On our mind : salience, context, and figurative language

Rachel Giora.

# ON OUR MIND

Salience, Context, and Figurative Language



# OXFORD

Oxford New York
Auckland Bangkok Buenos Aires Cape Town Chennai

Dar es Salaam Delhi Hong Kong Istanbul Karachi Kolkata

Kuala Lumpur Madrid Melbourne Mexico City Mumbai Nairobi

São Paulo Shanghai Taipei Tokyo Toronto

Copyright © 2003 by Oxford University Press, Inc.

Published by Oxford University Press, Inc. 198 Madison Avenue, New York, New York 10016

www.oup.com

Oxford is a registered trademark of Oxford University Press

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior permission of Oxford University Press.

Library of Congress Cataloging-in-Publication Data
Giora, Rachel, date
On our mind: salience, context, and figurative language / Rachel Giora.
p. cm.
Includes bibliographical references and index.
ISBN 0-19-513616-0
1. Psycholinguistics. 2. Figures of speech. 3. Context (Linguistics) I. Title.
BF455 .G525 2002
401'.9—dc21 2001006646

For my mother and father, Esther and Arie Shapira

987654321

Printed in the United States of America on acid-free paper

## **PRFFACE**

This book has been fun to write. It is not an individual enterprise. Some of my best friends have had a hand in shaping my ideas and nourishing the thoughts I tried to flesh out here. Whether they like it or not, I do hold them accountable. The list is long and unordered. I am especially indebted to Efrat Ben-Menachem, who first introduced me to the pleasures of researching. My dearest friends and colleagues—Mira Ariel, Asa Kasher, Tanya Reinhart, and Yeshayahu Shen—are particularly culpable. They are among the people who influenced me most, not the least by being such caustic critics, yet relentlessly helpful.

A great portion of this book was written while I was on my sabbatical in Santa Cruz, visiting with Ray Gibbs. Being in Santa Cruz was a marvelous experience. It allowed me, among other things, to share my thoughts with Ray, who is a most wonderful person and an admirable colleague and mentor. I cherish the countless hours we spent talking about our views and differences and about the topics that pervade the field. Ray Gibbs has been and still is a great source of inspiration, and his ideas certainly are reflected in my work. My Santa Cruzian experience also included my very dear friends Adam Ussishkin and Andrew Wedel, who always lent me their linguistic ear and never failed to be supportive and critical.

I have also benefited greatly from invaluable discussions with Ofer Fein—a colleague and co-author for several years—and from discussions and very valuable comments by Salvatore Attardo, Brian Bowdle, Hugh Bredin, Dan Chiappe, Herb Clark, Herb Colston, Seana Coulston, Carmen Curcó, Jack Du Bois, Steven Frisson, Dedre Gentner, Morton Gernsbacher, Paul Hopper, Dick Janney, Albert Katz, John Kennedy, Boaz Keysar, Teenie Matlock, Hoang Vu, Ellen Winner, Francisco Yus, Eran Zaidel, four anonymous reviews of this

book, and numerous other people who I have read or talked to over the past few years. My discussions with several students—Ann Kronrod, Noa Shuval, and particularly Orna Peleg—have been inspiring and helpful in molding many of my views. I am no less indebted to Catharine Carlin, Robin H. Miura, and John Rauschenberg from Oxford University Press, and to copyeditor Cynthia Garver, for all their support and help. Particularly dear to me are the comments of an old friend and colleague, Jacob Mey, who read the final version with a hawkish but most benevolent eye.

My heartfelt thanks go also to my sister, Galia Shaham, and my sons, Ran, Etay, and Yoav Giora for their love and support.

The research reported here could not have been performed without support afforded by the Israel Science Foundation, the Lyon Foundation, and the Tel Aviv University Basic Research Fund.

It is a great relief, though, that the quest for truth must always fail so that any new theory is bound to be improved, reversed, or replaced by new thinking. The thoughts expressed here are also improvements on previous articulations of them in my earlier writings. The accomplishments of this book, therefore, are an invitation for a reform.

Asim Abu Shakra's art was selected for the book's jacket as a tribute to the Palestinian people and their struggle for a land of their own. The different ways the Palestinians and the Israelis perceive the cactus shed light on the themes reverberating in this book. For the Palestinians, the cactus is a salient symbol of their homeland. To the Israelies, its salient meaning is associated with their renewed identity as people native to Eretz Yisra'el (the land of Israel).

