On the Graded Salience Hypothesis

RACHEL GIORA

Kecskes: *Professor Giora, in your recent papers and book (1997, 1999, 2003) you offer an alternative to the interactionist view that assumes the priority of context and selective compliance with contextual information. Your "Graded Salience Hypothesis" (GSH) claims that lexical processing takes priority over contextual processing. In the initial phase of language comprehension, contextual and lexical processes do not interact but run parallel, and this stage is dominated by lexical access. Salient (coded and prominent) meanings of lexical units are accessed automatically and are then revised in the case of a misfit with context. In what follows I would like you to answer three questions concerning your theory.*

Giora: At the outset, let me first elaborate a bit on the "priority" of salient meanings (Giora 1997, 1999). Following Fodor (1983), the Graded Salience Hypothesis assumes two different types of mechanisms that run parallel: a modular, bottom-up mechanism (e.g., lexical access) and a nonmodular, top-down machinery (e.g., contextual processes). There is, however, no claim that bottom-up processes are *necessarily* superior to top-down processes as far as speed is concerned. In fact, context may be highly informative and specific so that it can predict the appropriate meaning of the oncoming (linguistic) stimulus quite early on (Fodor 1983: 75–78; Peleg, Giora, and Fein 2001, in press; Rayner, Binder, and Duffy 1999). However, even when context is fast enough to come up with the contextually appropriate meaning before the processor encounters the relevant stimulus, it would not block the contextually inappropriate but salient responses related to that stimulus. In this sense, then, salient meanings are privileged, because they are context resistant. They would pop up regardless of contextual fit. At times, when contextually inappropriate, they wouldn't even be suppressed (Giora 2003, Chapter 2; Peleg et al. 2001, in press; for a different view, see Vu, Kellas, and Paul 1998).

Consider, for instance, the following art (1) by Tartakover—a Jewish Israeli artist. This art features a Palestinian boy with his toy. The text

Intercultural Pragmatics 1-1 (2004), 93–103 1612-295X/04/0001–0093

cites the Israeli military spokesperson informing us that a 6 year old Palestinian was shot dead "according to regulations." The artwork is titled *Childhood is not child's play*. In Hebrew and English, the salient (coded and prominent) meaning of *child's play* is nonliteral and has to do with "ease and simplicity." The contextually appropriate meaning, however, invites a compositional "risk-free activity" interpretation of the collocation. This literal, contextually appropriate interpretation, however, would neither block nor suppress the "ease and simplicity" sense of *child's play* that has been accessed on account of its salience.

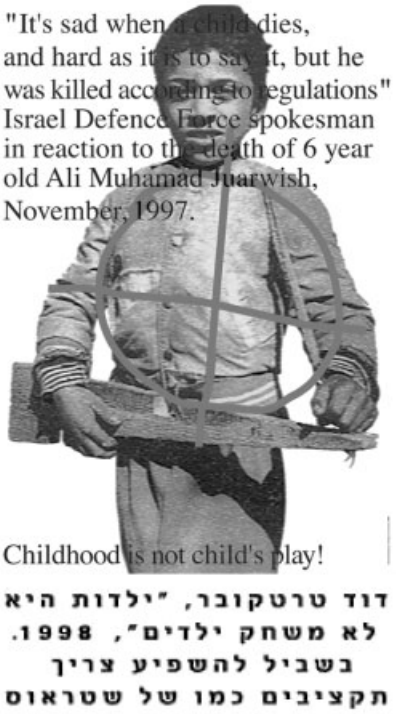

Or consider another example, taken from an article dealing with the killing and maiming of the Palestinians and destroying their livelihood and infrastructure (cited and analyzed in Giora 2003: 3). The paragraph below compares the Israeli brutal practices in the occupied territories to those of the German Nazis. Initially, however, the analogy is not made explicit (*There aren't six million Palestinians in the occupied territories*). Still, it is not lost on the readership, because, in the mind of Israelis, and Jews in particular, *six million* is strongly associated with the six million Jewish victims of the holocaust. No wonder then that the holocaust

meaning of *six million* is not inhibited by the preceding negation marker. Neither is it suppressed by the late context, which makes it clear that *six million* refers here to Palestinians rather than to Jews:

(2) There aren't six million Palestinians in the occupied territories, and the ideology of evil is different as well. Blunt and direct Nazi ideology is only found in the Messianic centers of the settlers in the territories. (Reinhart 2001)

