing contextual information does not produce selective activation of the subordinate interpretation of a biased ambiguous word. Kellas and Vu (1999) have expressed some concerns about our study, and we will respond to their concerns in this article.

Kellas and Vu were concerned about why we eliminated certain items from their stimulus set. The reasons for elimination, as we outlined (Binder & Rayner, 1998), were as follows: (1) both the subordinate and dominant biasing contexts biased the same meaning, (2) the dominant and subordinate meanings were reversed, and (3) the ambiguous words were balanced. We will discuss each of these problems in turn, but it is important to note that if the eliminated items had been used in the stimulus set, each would have reduced the subordinate bias effect, but not because of a strong biasing context.

With respect to Point 1, three of Kellas et al.'s (1995) items (see Table 1) instantiated either the same meaning or a sense variation of the same meaning, despite what Kellas and Vu claim in their reply. For example, the meaning of the word *date* does not truly differ in "He forgot the *date* on the final exam" and "He forgot the *date* of the birthday party." The same word, when used to refer to the fruit known as a *date*, does have a different meaning.¹ Likewise, when we initially evaluated Kellas et al.'s (1995) items, we felt that the contexts for *field* and *trade* were both sense variations of the same meaning—that is, both contexts for *field* biased an area of space, and both contexts for *trade* biased an act of commerce. In the Twilley, Dixon, Taylor, and Clark (1994) norms, the only associate listed for *field* is *dreams*, which obviously refers to an area of space, as in the movie *Field of Dreams*. In addition, ac-

Preparation of this article was supported by Grant HD17246 from the National Institutes of Health and by a Research Scientist Award (MH01255) to the second author. We thank Susan Duffy, Chuck Clifton, and Greg Simpson for comments on an earlier version of this paper. Requests for reprints should be addressed to K. S. Binder, Department of Psychology and Education, Mount Holyoke College, 50 College St., South Hadley, MA 01075-1462 (e-mail: kbinder@mtholyoke.edu).

Table 1
Contexts From Kellas et al. (1995) That Biased the Same Meaning

Ambiguous Word	Dominant Biasing Context	Subordinate Biasing Context
date	The student made an error. He forgot the date on the final exam.	The boyfriend was thoughtless. He forgot the date of the birthday party.
field	The planter sprayed the crops. He rechecked the field for signs of insects.	The referee miscounted. He rechecked the field for the extra player.
trade	The principal overheard the plan. He banned the trade between the two students.	The king declared war. He banned the trade with the rival country.

cording to the Twilley et al. norms, the word *trade* has two associated meanings listed: *barter* and *mark*. *Mark* refers to *trademark*, which is clearly just a free association and not a distinct meaning.² Most critically, Frazier and Rayner (1990) demonstrated that a different pattern of results emerges for sense ambiguities (such as the two meanings of *field* used by Kellas et al., 1995) when they are compared with true lexical ambiguities. Therefore, we believe that we were justified in throwing out these three items because the contexts for them instantiated the same meaning or because they had only ambiguities of sense, rather than true ambiguities of meaning.

With respect to Point 2, we found seven homonyms in which the dominance was reversed (a dominant meaning according to our local norms and/or the Twilley et al., 1994, norms was counted as a subordinate meaning, and vice versa). For example, consider the contexts for the ambiguous word *pot*:

Dominant biasing context

The addict saw the police. He burned the pot to destroy the evidence.

Subordinate biasing context

The husband roasted the chicken. He burned the pot on the back burner.

According to both the Twilley et al. norms and our local norms, the ambiguous word *pot* is a biased ambiguous word. However, according to both sets of norms, the dominant meaning is associated with *pans*, and the subordinate interpretation is associated with *drugs* (a complete listing of these items is in Table 2). Clearly, the dominance was reversed for this item, as well as for the other six.

Thus, these items, in their original state, were invalid for our subjects.

Finally, with respect to Point 3, we eliminated 14 items that were classified as balanced ambiguous words according to the Twilley et al. (1994) norms or local norms. Kellas and Vu admitted that for their experiment they "deliberately selected [items with] a wide range of meaning frequencies." Unfortunately, as a result of doing that, they failed to examine the subordinate bias effect (which by definition deals with the subordinate interpretation of a biased ambiguous word). Previous eye movement experiments have demonstrated that when the biasing context precedes a balanced ambiguous word, the processing time for the ambiguous word is not different from the processing time for an unambiguous control word (Duffy et al., 1988). Obviously, to understand the subordinate bias effect, we must examine the processing of *biased* ambiguous words. Thus, given that Duffy et al. used a .69/.31 criterion, and Kellas and colleagues were challenging the Duffy et al. model, it made sense to us to use the same standards.³ In addition, the strongest possible test of whether or not the context overrides the meaning frequency would be to use words that are strongly biased toward one meaning.