In the preface to Asim Abu Shakra's catalogue, Anton Shammas laments the tragedy of the Palestinian people by spelling out the appropriation of the indigenous symbol by the newcomers:

And since the Palestinian cactus, if we don't let bygones be bygones, has long ceased to function as a thorny, succulent plant with thick, leave-shaped stems, as the dictionary would have it, and has become instead "a common nickname for those born in Eretz Yisra'el and raised in it [tsabar]", and nothing more—

Someone had to transfer the cactus to the pot.

Someone had to shake us and tell us that it's over and done with; that from this point on there is no return, neither to the map nor to the land—just being let on the windowsill....

It never occurred to anyone before:
That here is the homeland for you, hanging on the wall.
That only in this way it is possible to go on.
That now it is possible to go on.

—Anton Shammas, preface to Asim Abu Shakra (Catalogue, Tel Aviv Museum of Art, 1994, in Hebrew; my translation).

#### Credits

The poem appearing in chapter 2 is from Yona Wallach (1997), "Absalom" in Wild Light, p. 4; translated by Linda Zisquit; copyright 1997 by the Sheep Meadow Press; reproduced with permission.

The poem appearing in chapter 7 is by Efrat Mishori (1999), "The Wall of Motherhood" in S. Kaufman, G. Hasan-Rokem, & T. S. Hess (Eds.), The Defiant Muse: Hebrew Feminist Poems from Antiquity to the Present, p. 237; copyright by the translator, Rachel Tzvia Back, and the Feminist Press; reproduced with permission.

All citations from Santa Barbara Corpus of Spoken American English are copyright 2000 by John Du Bois; reproduced with permission.

Chapter 4, figures 4.1, 4.2, 4.3, 4.4, 4.5, 4.6, and 4.7 are from R. Giora & O. Fein (1999), "Irony: Context and salience," *Metaphor and Symbol*, 14, 241–257; R. Giora, O. Fein, & T. Schwartz (1998), "Irony: Graded salience and indirect negation," *Metaphor and Symbol*, 13, 83–101; reproduced with the permission of Lawrence Erlbaum Associates, Inc.

Chapter 5, figure 5.2 is from O. Peleg, R. Giora, & O. Fein (2001), "Salience and context effects: Two are better than one," *Metaphor and Symbol*, 16, 173–192; reproduced with the permission of Lawrence Erlbaum Associates, Inc.

Chapter 5, figures 5.3 and 5.4 are from R. Giora & O. Fein (1999), "On understanding familiar and less-familiar figurative language," *Journal of Pragmatics*, 31, 1601–1618; reproduced with the permission of Elsevier Science.

All citations from R. Giora (2002), "Masking one's themes: Irony and the politics of indirectness," in M. M. Louwerse & W. van Peer (Eds.), *Thematics in psychology and literary studies* are copyright 2002 by John Benjamins; reproduced with permission.

# **CONTENTS**

- 1 Prologue 3
- 2 Salience and Context 13
- 3 Lexical Access 39
- 4 Irony 61
- 5 Metaphors and Idioms 103
- 6 Jokes 167
- 7 Innovation 176
- 8 Evidence from Other Research 185
- 9 Coda: Unaddressed Questions— Food for Future Thought 196

Notes 201

References 213

Author Index 243

General Index 251

ON OUR MIND

# Prologue

## We Didn't See; We Didn't Know

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. The army and the government are just protecting the living space of the settlers. And the rest are just deeply disappointed with the Palestinians, who failed to grasp how profound our desire for peace is.

But in Germany, too, most of the Germans were not Nazis. The majority just chose not to know.

-Tanya Reinhart, Yediot Aharonot [an Israeli daily], 14 March 2001

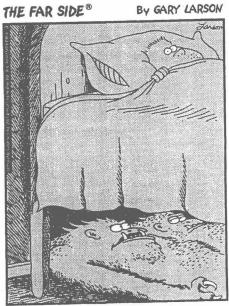
For Israeli Jews, the coin six million is strongly associated with the six million Jewish victims of the Holocaust. In spite of the negation marker (There aren't six million Palestinians), the accessibility of the collocation six million Jews is not diminished and allows for the analogy between the Palestinian and the Jewish victims to be established rapidly, even before it is made explicit by the late context. This is just one example shedding light on how privileged meanings shape the way we understand language. In this book, I explore the extent to which such prominent meanings affect our linguistic and psycholinguistic behavior.