It is because of the superiority of salience effects that the reference in the following sentence to Nazi ideology (*the ideology of evil is different as well*) needs no introduction and can be treated as "given." This reference demonstrates context ineffectiveness in blocking and even in suppressing so-called irrelevant but salient meanings. (For experimental findings, see Giora 2003; for the ineffectiveness of negation in inhibiting negated information, see Giora, Balaban et al. 2004; Giora, Fein, et al. 2004). In addition to salience resistance to context effects, "the priority of salient meanings" further pertains to the superiority of salient over less salient meanings: salient meanings would be accessed faster than less salient meanings (for a review, see Giora 2003, chapter 2).

Importantly, "the priority of salient meanings" also alludes to the superiority of "salience" over alternatives such as "literality" and "nonliterality" when processing is in question (for a review, see Giora 2002a). For instance, in Giora and Fein (1999a) we show that familiar metaphors, whose literal and figurative meanings are similarly salient, were processed along similar comprehension routes. In contrast, novel metaphors, whose metaphoric meaning is nonsalient, took longer to read than their more accessible literal interpretation. Similarly, familiar ironies and their familiar literal interpretations were processed initially both literally and ironically. Unfamiliar ironies were initially processed only literally (Giora & Fein 1999b). Such findings support the view that, rather than the literality-nonliterality divide, it is the salience continuum that plays a significant role in language comprehension.

Kecskes: *What prior knowledge is encoded in the lexical units? Is it fair to say that lexical units encode prior standard contexts as a result of their extensive use? If so, it would make sense to differentiate between prior standard context and the actual context. World knowledge is given to people in two forms: encapsulated in lexical items based on prior encounters and experience (prior standard context), and in the actual linguistic and extralinguistic context framed by the given situation. What do you think about the idea that actual contextual meaning is the result of the "colli-*

sion'' of prior contexts encoded in the lexical units and actual linguistic context supported by the actual socio-cultural background (Kecskes 2002)?

Giora: In part, what is encoded in a lexical item is the processor's past experience with that item in various contexts. Indeed, exposure to the meanings of a stimulus is one factor that determines degree of salience within the lexicon. There are, however, factors that might not be related to experiential familiarity, frequency, or conventionality but would nonetheless contribute to a concept's priority. These are ''unspoken'' issues, whose salience is high because they enjoy cultural or individual prominence. For instance, we might have little experiential familiarity with taboo words, simply because they don't occur often in conversations. Nonetheless, the concepts denoted by these stimuli preoccupy speakers and are foremost in their mind. As a result, upon encountering these words, even in a context that is heavily biased toward their more frequent but non-taboo meaning, their less frequent but ''forbidden'' meaning would spring to mind on account of its salience. For instance, in Hebrew, *ligmor* is an ambiguous word denoting both ''to come sexually'' and ''to finish/end.'' The most frequent use of *ligmor* is its nonsexual sense. Nonetheless, because *ligmor* is sexually loaded, and because, though ''unspoken,'' sex and sexual connotations, if not frequently on our mind, have a special prominent status, there is hardly any chance this sense would not spring to mind when *ligmor* is encountered. As a result, Hebrew speakers now refrain from using *ligmor* in nonsexual contexts. Instead, they use a non-sexually laden alternative—*lesayem*. Regardless, since *lesayem* is used euphemistically to avoid using *ligmor*, its use almost always brings to mind the to-be-avoided sexual connotations of the replaced word *ligmor*.

The following English example might be illustrative of the high salience of less frequent senses (Giora 2003: 175):

(3) A bus stops and two Italian men get on. They sit down and engage in an animated conversation. The lady sitting behind them ignores them at first, but her attention is galvanized when she hears one of the men say the following:

> ''Emma come first. Den I come. Den two asses come together. I come once-a-more. Two asses, they come together again. I come again and pee twice. Then I come one lasta time.''
> ''You foul-mouthed swine,'' retorted the lady indignantly. ''In this country we don't talk about our sex lives in public!''
> ''Hey, coola down lady,'' said the man. ''Who talkin' abouta sexa? I'm a justa tellin' my frienda how to spella 'Mississippi.'''