At the heart of both Points 2 and 3 is the issue of the appropriateness of using local instead of published norms. We have thoroughly examined our local norms in conjunction with the Twilley et al. (1994) norms and the Nelson, McEvoy, Walling, and Wheeler (1980) norms, and it is quite clear to us that the UMass norms are more consistent with the Twilley et al. norms than they are with the Nelson et al. norms. We base this statement on our examination of 70 items that were in all three sets of norms.

Table 2
Items From Kellas et al. (1995) in Which the Meaning Dominance Was Reversed

Ambiguous Word	Subordinate Biasing Context*	Associate for Subordinate Meaning†
plot	The writer changed the setting. He developed the plot to reveal the victim.	land
pot	The husband roasted the chicken. He burned the pot on the back burner.	drugs
range	The gunner saw the target. He adjusted the range before firing the cannon.	stove
bar	The owner made a decision. He painted the bar during the slow season.	metal
screen	The usher saw the hole. He examined the screen after the patrons left.	door
tip	The man pondered the situation. He considered the tip to be sound advice.	money
trace	The student examined the fluid. He uncovered a trace of the deadly poison.	trail

*The subordinate biasing context was the context that was provided by Kellas et al. (1995). †The associate for subordinate meaning was the associate that was provided by our local norms.

When we correlated our local norms with the norms of Twilley et al., we obtained a correlation of $+.83$. However, when we performed the same analysis on our norms and the Nelson et al. norms, the correlation coefficient was $+.61$ (also see Table 3 for a comparison of the items that were eliminated across the three sets of norms). Given the apparent differences in meaning dominance, it is clearly important to obtain local norms. The fact that some norms have been published does not seem to be a compelling reason to use them if they are not consistent with the interpretations used by subjects in a study. Specifically, it makes no sense to use items that are not consistent with local lexicons.

In addition, norms should be as current as possible.⁴ For exactly this reason, we have continually updated our norming database. This fact accounts for possible discrepancies in which items are used as balanced or biased words across various studies in our lab. Although Kellas and Vu raised some important questions about local norms, it is not clear to us that the use of published norms is the automatic answer. We agree with Kellas and Vu that published norms have certain advantages over local norms and believe that it is preferable to use items that have consistency across published and local norms. However, when the published norms for items are inconsistent with the local norms, we believe that reliance on the local norms makes more sense. We also note that we have typically reported relevant information about the norming process we have used and have also published the items in an appendix (see Duffy et al., 1988; Rayner & Duffy, 1986; Rayner & Frazier, 1989).

As stated above, the original Kellas et al. (1995) items were inappropriate for our subjects, based on our local norms. How can we know whether the same items that were problematic for us were, in fact, problematic for the Kellas et al. subjects (students at the University of Kansas)? Since we do not know how the Nelson et al. (1980) norms compare with the local lexicons in Kansas, we can use the

following line of reasoning to evaluate the appropriateness of the Nelson et al. norms for the Kansas subjects. When we initially tried to replicate the Kellas et al. finding with the original stimulus set, we obtained a pattern of results that was consistent with the one that they obtained (see Binder & Rayner, 1998). When we subsequently eliminated the items that were problematic according to our local norms or to the norms of Twilley et al. (1994), we no longer obtained evidence of the selective access of the subordinate interpretation of a biased ambiguous word. Therefore, we continue to believe that the discrepancy between the two findings is attributable to the stimulus set.⁵

Finally, Kellas and Vu appear to be confused about how we selected the stimuli that were used in our study. Initially, they provided us with 56 items, all of which contained an ambiguous target word. However, we felt that in order to have a stronger design, we should have a neutral baseline in the form of an unambiguous control word. Kellas et al. (1995) argued that a strongly biasing context can produce selective access of a subordinate meaning. That is, only one meaning (i.e., the subordinate interpretation) is ever accessed in this condition. Thus, it makes sense to compare this condition with one in which only one meaning is accessed—namely, an unambiguous control word.⁶ Therefore, we created a control condition for all 56 items. However, when one uses an unambiguous control word, the word must be matched along various dimensions, such as length, frequency, and the word's fit with the context. Since the Kellas et al. items were not originally written for control words to fit into the context, some items worked better than others; thus these factors (length, frequency, and fit with context) were taken into consideration when we selected the final stimulus set. As noted above, there were instances in which Kellas et al. had the dominance reversed for our subjects (a dominant meaning according to our norms was counted as a subordinate meaning, or vice versa), but the items had control words that were matched quite well on length, frequency, and fit with context. In these instances, we simply coded the item according to our norms (i.e., the dominant biasing context was, in fact, biasing the dominant interpretation, and the subordinate biasing context was biasing the subordinate meaning), and then used it in the study. Thus, despite the fact that an item's meaning dominance was initially coded inappropriately for our subjects, we found no reason to eliminate the item, given that we could easily re-code it and use it in our study.