To give you an immediate grasp of the topics I address in this book, consider a few mundane anecdotes related to familiar routines and meanings foremost on our mind. For instance, although my telephone number is highly familiar to me, I find it difficult to recognize if someone segments it differently or reads it to me from right to left. I might dial an unintended number—that which I use frequently and which is consequently more accessible to me. Other exam-

ples may be even more embarrassing, such as when a person wishes to refer to her husband but uses her ex-husband's name instead.

Another commonplace experience I have with people's faces may also be suggestive. For example, in a movie theater I might bump into an old student of mine, who is not highly familiar to me: I might have seen her in the classroom, among 40 other students, once a week, a while ago. Though her face "rings a bell," I might not be able to identify her. Out of her "natural" environment (the campus? the classroom?), I have difficulties telling who she is. The memory of her face is not salient enough to me-not foremost on my mind-to allow a rapid recognition. In contrast, however, I should experience no such difficulties if, under the same circumstances, I came across a highly familiar face such as my son's or my best friend's. I would surely identify them instantly, without having to rely on contextual information for help. Their faces are so familiarso salient—I can tell who they are regardless of appropriateness or naturalness of context.

The cartoon by Gary Larson (1984) shown in figure 1.1 may also be illustrative. It portrays the prototypical fear of monsters that almost every child has: a person is lying on the bed with some unfamiliar creatures underneath. However, while this might be our first preferred "reading" of the picture, this reading



"I've got it again, Larry ... an eerie feeling like there's something on top of the bed."

Figure 1.1 The monster under the bed. The Far Side® by Gary Larson © 1983 FarWorks, Inc. All Rights Reserved. Used with permission.

will have to be revisited when we see the caption: I've got it again, Larry . . . an eerie feeling like there's something on top of the bed. At this point, we have to discard our first adopted familiar (egocentric, human) perspective and take a less familiar, less frequent perspective—that of the monster who is under the bed. Having to reject a highly salient point of view in favor of an obscure one, which now seems just as "plausible," indeed allows for the funniness of the cartoon to get across.

The following poem, recently distributed on the Internet, explores the relations between such imbalanced perspectives. As the cartoon does, the poem ridicules our lack of flexibility in the face of variability and instability. It derides our automatic adoption of a "close to home" perspective, regardless of other likely possible alternatives:

(1) A Poem Written by an African Shakespeare Dear white fella

> Couple things you should know When I born, I black

When I grow up, I black

When I go in sun, I black

When I cold, I black

When I scared, I black

When I sick, I black

And when I die, I still black.

You white fella When you born, you pink When you grow up, you white When you go in sun, you red When you cold, you blue When you scared, you yellow When you sick, you green And when you die, you grey.

And you have the cheek to call me colored?

Setting out from a salient, egocentric perspective results in reducing the multifacetedness and versatility of an outgroup member to a single feature ('black'). Even in the face of fluctuating circumstances, for the ('white') Self, the ('black') Other is always homogenous and invariable.

Interestingly, this inflexibility of our salience-bound mind is also conveyed linguistically, through exposing the rigidity of word and concept meaning. Indeed, the first stanza challenges the conventional, most frequent, 'black' sense of colored. The second stanza amplifies the challenge by introducing the alternative, less salient (and also less euphemistic) senses of colored ('colorfulness', 'versatility'). Nevertheless, in spite of this massive attempt at revising the 'black' sense of colored, this sense is not shut out. When, at the end of the poem, we encounter colored ("And you have the cheek to call me colored?"-i.e., colorful), we do not activate 'colorful' exclusively. The salient 'black' sense is also invoked.

Language users are sensitive to degrees of salience. No wonder puns tend to spell out the less salient meaning of ambiguous words or expressions, trusting the lexical processor to activate the more salient meaning on its own accord and make the interplay between these meanings possible. Consider the advertisement *Don't leave without a good buy*, inscribed on an airport store, which calls for some interplay between the explicit, less salient meaning ('good bargain') of the homophone (a word having multiple meanings that share the same sound) and its more salient meaning ('good bye').

