8

Contextual Strength: the Whens and Hows of Context Effects

Orna Peleg, Rachel Giora and Ofer Fein*

1 Introduction

Highlighting the role context plays in shaping our linguistic behaviour is the major contribution of pragmatics to language research. Indeed, pragmatics has shifted the focus of research from the code to contextual inference (Carston, 2002; Sperber & Wilson, 1986/1995). It is widely agreed now that contextual information is a crucial factor determining how we make sense of utterances. The role of context is even more pronounced within a framework that assumes that the code is underspecified allowing for top-down inferential processes to narrow meanings down and adjust them to the specific context.

There is, however, ample evidence suggesting that the acknowledged supremacy of context should be qualified. Findings show that, at times, even a strong context does not filter out incompatible meanings and therefore does not allow frictionless processing. The conversation in (1), which took place between M and W who is interested in biology and genetics, is a case in point:

(1) *M*: I wanted to talk with you about something, but I can't remember what.
W: [NOTES SEWING THREADS ON THE TABLE] It must have to do with thread (Joking).
M: Yea, I wanted you to do Maya's jeans (Joking).
W: You know, the first interpretation I got was genes with a g.
(15 June 2001, reconstructed from memory, Ariel, in press.)

In spite of strong contextual evidence to the contrary – the cognitive environment manifest to the hearer strongly supports the 'jeans' meaning of the homophone – coupled with the implausibility of 'doing Maya's genes' as

*Corresponding author

opposed to 'doing Maya's jeans', the interpreter came up with the less likely interpretation first. Context did not inhibit what was foremost on his (genetics-oriented) mind – the 'genes' meaning of the linguistic code.

How come the hearer did not activate the relevant interpretation first? In what follows we will lay out the whens and hows of context effects: we will specify the conditions under which context may be more or less powerful and question the hypothesis that a strong context may affect comprehension entirely. Specifically, we will focus on the distinction between lexical processes involving coded but contextually inappropriate meanings versus contextual processes involving appropriate interpretations.

2 Effects of contextual strength on initial processing

Though no theory which accounts for comprehension denies the effect of context on how we make sense of utterances, various theories have different views on the speed and locus of these effects. Particularly, they diverge with regard to the very early moments of comprehension.

2.1 The direct access view

Proponents of the direct access view assume that context affects comprehension entirely. According to this view, top-down (contextual) processes interact with bottom-up (lexical) processes rather early on. If context is sufficiently rich and specific, it penetrates lexical processes and selects the appropriate meaning exclusively so that initial comprehension is effortless and seamless, involving no incompatible phase at all (e.g., Marslen-Wilson and Tyler, 1980; Martin, Vu, Kellas and Metcalf, 1999; McClelland and Rumelhart, 1981; Vu, Kellas and Paul, 1998; Vu, Kellas, Metcalf and Herman, 2000). Thus, upon processing

(2) The gardener dug a hole. She inserted the *bulb*.

comprehenders activate only the compatible 'flower' meaning of *bulb*, since this is the only interpretation of *bulb* that would be relevant in the given context. In contrast, the 'light' sense of *bulb*, though salient, should not be activated, since, in the set of accessible assumptions, it is irrelevant (Vu et al., 1998; Vu et al., 2000).

A more moderate version of the direct access view, while assuming that context affects comprehension significantly, also acknowledges the influence of meaning salience on comprehension. In this view, contextual processes are of primary effect: They interact with lexical processes and select the contextually appropriate meaning instantly. However, they do not inhibit irrelevant meanings, which get activated upon encounter of the lexical stimulus (Rayner, Pacht and Duffy, 1994; Kawamoto, 1993). Importantly, however, though lexical processes operate regardless of contextual processes

and allow activation of various meanings, the appropriate interpretation always reaches sufficient levels of activation first (Bates, 1999, personal communication, July 2001; Gibbs, 1994; McRae, Spivey-Knowlton and Tanenhaus, 1998).