It is important to note that salience is not just a function of experience. It has cognitive components as well. For instance, prototypical meanings need not be the most frequent ones, but they would be more salient than less prototypical, though equally frequent concepts. That is, I may be similarly exposed to crows and sparrows. But when encountering *bird*, it would be the "sparrow" that would come to mind first, because it is a more salient member of the category of birds than crows (Rosch 1973). Similarly, because my cognition is better adapted to recognize people from their faces than from any other part of their body, I would more easily recognize my son from his face than his back even though I know him very well from top to toe.

When perceiving a stimulus in an actual context, that context need not "collide" with the most salient response of that stimulus, because, in most cases, the most salient, in fact, the most probable, response of a stimulus would also be the one invited by the very same context. At times, of course, there would be a mismatch between salient responses and contextually appropriate ones. This mismatch would involve further activation processes or inferencing.

Kecskes: *The idea of underspecified word meaning figures in several linguistic theories including Bierwisch's (1996, 1997) two level conceptual semantics, Pustejovsky's (1995) theory of the generative lexicon and Sperber and Wilson's (1986/1995) relevance theory. These approaches claim that lexical units get into the context underspecified. It is argued that the specification of word meaning in context is achieved by conceptual shift, which "shifts" the core meaning into various conceptual fields, and by conceptual differentiation, which only "differentiates" the core meaning in different ways within one and the same conceptual domain, yields literal meanings. In other words, interpretations mean mapping underspecified semantic meanings onto fully determined conceptual meanings on the basis of our encyclopedic knowledge (e.g., Bierwisch 1996, 1997; Bibok and Németh 2001). What do you think about these approaches? Do they completely contradict the GSH or is there some way of reconciling the two approaches?*

Giora: Assuming that the lexicon's entries are underspecified assigns to context a role in homing on a (contextually appropriate) specific interpretation. It is important to note, however, that speakers do not always intend addressees to look for a fully specified interpretation and that comprehenders, on their part, are not always after a fully specified interpretation. The following conversation, taken from Tsuyoshi & Thompson (1996: 77; numbers in square brackets indicate overlap) shows that the

speaker responds by agreeing even before she hears the full propositional content:

(4) D: **. . she [4 always was, . . you know 4].**
G: **[4 yeah. . . exactly 4].**
D: . . . pretty much uh, . . . able to do anything that I wanted to do.

Of course, all the theories assume that the pragmatically appropriate interpretation is a result of some adjustment to contextual information. They vary only with regard to when and how, in the course of interpretation, this adjustment takes place and affects interpretation. Primarily, they diverge with regard to whether it is a pre or a post-lexical effect. It is apparent, then, that the Graded Salience Hypothesis need not take issue with a view of the lexicon as comprising of underspecified entries. It should, however, take issue with theories that assume that lexical access is not autonomic but can be affected by contextual processes so that context pre-selects the appropriate (core) meaning while blocking inappropriate ones.

Again, whether meanings or senses of a stimulus are more or less specified is less crucial for the Graded Salience Hypothesis, because, even on a specification view, we are dealing with some level of abstraction. And even on the Underspecification Model, it is not quite clear how underspecified the core is, nor how specific the final interpretation is.

I am not sure that Relevance Theory (Carston 2002; Sperber & Wilson 1986/1995) indeed assumes an underspecification view of the lexicon. Given that context shapes interpretation following initial access of logical forms and linguistic meanings, final interpretation is achieved via processes such as loosening and narrowing down of initial outputs (Carston 2002: 323–359). Narrowing down these outputs involves the addition of conceptual material. This enrichment, which reduces the things in the world the concept can pick out (5; added features in square brackets), is consistent with an underspecification view of the lexicon. Loosening such outputs, which involves the subtraction of conceptual material so that it becomes more inclusive (6; reduced features in square brackets), seems less so. In fact, discarding features of a concept that has already been accessed seems more in line with a view of a more rather than a less specified core.

(5) [Palestinian] Villagers continue resistance [against the construction of the so-called ''security fence'' built by the Israeli government] despite [Israeli] military repression in Beit Surik. (International Solidarity Movement's email message, 27.02.04)

(6) You murdered [literally killed] our sister. (Said by the sister of a victim of a suicide bomber to Mofaz, the Israeli Defense Minister, who visited their home; Channel 10, 15.6.03. Originally in Hebrew).

The Graded Salience Hypothesis should have no problem with a view of the lexicon as involving encoding the very same concept at various levels of specificity abstraction (see Armstrong, Gleitman, & Gleitman 1983), or with a hierarchical view of a concept's structure involving more and less specified features. It has also no problem with the view that assumes that the lexicon comprises entries at various levels of specification for different concepts. It is quite clear, though, that it assumes some level of specificity and therefore predicts processing difficulties when these meanings or senses are salient but contextually inappropriate.