We will now address Kellas and Vu's comments and concerns about context. They argued that the point of distinction between the reordered access model and the context-sensitive model lies in the relative importance of meaning frequency and context. They suggest that whereas the reordered access model maintains a special status for meaning frequency, the context-sensitive model maintains a special status for context. So it was reasonable for Kellas and colleagues to predict that when readers were provided with a context that strongly biased the subordinate mean-

Table 3
Relative Meaning Frequencies for the Dominant Meaning for Items That Were Eliminated Because They Were Balanced

Ambiguous Word	Nelson et al. (1980)	Twilley et al.	
		(1994)	UMass
court	.80	.59	.55
crane	—	.52	.63
draft	.92	.52	.57
express	.85	.65	—
fan	.82	.80	.64
marble	.67	.59	—
speaker	.65	.65	.55
stock	.77	.50	.53
terminal	.98	.55	.58
bark	.84	.72	.58
board	.77	.91	.65
bow	.79	.57	.52
tag	.78	.51	—
range	.60	.57	.58
<i>M</i>	.79	.61	.58

ing of a biased ambiguous word, the readers would selectively access the subordinate meaning of that word. Thus, they argued that the discrepancy between their findings (Kellas et al., 1995) and those obtained by Duffy et al. (1988) occurred because their contexts were more strongly biased than those used by Duffy et al. Given that their question is an empirical one, we decided that the best thing to do was to collect the appropriate norming data. Since we had already collected context strength norming data on the Kellas et al. materials,⁷ we subjected the Duffy et al. contexts to a norming procedure that was identical to the one that we used to collect the contextual strength norming information on the Kellas et al. materials.

It is important to note that the results of these norming tasks demonstrate that the contexts from the Duffy et al. (1988) items were *not* rated as being weaker in contextual strength in comparison with the Kellas et al. (1995) materials. Rather, the Duffy et al. items received *stronger* ratings than did the Kellas et al. items (5.9 vs. 5.1, respectively, on a 7-point scale). However, Kellas and Vu noted that we did not do the norming task the same way they did, and they suggest that the procedure we used may have inflated the strength ratings for the Duffy et al. items. Specifically, they take us to task for providing the ambiguous word along with a short definition of the meaning that was intended by the sentence before the presentation of the actual item that was being rated. They argue that the definition provided a "framework for interpreting the ambiguous word," and thereby inflated the strength rating. In order to avoid this problem, they provided subjects with a 9-point scale,⁸ anchored on each side with an associate for both the dominant and subordinate interpretation. It is unclear how and why the associates that they provided are different from the definitions that we provided. Therefore, we are not quite sure why Kellas and Vu argue that our provision of the definitions influenced the norming procedure.

Regardless of the methodological differences in the collection of the norming data, the fact remains that we collected data for *both* sets of stimuli, using the *same* procedure for both sets. Thus, a direct comparison between the ratings we obtained for the Duffy et al. (1988) items and the ratings we obtained for the Kellas et al. items was warranted. Perhaps our procedure did inflate the strength ratings, but if it did, it should have done so for *both* sets of items. There is absolutely no reason that the procedure would inflate the ratings of the Duffy et al. items and not of the Kellas et al. items. Further, when we initially collected the norming data for the Kellas et al. items only, we replicated their results with our procedure. That is, items that they claimed were strongly biasing when their method was used were also strongly biasing when our method was used, and items that they had rated as weakly biasing were weakly biasing by our standards as well. However, given Kellas and Vu's concern with the procedure we used to obtain the norming data, it is unclear, and very puzzling, why they did not simply collect their own norming data, using

their preferred procedure on the Duffy et al. materials, given that all of the items are published in the appendix of that paper.