2.2 The modular view

Unlike the direct access view, the modular view assumes independent – modular and non-modular – systems that do not interact initially (Fodor, 1983). A modular system (lexical access) is sensitive only to its domain-specific (lexical) information. It is encapsulated and does not have access to information outside the module. Rather, initial input analyses are stimulus driven. They are automatic, rapid and on some traditional interpretations, exhaustive: all the responses (meanings) to a stimulus (word) are activated upon its encounter. In contrast, non-modular systems (contextual processes) are sensitive to all kinds of information (linguistic and nonlinguistic) and integrate various outputs into a coherent representation. Non-modular, contextual processes thus affect comprehension post-lexically: they operate after all the meanings of a linguistic stimulus have been activated. Within this framework, context effects are limited. They are slower than lexical processes and either integrate contextually appropriate outputs or suppress them as irrelevant and interfering with comprehension (Swinney, 1979). As a result, initial processes are not always smooth and may involve contextually inappropriate responses that would trigger sequential processes.

The Gricean model (1975) can be viewed as compatible with this view. For Grice, linguistic processes are primary. Context affects comprehension only after the initial (literal) interpretation of the (sentence) unit has been accomplished. If this interpretation reaches contextual fit, no more processes are required. If, however, it fails, further inferential processes follow, involving suppression of irrelevant meanings and derivation of contextually appropriate interpretations (implicatures). In this view, then, comprehension may initially go astray, with a later revision and adjustment stage.

2.3 The graded salience hypothesis

The graded salience hypothesis (Giora, 1997, 2003; Peleg, 2002, Peleg, Giora and Fein, 2001) shares a number of assumptions with the modular view. It too assumes distinct mechanisms: one bottom-up, sensitive only to domain-specific (linguistic) information; and another, top-down, sensitive to all kinds of (linguistic and extra-linguistic) knowledge. Unlike the traditional modular assumption, however, it assumes that the modular (lexical access) mechanism is itself ordered:[1] more salient responses (meanings) are accessed faster than and reach sufficient levels of activation before less salient ones.

[1] For a similar view see Duffy, Morris and Rayner, 1988; Rayner and Frazier, 1989; Rayner and Morris, 1991; Sereno, Pacht and Rayner, 1992, among others.

A response is salient to the extent that it is coded. The relative salience of the coded meaning is a function of its prototypicality, or amount of experiential familiarity induced by exposure (frequency). Uncoded responses (implicatures) are nonsalient. According to the graded salience hypothesis, then, salient meanings would be activated automatically upon encounter of the lexical stimulus, regardless of contextual information.[2]

In this framework, contextual information may also affect comprehension immediately. Particularly, a highly informative context may be predictive enough to avail meanings on its own accord very early on without even penetrating lexical access. Indeed, strong contextual information may be faster than lexical processes, so much so, that it may avail meanings before the relevant stimulus is even encountered (fostering an impression of direct access). This may be particularly true when the stimulus is placed at the end of a strong sentential context, after most information has been accumulated and integrated, allowing for effective guessing based on inferential processes. Importantly, however, context does not interact with lexical processes but runs in parallel (Giora, Peleg and Fein, 2004; Peleg et al., 2001). According to the graded salience hypothesis, then, even a strong context has limited effects initially. It may be predictive but it cannot block salient meanings.

Assuming a simultaneous operation of the encapsulated, linguistic mechanism on the one hand and the integrative, central system mechanism on the other allows the graded salience hypothesis to predict when contextual information may be faster than, coincidental with, or slower than linguistic processes. Unlike the modular view, then, the graded salience hypothesis does not always predict slower contextual effects that result in sequential processes.[3] Neither does it assume that activation of a whole linguistic unit should be accomplished before contextual information comes into play (as assumed by Grice, 1975). Rather, along the communication path, context and linguistic effects run in parallel, with contextual information availing meanings on its own accord, affecting only the end product of the linguistic process.