Let's consider the psycholinguistic angle of the Underspecification View. For instance, according to one version of the Underspecification Model, there are more and less specific entries in the lexicon. Initially, however, comprehension would involve only an underspecified core, which is compatible with all the specific senses related to that core. For instance, items such as metaphors and metonymies, whose multiple (literal and nonliteral) senses are related, will activate an abstract core, which is compatible with both senses. Initial processing should therefore involve no frictions. However, at a later phase, context would home in on the specific, contextually appropriate sense. At this stage, processing difficulties might emerge, particularly for the less salient sense (Frisson & Pickering 2001). In contrast, items consisting of multiple but unrelated meanings (e.g., homonyms) cannot collapse into one core compatible with all its meanings. As a result, the processor is forced to make an early commitment when it encounters such ambiguity. It thus selects the contextually appropriate meaning very early on while discarding the inappropriate one. However, since the inappropriate meaning has surfaced on account of its salience, it would cause contextual misfit resulting in early processing difficulties (Frazier & Rayner 1990).

The assumption that initial processing involves underspecified core meaning might be particularly true when the various senses are similarly salient. However, when salience imbalance is significant, the Graded Salience Hypothesis predicts the superiority of one response over the other even when related senses are concerned. As mentioned earlier, it indeed predicts that highly familiar responses such as the idiomatic meaning of highly familiar idioms would be initially processed idiomatically. Literal interpretation will lag behind. In a literally biasing context, this will affect

(V7(M) 21/6/04 15:54) WDG/G J-1113 Intercultural, 1-1 PMU: WSL 21/6 Times_M (0).3.04.05 (148×225mm) pp. 93–104 ch05_P (p. 99)

early processing difficulties (van de Voort & Vonk 1995; for similar findings regarding salient and less or nonsalient senses, see also Gibbs 1980; Janus & Bever 1985; Giora, Fein, et al. 2004; Pexman, Ferretti, & Katz 2000). If, however, initial processing involves multiple coded meanings, the predictions of the Graded Salience Hypothesis are similar to those of the Underspecification Model.

Kecskes: *One of the most important parts of the Graded Salience Hypothesis is the dynamics of salience. Salient meanings change diachronically. You say in your book (Giora 2003) that "how salience shifts is still a mystery." What do you mean by this? Do you think that there is no theoretical explanation for why lexical units such as "piece of cake," "cool," "gay," and the like have changed their salient meaning. How can historical semantics help us with this "mystery"?*

Giora: I was basically wondering about the evolution of meanings: what motivates the dynamics of salience. In our recent work we suggest that, among other things, linguistic change might be motivated by the pursuit of pleasure (Giora 2003, chapter 7; Giora, Fein, Kronrod, Elnatan, Shuval & Zur 2004).[1] Our findings show that pleasure is sensitive to optimal novelty—innovation that evokes both a novel and a familiar response. Pleasure thus seems sensitive to a stimulus that allows a novel insight into the salient. You might as well say that change is motivated by subversiveness, because it is the novelty that de-automatizes the familiar that pleases most (see also Mukařovský 1932/1964).

Consider, for instance, the novelty of *weapons of mass distraction* (Butcher 2003)[2], which, while evoking the all too familiar *weapons of mass destruction*, further questions it. It exposes it as a piece of propaganda—as a catchphrase used by governments to lure the masses into accepting the occupation of Iraq. Or consider the ironic, all too familiar *Read my lips*, which started its career as a literal expression only to be later undermined by an ironical, optimally innovative usage. Recently, however, this entrenched ironic meaning has been subjected to yet another optimally innovative manipulation in the form of *Read my lipstick* used ironically to deride a parliament candidate on account of her femininity.