Kellas and Vu argue, on a different level, that it is the structure of their passages that makes them more biasing than those used by Duffy et al. (1988). Their items are made up of two sentences. The first sentence consists of a character performing some action or a character with some described attribute. The second sentence contains a pronoun that refers to the character in the first sentence, and it also contains the ambiguous word. The context of the second sentence is neutral; thus, they argue that it is the process of assigning the pronoun to its referent that produces the strongly biasing context. Although the Duffy et al. items each consisted of a single sentence, the sentences were made up of two different clauses. The initial clause contained the contextually biasing information, and the second clause contained the target word. Thus, the only concrete difference between the two sets of stimuli is that the items constructed by Kellas et al. (1995) use a period to separate the biasing context from the ambiguous word, whereas a comma is used to separate this information in the Duffy et al. items. Yet, Kellas and Vu continue to argue that in terms of on-line processing, the contexts of Duffy et al. are less potent because they are presented in the first clause of a two-clause sentence, instead of in a separate sentence (as is the case in their materials). It would seem to us that information in a previous sentence would be more inaccessible than information in a previous clause. Therefore, it is not clear why Kellas and Vu believe that, structurally, their contexts are more biasing.⁹

Given the results of our two experiments (Binder & Rayner, 1998), we are left with the same conclusion that we reached before: The subordinate bias effect is not eliminated in the face of strongly biased context. The eye movement experiment and the self-paced reading experiment tell the same story. We have eliminated experimental method as the cause of the difference. Furthermore, we have verified that the contexts used by Duffy et al. (1988), who first observed the subordinate bias effect, are stronger than those used by Kellas et al. (1995). We are left with one inescapable conclusion: The materials set used by Kellas et al. was problematic. Finally, we believe that if it were convincingly demonstrated that the subordinate bias effect was eliminated by strong context, the reordered access model could accommodate the finding (see Rayner, Binder, & Duffy, in press, for further discussion). But, the important empirical point at the moment is that the effect is not eliminated by strong contextual constraint. We continue to think that this finding is important.

REFERENCES

- BINDER, K. S., & MORRIS, R. K. (1995). Eye movements and lexical ambiguity resolution: Effects of prior encounter and discourse topic. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 21, 1186-1196.
- BINDER, K. S., & RAYNER, K. (1998). Contextual strength does not mod-

- ulate the subordinate bias effect: Evidence from eye fixations and self-paced reading. *Psychonomic Bulletin & Review*, 5, 271-276.
- DOPKINS, S., MORRIS, R. K., & RAYNER, K. (1992). Lexical ambiguity and eye fixations in reading: A test of competing models of lexical ambiguity resolution. *Journal of Memory & Language*, 31, 461-476.
- DUFFY, S. A., MORRIS, R. K., & RAYNER, K. (1988). Lexical ambiguity and fixation times in reading. *Journal of Memory & Language*, 27, 429-446.
- FRAZIER, L., & RAYNER, K. (1990). Taking on semantic commitments: Processing multiple meanings versus multiple senses. *Journal of Memory & Language*, 29, 181-200.
- KELLAS, G., MARTIN, C., YEHLING, K., HERMAN, R., & VU, H. (1995). *Contextual strength as a determinant of the subordinate-bias effect*. Poster presented at the 36th Annual Meeting of the Psychonomic Society, St. Louis.
- KELLAS, G., PAUL, S. T., MARTIN, M., & SIMPSON, G. B. (1991). Contextual feature activation and meaning access. In G. B. Simpson (Ed.), *Understanding word and sentence* (pp. 47-71). Amsterdam: Elsevier, North-Holland.
- KELLAS, G., & VU, H. (1999). Strength of context *does* modulate the subordinate bias effect: A reply to Binder and Rayner. *Psychonomic Bulletin & Review*, 6, 511-517.
- NEILL, W. T. (1989). Lexical ambiguity and context: An activation-suppression model. In D. S. Gorfein (Ed.), *Resolving semantic ambiguity* (pp. 63-83). New York: Springer-Verlag.
- NEILL, W. T., HILLIARD, D. V., & COOPER, E. (1988). The detection of lexical ambiguity: Evidence of context sensitive parallel access. *Journal of Memory & Language*, 27, 279-287.
- NELSON, D. L., MCEVOY, C. L., WALLING, J. R., & WHEELER, J. W., JR. (1980). The University of South Florida homograph norms. *Behavior Research Methods & Instrumentation*, 12, 16-37.
- PACHT, J. M., & RAYNER, K. (1993). The processing of homophonic homographs during reading: Evidence from eye movement studies. *Journal of Psycholinguistic Research*, 22, 251-271.
- PAUL, S. T., KELLAS, G., MARTIN, C., & CLARK, M. B. (1992). Influence of contextual features on the activation of ambiguous word meanings. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 18, 703-717.
- RAYNER, K., BINDER, K. S., & DUFFY, S. A. (in press). Contextual strength and the subordinate bias effect. *Quarterly Journal of Experimental Psychology*.
- RAYNER, K., & DUFFY, S. A. (1986). Lexical complexity and fixation times in reading: Effects of word frequency, verb complexity, and lexical ambiguity. *Memory & Cognition*, 14, 191-201.
- RAYNER, K., & FRAZIER, L. (1989). Selection mechanisms in reading lexically ambiguous words. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 15, 779-790.
- RAYNER, K., PACHT, J. M., & DUFFY, S. A. (1994). Effects of prior encounter and global discourse bias on the processing of lexically ambiguous words: Evidence from eye fixations. *Journal of Memory & Language*, 33, 527-544.
- SIMPSON, G. B. (1984). Lexical ambiguity and its role in models of word recognition. *Psychological Bulletin*, 96, 316-340.
- SIMPSON, G. B. (1994). Context and the processing of ambiguous words. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics* (pp. 359-374). San Diego: Academic Press.
- TWILLEY, L. C., DIXON, P., TAYLOR, D., & CLARK, K. (1994). University of Alberta norms of relative meaning frequency for 566 homographs. *Memory & Cognition*, 22, 111-126.