2.4 Predictions

The various theories have different predictions with regard to the whens and hows of context effects. According to the direct access view, a strong context will always win over initially even if lexical effects are strong (as when it biases a polar ambiguity toward the less-salient meaning, having thus to inhibit or be faster than a highly accessible response). According to

[2] Coded meanings of low salience, however, may not reach sufficient levels of activation and may not be visible in a context biased toward the more salient meaning of the word (but see Hillert and Swinney, 2001, for a different view).

[3] Note, however, that Fodor (1983: 75) did not exclude predictive effects.

the modular view, lexical processes will always be faster, since they are automatic and encapsulated. Salience imbalance would not affect processing either, since response is exhaustive and unordered. The graded salience hypothesis takes both strength of context and salience effects into consideration. While lexical effects are constant across sentence position, being sensitive only to degree of (coded) salience, contextual effects may vary with respect to predictability and sentential position. Given these variables, the graded salience hypothesis predicts that:

a. Context effects might precede lexical effects when the stimulus is placed in sentence final position, provided the preceding context is highly predictive. Under this condition, guessing the compatible concept(s) would be fast and often occur before the lexical stimulus is encountered.
b. Contextual effects would not precede lexical effects in sentence initial position. In this position, even a strong prior context will not have speedy enough effects to enable it to predict oncoming concepts long before lexical accessing occurs. The assumption is that in initial position, predictive effects are less pronounced than in final position, since beginnings are less constrained than ends.[4]
c. Under all conditions, the incompatible coded meanings will be activated upon encounter of the stimulus, albeit at different levels of activation, determined by their relative salience (see note 2).

3 Findings

Findings in Giora and colleagues (2004), and Peleg and colleagues (2001) support the graded salience hypothesis. In all, our studies demonstrate that lexical and contextual processes make up independent mechanisms that do not interact initially. Specifically, we showed that, as predicted (see (a) above) when placed in final position, constraining contexts can predict the appropriate meaning of a lexical stimulus even before that stimulus is encountered, thus availing appropriate concepts without interacting with lexical processes (Experiment 1 below). We then compared access of coded meanings in sentence initial versus final position. As predicted (see (b) above), we showed that while final position favours context effects, initial position does not (Experiment 2 below). This is true even when the preceding context is highly predictive (as when the target sentence features the previous sentence topic; Experiment 4 below). However, even in final position, coded meanings get activated, despite contextual information to the contrary, as predicted (see (c) above; Experiments 3 and 4 below).

[4] See also Gernsbacher (1990) on the processes involved in building a new substructure that are initially insensitive to information in prior context.

3.1 Experiment 1

Review of the literature reveals that experimental data suggestive of selective access induced by prior context was based on materials whose targets were placed at the end of strong sentential contexts. For instance, Vu, Kellas and Paul (1998) and Vu, Kellas, Metcalf and Herman (2000) showed that homonyms such as *bat* activated contextually appropriate meanings exclusively when placed at the end of a highly constraining context such as (3)–(4):

(3) The slugger splintered the *bat*.*

(Probes displayed at *: salient-wooden; unrelated-safe; less-salient-fly; unrelated-station.)

(4) The biologist wounded the *bat*.*

(Probes displayed at *: salient-wooden; unrelated-safe; less-salient-fly; unrelated-station.)

In their studies, subjects read such sentences and named one of four probes (presented in (3)–(4). Findings demonstrated that they always named the contextually compatible probe faster than the unrelated one. On the face of it, then, such findings support the direct access view. They show that only contextually appropriate meanings were tapped initially, irrespective of meaning salience. Indeed, if these findings were a result of context penetrating lexical accessing, they would question the graded salience hypothesis.

To support the alternative view proposed by the graded salience hypothesis that lexical and contextual processes do not interact initially, one should be able to show that results, accounted for by an interactive system, can also be accounted for by non-interactive machinery. Thus, if Vu and collaborators' findings are replicated in the absence of the relevant lexical stimulus, this would support the view that these results are the end-product of contextual processes alone.