In the same manner, the caricature in (7) questions the legitimacy of the Israeli military (Zahal). It translates into "Let Zahal lie" and subverts the familiar "Let Zahal lead"—a recent slogan of the Jewish settlers in the occupied territories who are trying to silence any criticism of the military coming from the left:

In addition to finding that subversiveness is likable, we have also found that the entirely novel is least preferred while the most salient is second best in liking (for a different view see Berlyne 1971; Wundt 1874; Zajonc 1968, 1980, 2000). Such findings highlight the role of salience in aesthetics and preference. Across the board, stimuli involving salient responses were preferred over those that evoked only novel ones (Giora et al. 2004). This might explain an intercultural phenomenon. In many languages, coining new words when these are in need tends to echo salient meanings and sounds. For instance, the Hebrew *dibuv* ('dub') has been selected so as to match the English *dub*, which is familiar to many Hebrew speakers. Similarly, the Hebrew *ashaf* echoes the nativized English *chef*. (Zuckerman 1999). Although these innovations are not optimal, they rely on salience to sugar the pill of the novel. (For the definition of Optimal Innovation, see Giora et al. 2004).

The finding that we are inclined toward the familiar may explain why propaganda is so effective. It keeps repeating messages until they are familiar enough to make us feel comfortable about them. And because what is pleasing gets across as true (see McGlone & Tofighbakhsh 2000), we are almost defenseless against the familiar. Still, if it is indeed the pursuit of pleasure that, among other things, motivates change, optimal novelty in language and mind may eventually override the non-costly pleasure effects of familiarity and consequently become a redeeming feature of language, science, and art.

(V7(M) 21/6/04 15:54) WDG/G J-1113 Intercultural, 1-1 PMU: WSL 21/6 Times_M (0).3.04.05 (148×225mm) pp. 93–104 ch05_P (p. 101)

Notes

1. For information on how social identity and awareness motivate linguistic change, see Ariel & Giora (1998); Coates (1986); Giora (2002b); Peleg (1992) among others. On semantic change, see Traugott, & Dasher (2002b).
2. This neologism is taken from *The madness of George Dubya*—a recent theater show by Justin Butcher (London: May, 2003).

References

Ariel, Mira and Rachel Giora. 1998. A Self versus Other point of view in language: Redefining femininity and masculinity. *International Journal of the Sociology of Language* 129: 59–86.

Armstrong, Sharon, Lila Gleitman, and Henry Gleitman. 1983. What some concepts might not be. *Cognition* 13: 263–308.

Berlyne, Daniel E. 1971. *Aesthetics and Psychobiology*. New York: Century Psychology Series.

Bibok, Károly and Enikö Németh. 2001. How the lexicon and context interact in the meaning construction of utterances. In Enikö Németh and Károly Bibok (eds.), *Pragmatics and the Flexibility of Word Meaning.* Amsterdam: Elsevier. 289–321.

Bierwisch, Manfred. 1996. How much space gets into language? In Bloom, Paul, Mary A. Peterson, Lynn Nadel, and Merrill F. Garrett (eds.), *Language and Space.* Cambridge, MA: MIT Press. 31–77.

—. 1997. Lexical information from a minimalist point of view. In Chris Wilder, Hans-Martin Gartner and Manfred Bierwisch (eds.), *The Role of Economy Principles in Linguistic Theory.* Berlin: Academie Verlag. 227–267.

Carston, Robyn. 2002. *Thoughts and Utterances*. Oxford: Blackwell.

Coates, Jennifer. 1986. *Women, Men and Language*. London: Longman.

Fodor, Jerry. 1983. *The Modularity of Mind*. Cambridge, MA: MIT Press.

Frazier, Lyne and Keith Rayner. 1990. Taking on semantic commitments: Processing multiple meanings vs. multiple senses. *Journal of Memory and Language* 29: 181–200.

Frisson, Steven and J. Martin Pickering. 2001. Obtaining a figurative interpretation of a word: Support for underspecification. *Metaphor and Symbol* 16: 149–172.

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

—. 1999. On the priority of salient meanings: Studies of literal and figurative language. *Journal of Pragmatics* 31: 919–929.

—. 2002a. Literal vs. figurative language: Different or equal? *Journal of Pragmatics* 34: 487–506.

—. 2002b. Theorizing gender: Feminist awareness and language change. In Bettina Baron and Helga Kotthoff (eds.), *Gender in Interaction.* Amsterdam: John Benjamins. 329–347.

—. 2003. *On Our Mind. Salience, Context, and Figurative Language.* Oxford, New York: Oxford University Press.

Giora, Rachel, Noga Balaban, Ofer Fein, and Inbar Alkabets. 2004. Negation as positivity in disguise. In Herbert L. Colston, and Albert Katz (eds.), *Figurative Language Comprehension: Social and Cultural Influences.* Hillsdale, NJ: Erlbaum. 233–258.