NOTES

1. For the ambiguous word *date*, both the Twilley, Dixon, Taylor, and Clark (1994) and the Nelson, McEvoy, Walling, and Wheeler (1980) norms provide three meanings: *girl*, *time*, and *fruit*. For the Twilley et al. norms, the *girl* and *time* meanings are fairly equal (.43 and .40, respectively), whereas the *fruit* meaning is clearly less frequent (.13). Similarly, in the Nelson et al. norms, the meaning frequencies for *girl* and *time* are .52 and .30, respectively, and *fruit* is also the least frequent meaning (.13). We think that, because of the way the sentences were written by Kellas et al. (1995), both contexts are consistent with the *time* meaning. If, however, they had written "He forgot *his* date," this would have been more consistent with the *girl* interpretation. But more importantly, the *fruit* meaning is the subordinate meaning, and neither of their contexts biases this interpretation.

2. We were initially informed by Kellas and colleagues that the Twilley et al. (1994) norms were the primary source used in the study by Kellas et al. (1995). However, after we had categorized the items and collected the data reported in Binder and Rayner (1998), we were informed that the Nelson et al. (1980) norms were, in fact, the primary source of norms.

3. Kellas and Vu say that the criterion for what constitutes the upper limit of balanced words "has never been previously reported." However, on p. 435 of Duffy et al. (1988), it is stated that a .69/.31 criterion was used. Also, it should be noted that most of the words in Table 3 are biased ambiguous words according to the Nelson et al. (1980) norms, but three of them arguably are not biased words, even by those norms.

4. We thank a reviewer of the Binder and Rayner (1998) article for noting this point.

5. We are not quite sure why Kellas et al. (1995) have not simply eliminated the items that we eliminated and rerun the analyses to see how those results compare with their original results. Presumably, if the items were not problematic for their subjects, the pattern of data should remain unchanged.

6. As stated in Binder and Rayner (1998), regardless of whether the subordinate bias effect was assessed by comparing the time spent reading the biased ambiguous word compared to itself (in the weak context condition) or to an unambiguous control word, the subordinate bias effect was not eliminated.

7. When we had to eliminate several items from the Kellas et al. (1995) stimulus set, it was important to make sure that the remaining items were representative, in contextual strength, to the entire original set of stimuli. Not only did we compare our selected items with the Kellas et al. contextual strength norming information, but we also collected our own norming data on their entire stimulus set.

8. In their Note 2, Kellas and Vu indicate that they are puzzled about why Binder and Rayner (1998) made reference to a 5-point rating scale. When we initially received the stimulus set, they enclosed a copy of an abbreviated method section that provided normative values for their strength manipulation. Nowhere in the document did they refer to a 9-point scale. Hence, we inferred that they used a 5-point scale, since the values were 3.7 and 1.99.

9. Kellas and Vu also state that "the contexts of the Duffy et al. (1988) items do not clearly eliminate the dominant meaning." We believe that this is a gross exaggeration, and interested readers are invited to consult the appendix of the Duffy et al. paper and make the determination for themselves.

(Manuscript received July 1, 1998;
revision accepted for publication January 15, 1999.)