To do that, we used Vu and colleagues' materials, but presented the probes in sentence pre-final position in order to see whether contextual processes could induce the appropriate meaning even *before* the target word is encountered. In our study (Peleg et al., 2001), 60 native speakers of English read the sentences off a computer screen and had to make lexical decisions as to whether a probe presented before the final (target) word was a word or a non-word:

(5) The slugger splintered the* *bat*.

(Probes displayed at *: salient-wooden; less-salient-fly; unrelated-station.)

(6) The biologist wounded the* *bat*.

(Probes displayed at *: salient-wooden; less-salient-fly; unrelated-station.)
(Manipulated items taken from Vu et al., 1998.)

3.1.1 Results and discussion

Results replicated those by Vu and his colleagues, indicating that contextually compatible responses were always faster than incompatible responses, regardless of whether the context was biased in favour of the less or more salient meaning of the target (see Table 8.1). This was true of both the participant and item analyses. Replication of Vu and colleagues' findings under conditions that disallow lexical accessing is consistent with our assumption that contextual processes are speedy toward the end of sentences and can predict the appropriate meaning on their own accord very early on, without interacting with lexical processes.[5]

Notwithstanding, it still remains to show that when the lexical stimulus is eventually encountered, lexical accessing proceeds automatically, irrespective of contextual information. Experiment 2 was therefore designed to show that even a strong and speedy context does not penetrate lexical access when this is triggered.

3.2 Experiment 2

In order to show that lexical processes are encapsulated with respect to contextual information, we attempted to replicate Vu and colleagues' (2000) results, manipulating targets' position in the sentence (Giora et al., 2004). We predicted that, at the beginning of sentences, their (Vu et al., 2000) results will not be replicated, since at this position, effects of a strong prior context would neither inhibit nor precede salient meanings. These predictions do not fall out of interactive models, which assume that, in a rich and supportive context, the appropriate meaning is tapped initially, directly and exclusively, or at least more rapidly than the inappropriate meaning. However, as before, probing targets in sentence final position would yield results similar to those obtained by Vu and colleagues (and by our first experiment). Unlike initial position, we argue, sentence final position allows contextual processes to be fast and obscure but not inhibit lexical processes.

Table 8.1 Mean response times (in milliseconds) to probes by context type

	Salient Probe		Less-Salient Probe		Unrelated Probe	
Context	M	SD	M	SD	M	SD
Salient	951	252	1003	243	1005	255
Less-salient	1057	275	927	237	994	231

[5] For an alternative critique of Vu et al.'s findings, suggesting that it is the choice of items that is responsible for their results, see Binder and Rayner (1999).

To test our hypotheses we used Vu and colleagues' (2000) materials. For example, in (7) the context is suggestive of the salient/dominant ('electricity') sense of *bulb*; in (8) it is strongly suggestive of the less-salient/subordinate ('plant') sense of the word:

(7) The custodian found the solution. She inserted the *bulb.**

(Probes displayed at *: salient-light; less-salient-flower; unrelated- cliff.)

(8) The gardener dug a hole. She inserted the *bulb.**

(Probes displayed at *: salient-light; less-salient-flower; unrelated- cliff.)

To manipulate sentence initial versus final position, we subjected the second sentence of their (2000) materials (*She inserted the bulb*) to passivization. In this experiment, we tested only the less-salient condition, because it involves lexical access of salient meanings that conflict with contextual processes inducing compatible but less-salient meanings:

(9) The gardener dug a hole. The *bulb** was inserted.

(Probe displayed at *.)

Sixty native speakers of English read the original and the passivized versions of Vu and colleagues' (2000) discourses off a computer screen and were administered lexical decision tasks. Relative salience of target meanings had been established by a pre-test.