Giora, Rachel and Ofer Fein. 1999a. On understanding familiar and less-familiar figurative language. *Journal of Pragmatics* 31: 1601–1618.

—. 1999b. Irony: Context and salience. *Metaphor and Symbol* 14: 241–257.

Giora, Rachel, Ofer Fein, Ann Kronrod, Idit Elnatan, Noa Shuval, and Adi Zur. 2004. Weapons of mass distraction: Optimal innovation and pleasure ratings. *Metaphor and Symbol* 19: 115–141.

Janus, Raizi A. and Thomas G. Bever. 1985. Processing of metaphoric language: An investigation of the three stage model of metaphor comprehension. *Journal of Psycholinguistic Research* 14: 473–487.

Kecskes, Istvan. 2002. *Situation-Bound Utterances in L1 and L2.* Berlin: Mouton de Gruyter.

McGlone, S. Matthew and Jessica Tofighbakhsh. 2000. Birds of a feather flock conjointly (?): Rhyme as reason in aphorisms. *Psychological Science* 11: 424–428.

Mukařovský, Jan. 1932/1964. Standard language and poetic language. In Paul L. Garvin (ed.), *A Prague School Reader on Esthetics, Literary Structure, and Style*. Washington, DC: Georgetown University Press. 17–30.

Peleg, Orna. 1992. The relationship between linguistic divergence and social identity among Israelis of oriental origin. Unpublished MA thesis. Tel Aviv University.

Peleg, Orna, Rachel Giora, and Ofer Fein. 2001. Salience and context effects: Two are better than one. *Metaphor and Symbol* 16: 173–192.

—. (In press). Contextual strength: The whens and hows of context effects. In Ira Noveck and Dan Sperber (eds.), *Experimental Pragmatics.* Basingstoke: Palgrave.

Ortony, Andrew, Richard J. Vondruska, Mark A. Foss, and Lawrence E. Jones. 1985. Salience, similes, and asymmetry of similarity. *Journal of Memory and Language* 24: 569–594.

Pexman, M. Penny, Tod Ferretti, and Albert Katz. 2000. Discourse factors that influence irony detection during on-line reading. *Discourse Processes* 29: 201–222.

Pustejovsky, James. 1995. *The Generative Lexicon.* Cambridge, MA: MIT Press.

Rayner, Keith, Katherine S. Binder, and Susan Duffy. 1999. Contextual strength and subordinate bias effect. *Quarterly Journal of Experimental Psychology* 52A: 841–852.

Sperber, Dan and DeirdreWilson. 1986/1995. *Relevance: Communication and Cognition.* Oxford: Blackwell.

Reinhart, Tanya. 2001, March 14. WE DIDN'T SEE; WE DIDN'T KNOW. *Yediot Aharonot* (English translation http://www.MiddleEast.Org).

Traugott, Elizabeth C. and Richard B. Dasher. 2002. *Regularity in Semantic Change*. Cambridge Studies in Linguistics 97. Cambridge: Cambridge University Press.

Van de Voort, Marlies E. C. and Wietske Vonk. 1995. You don't die immediately when you kick an empty bucket: A processing view on semantic and syntactic characteristics of idioms. In Martin Everaert, Eric-Jan van der Linden, Andre Schenk, and Rob Schreuder (eds.), *Idioms: Structural and Psychological Perspectives.* Hillsdale, NJ: Lawrence Erlbaum Associates. 283–299.

Vu, Hoang, George Kellas, and Stephen T. Paul. 1998. Sources of sentence constraint in lexical ambiguity resolution. *Memory and Cognition* 26: 979–1001.

Wundt, Wilhelm M. 1874. *Grundzung der Physiologischen Psychologie*. Leipzig, Germany: Engelmann.

Zajonc, Robert B. 1968. Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology* 9: 1–27.

—. 1980. Feeling and thinking: Preference needs no inferences. *American Psychologist*, 35, 151–175.

—. 2000. Closing the debate over the independence of affect. In J. P. Forgas (ed.), *Feeling and Thinking: The Role of Affect in Social Cognition*. Cambridge, UK: Cambridge University Press. 31–58.

Zuckerman, Gilad. 1999, December. 'Phono-semantic matching' as a means of lexical enrichment: Comparative analysis of Ivrit, Turkish, Chinese, and Japanese. Paper presented at the Linguistics Colloquium, Tel Aviv University.