3.2.1 Results

Results support the graded salience hypothesis. They show that, as predicted, in sentence initial position, responses were faster to the salient (incompatible) probes than to the less-salient (compatible) probes. In sentence final position, however, the picture was different. Responses to the less-salient (compatible) probes were faster than responses to the salient (incompatible) probes (see Figure 8.1).

3.2.2 Discussion

The above results support our view concerning the whens and hows of context effects. In sentence initial position, where only constraints from a previous discourse can be operative, context effects are slow and do not precede lexical processes. Their slow effects in initial position, then, do not allow it to conceal the effects of the lexical mechanism, thus attesting to the involvement of different, non-interactive mechanisms in discourse comprehension. In contrast, this expectation-driven mechanism is fast toward the end of

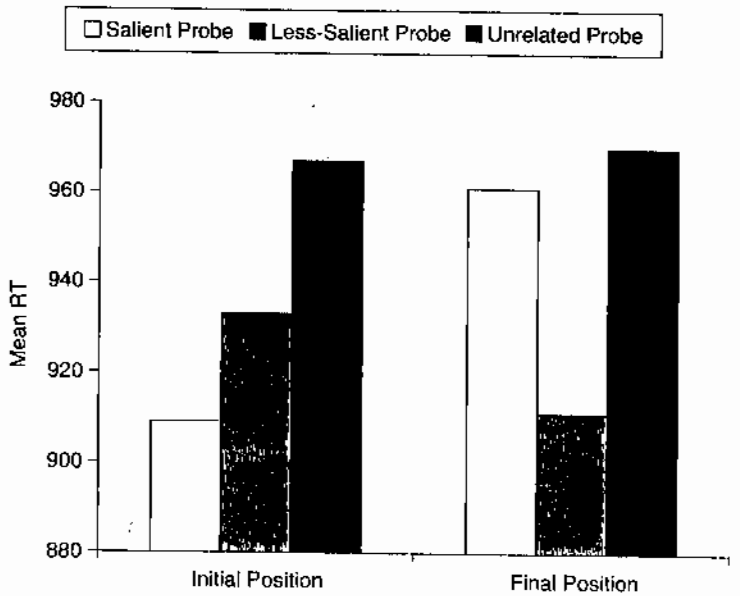

Figure 8.1 Mean response times (in milliseconds) to probes related to the salient (contextually incompatible) and less-salient (contextually compatible) meanings of the target words, and unrelated probes

sentences. At this point, different types of information enable it to predict an upcoming concept swiftly and obscure lexical processes.

Taken together, these findings cannot be accounted for by the context-sensitive, interactive models, which predict that, given enough constraints, either compatible meanings will be activated exclusively or they will be accessed first. These predictions do not hold for sentence initial position, in spite of a prior strong context. We want to further argue that the first prediction – regarding exclusive access of compatible meanings – does not in addition hold for either sentence position.

3.3 Experiment 3

To further demonstrate that, even in sentence final position, salient meanings are not blocked when incompatible, an additional experiment was designed (see Giora et al., 2004). The purpose of this study was to show that salient, but incompatible meanings are not inhibited even when context favours contextual effects. To do that, we compared sentences containing an ambiguous word whose less-salient meaning is contextually compatible (*The gardener dug a hole. She inserted the bulb*) with control sentences ending in a compatible but non-ambiguous word (*The gardener dug a hole. She inserted the flower*). We predicted that, following the ambiguous word (*bulb*), the salient but

Table 8.2 Mean response times (in milliseconds) to probes related to the salient (contextually incompatible) meaning

	Salient Probe	
Target word	**M**	**SD**
Ambiguous	993	285
Control	1070	283

incompatible meaning ('light') would be activated compared to the control condition. As in previous experiments, native speakers read the sentences and had to make a lexical decision as to whether related and unrelated probes were a word or a non-word in English.

3.3.1 Results and discussion

As predicted, inappropriate but salient meanings ('light') were activated following the ambiguous condition only (see Table 8.2). This was true for both the subject and item analyses. Such results demonstrate that salient though inappropriate meanings are activated even in a sentential position that benefits contextual processes. Placed in sentence final position, the ambiguous word (*bulb*) facilitated the activation of the probe related to the salient, but contextually inappropriate meaning ('light') compared to the control (*flower*). This finding is inconsistent with the predictions of the radical version of direct access view according to which interactive mechanism should have tapped the contextually compatible meaning exclusively, as allegedly shown by Vu and his colleagues (2000). However, as shown here, this was not the case. Their (Vu et al., 2000), findings are, therefore, more compatible with the assumption that, under conditions that favour contextual processes, a strongly biasing context can avail the appropriate meaning very early on without penetrating lexical accessing that might occur independently somewhat later.

Though our findings so far demonstrate that even in sentence final position, salient but incompatible meanings get activated, they nevertheless show that, in that position, context effects may be faster. The appropriate though less-salient meanings reach sufficient levels of activation faster than salient but incompatible meanings. Only in initial position, is this not the case. Would findings in sentence initial position be subverted if information in sentence initial position is exceedingly predictive? To test this possibility, we designed Experiment 4.

3.4 Experiment 4

Experiment 4 aimed to show that in initial position, contextual effects will not supercede lexical effects even when prior context is highly predictive of

oncoming concepts. One kind of high-predictability concepts is topical referents. We therefore compared activation levels of salient but irrelevant meanings with nonsalient but topically compatible interpretation of targets placed in sentence initial position preceded by a context substantiating this topical information. Indeed, initial position is known to be the preferred position for topics (see Giora, 1985a, 1985b, and Reinhart, 1980, and references therein).

Sixty native speakers read Hebrew sentences in which the target word (*delinquent*) appeared either in initial (10) or final (11) position. A prior context strongly biased these sentences toward a nonsalient (metaphorical) meaning which, in all cases, was the topic of the previous context as well as the topic of the target sentence. We took advantage of the relative free word order in Hebrew:

(10) Sarit's sons and mine went on fighting continuously. Sarit said to me: These delinquents* won't let us have a moment of peace.

(Probes displayed at *: salient-criminals; contextually compatible-kids; unrelated-painters.)

(11) Sarit's sons and mine went on fighting continuously. Sarit said to me: A moment of peace won't let us have these delinquents*.[6]

(Probes displayed at *: salient-criminals; contextually compatible-kids; unrelated-painters.)

Readers had to make a lexical decision as to whether the probe was a word or a non-word in Hebrew.

3.4.1 *Results and discussion*

Results show that context effects were not faster than lexical effects in sentence initial position. Though contextually compatible nonsalient meanings were made available immediately, these effects were not strong enough to supercede lexical effects. However, in final position, results replicated those of Vu and colleagues and of our own (Experiments 1 and 2). In sentence final position, contextual effects were somewhat faster than salience effects, emerging probably before the target word was encountered and processed (cf. Experiment 1). As before, these effects did not inhibit salient though inappropriate meanings in either position (see Figure 8.2 and Peleg et al., 2001).

Such findings support our view that language comprehension involves two distinct mechanisms that run in parallel: one sensitive to contextual

[6]The word order in Hebrew is such that the target NP occupies initial position, preceding the demonstrative.

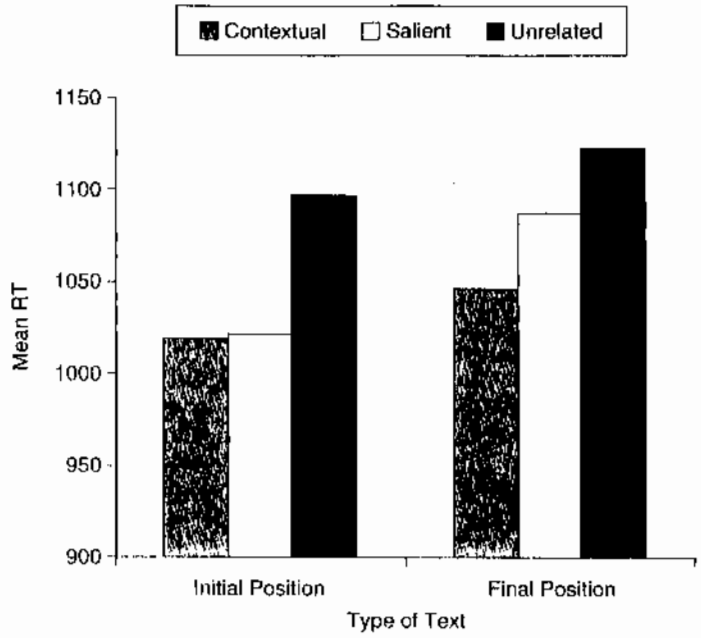

Figure 8.2 Mean response times (in milliseconds) to probes related to the salient (contextually incompatible) and less-salient (contextually compatible) meanings of the target words, and unrelated probes

information and one sensitive to coded, salient information. Thus while contextual information may have fast effects, they do not filter out salience effects. Salient meanings are activated upon encounter of the verbal stimulus, irrespective of context predictiveness. While salience effects are constant across position, speed of context effects varies as a function of the targets' location. They are faster toward the end of sentences, and less pronounced at the beginning of sentences.

4 General discussion

The involvement in comprehension of distinct mechanisms that do not interact initially enables comprehenders to resist exclusive conformity with contextual information. Contextual information, though effective, is limited. This intelligent, integrative mechanism is very powerful, particularly toward the end of discourse units. Still, it does not control other processes entirely. Contrary to appearances, it does not penetrate lexical accessing and it does not activate meanings selectively. Experiment 1 suggests that previous

findings supporting selective access (Vu et al., 1998; Vu et al., 2000) might have been affected by contextual processes which did not interact with lexical processes. Experiment 2 indeed demonstrated that contextual processes did not interact with lexical processes, which are automatic and sensitive to lexical stimuli only. Though sentence final position favours contextual processes and allows them to occur even before lexical access is initiated, this is not true of sentence initial position in which contextually compatible meanings are not faster than salient but incompatible meanings. This has been further demonstrated with information that is highly predictive pragmatically. Sentence initial position did not favour contextual information over lexical accessing of salient but incompatible meanings even when such contextual information was highly accessible and useful. Lexical processes, then, are uninterrupted initially, even when context is highly powerful location-wise (Experiment 3) and content-wise (Experiment 4).

The impenetrability of the lexical mechanism allows humans to have access to meanings not invited by information accumulated outside the module. This multiplicity of sources of meanings (originating in the context and the lexicon) allows for non-standard choices. Indeed, findings in Giora (2003) attest that comprehenders do not always suppress salient but contextually incompatible information (as assumed by Fodor, 1983, and Grice, 1975), but occasionally utilize it for various purposes such as humour, pleasure, innovativeness and subversion. The existence of a mechanism and a set of privileged meanings that resist immediate compliance with contextual information even when it is very strong provide for 'a variety of situations' which allow the individual an insight into different alternatives and a second (critical?) thought.

References

Ariel, M. (in press). *Pragmatics and Language*. Cambridge: Cambridge University Press.

Bates, E. (1999). On the nature and nurture of language. In E. Bizzi, P. Calissano and V. Volterra (eds), *Frontiere della biologia* [Frontiers of Biology]. *The Brain of Homo Sapiens*. Rome: Giovanni Trecanni.

Binder, K. S., and Rayner, K. (1999). Does contextual strength modulate the subordinate bias effect? A reply to Kellas and Vu. *Psychonomic Bulletin and Review* 6: 518–22.

Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.

Duffy, S. A., Morris, R. K., and Rayner, K. (1988). Lexical ambiguity and fixations times in reading. *Journal of Memory and Language* 27: 429–46.

Fodor, J. (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.

Gernsbacher, M. A. (1990). *Language Comprehension as Structure Building*. Hillsdale NJ: Lawrence Erlbaum Associates.

Gibbs, R. W. Jr. (1994). *The Poetics of Mind*. Cambridge: Cambridge University Press.

Giora, R. (1985a). Towards a theory of coherence. *Poetics Today* 6: 699–716.

Giora, R. (1985b). A text-based analysis of nonnarrative texts. *Theoretical Linguistics* 12: 115–35.

Giora, R. (1997). Understanding figurative and literal language: The graded salience hypothesis. *Cognitive Linguistics* 7: 183–206.

Giora, R. (2003). *On Our Mind: Salience, Context, and Figurative Language*. New York: Oxford University Press.

Giora, R., Peleg, O., and Fein, O. (2004). Resisting contextual information: You can't put a salient meaning down. Paper submitted for publication.

Grice, P. H. (1975). Logic and Conversation. In P. Cole and J. Morgan (eds), *Speech Acts: Syntax and Semantics* vol. 3: 41–58. New York: Academic Press.

Hillert, D., & Swinney, D. (2001). The processing of fixed expressions during sentence comprehension. In A. Cienki, B. Luka and M. Smith (eds), *Conceptual and Discourse Factors in Linguistic Structure*: 107–22. Stanford, CA: CSLI Publications.

Kawamoto, A. H. (1993). Nonlinear dynamics in the resolution of lexical ambiguity: A parallel distributed processing account. *Journal of Memory and Language* 32: 474–516.

Marslen-Wilson, W. D., and Tyler, L. K. (1980). The temporal structure of spoken language understanding. *Cognition* 8: 1–71.

Martin, C., Vu, H., Kellas, G., and Metcalf, K. (1999). Strength of discourse context as a determinant of the subordinate bias effect. *The Quarterly Journal of Experimental Psychology* 52A: 813–39.

McClelland, J. L., and Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception. Part 1: An account of basic findings. *Psychological Review* 88: 375–407.

McRae, K., Spivey-Knowlton, M. J., and Tanenhaus, M. K. (1998). Modeling the influence of thematic fit (and other constraints) in on line sentence comprehension. *Journal of Memory and Language* 38: 283–312.

Peleg, O. (2002). Linguistic and nonlinguistic mechanisms in language comprehension. Unpublished Ph.D. dissertation. Tel Aviv University.

Peleg, O., Giora, R. and Fein, O. (2001). Salience and context effects: Two are better than one. *Metaphor and Symbol* 16: 173–92.

Rayner, K., and Frazier, L. (1989). Selection mechanisms in reading lexically ambiguous words. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 15: 779–90.

Rayner, K., and Morris, R. K. (1991). Comprehension processes in reading ambiguous sentences: Reflections from eye movements. In G. B. Simpson (ed.), *Understanding Word and Sentence*: 175–98. Amsterdam: North Holland.

Rayner, K., Pacht J. M., and 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 and Language* 33: 527–44.

Reinhart, T. (1980). Conditions for text coherence. *Poetics Today* 1: 161–80.

Sereno, C. S., Pacht, J. M., and Rayner, K. (1992). The effect of meaning frequency on processing lexically ambiguous words: Evidence from eye fixations. *Psychological Science* 3: 269–300.

Sperber, D., and Wilson, D. (1986/1995). *Relevance: Communication and Cognition*. Oxford: Blackwell.

Swinney, D. A. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*, 18: 645–59.

Vu, H., Kellas, G., Metcalf. K., and Herman, R. (2000). The influence of global discourse on lexical ambiguity resolution. *Memory and Cognition* 28: 236–52.

Vu, H., Kellas, G., and Paul, S. T. (1998). Sources of sentence constraint in lexical ambiguity resolution. *Memory and Cognition* 26: 979–1001.