Or, consider how neologism in spelling by Maude Meehan (1991) (motivated by feminist awareness), inducing novel, nonsalient meanings (men oh pause, men struate), assumes that the salient meanings of these homophones ('menopause', 'menstruate') will pop up and highlight the differences between the various meanings:

#### (2) MEN OH PAUSE

and please take note I pen this semi-demi dithyramb in hope that I may set you men struate That sparse and very final show of red says only that a woman has bled not that she's dead or past delight in those hot-blooded lusty rites witch still transpire on any given night in grey and white haired womens beds

Some puns display the less salient meaning graphically, again, counting on the processor to avail the salient meaning spontaneously. For instance, the advertisement *GET A'HEAD* (fig. 1.2) portrays the less salient literal meaning of 'a head' graphically (and also via the misspelling), assuming that the salient meaning ('advance yourself') of *get ahead* will get across independently and automatically.

Jokes, as well as point stories—stories with a surprise ending—manipulate the very same tendency to opt for the more salient information first, although this time the process occurs in reverse order. Such discourses deliberately re-

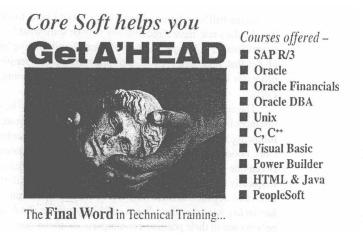


Figure 1.2 Get ahead.

frain from spelling out the intended, less salient meaning, counting on its slower activation. No wonder we are caught off guard, so to speak, ignoring the less salient yet contextually appropriate meaning while accessing the salient meaning first, only to find out that this interpretation requires reinterpretation. For instance, when, in *And Now for Something Completely Different*, Monty Python (McNaughton, 1972) talked about "baby snatchers," they manipulated our tendency to activate the more salient interpretation first in order to catch us by surprise when eventually we find out that it was the babies who did the snatching (of adults).

Note that the practice of deliberately refraining from making the less salient meaning explicit (as in jokes) relies on the same assumption underlying discourses that make it manifest (as in advertisements). They all assume that, while less salient meanings are slow, salient meanings are accessed rapidly. Indeed, this sort of "reflex" is pervasive, irrespective of the fact that it may, at times, be "stupid" (Fodor, 1983), ignoring or resisting sensible behavior.

The following artifact conversation, taken from *The Opposite of Sex*—a movie by Don Roos (1998)—may exemplify our tendency to attend to what is salient to us first, even at the cost of erring. (On a somewhat similar view, see also Fodor, 1983: 70.) The conversation in (3) takes place between Lucia and Bill. Bill is Lucia's late brother's lover, and since Lucia is practically his housemate, Bill's familiarity with her lifestyle is detailed and extensive, inclusive of the fact that she does not date at all:

(3) Lucia: I don't get sex.

Bill: You should get out more.

Lucia: No, I mean I don't understand sex. I don't get it, get it? Just it seems like a lot of trouble for not much.

Given Bill's background knowledge about Lucia, to interpret her as saying that she does not 'have'/'receive' sex may be somewhat less informative or relevant (Sperber & Wilson, 1986/1995) than the alternative 'understand' meaning of get that may make Bill's interpretation more "sensible." Bill, however, processes the meaning that is most accessible to him and fails to understand Lucia.

The superiority of salience over "sensibility" may be better illustrated by instances in which what we say falls short of what we mean. For example, in the following instances, the presence or absence of an explicit negation marker ('not') does not make a difference. In spite of the apparent inconsistency between (the negative) couldn't care less and (the positive) could care less, these frequently used collocations are interchangeable. Similarly, regardless of the (Hebrew) negation marker (lo), which appears in utterances such as eix she-lo kor'im lo ('whatever they do not call him') and ma she-lo iye ('whatever won't be') but not in their positive counterparts eix shekor'im lo ('whatever they call him') and ma she-iye ('whatever will be'), the meaning of the different members of each pair is identical.

Another instance of a linguistic behavior that "does not make sense" but which demonstrates that meaning salience reigns supreme, defying even "logic," is the evolution of the (Hebrew) masculine adjectives koosi and kooson. These were derived historically from the feminine koosit (meaning colloquially 'sexy'), originally a (chauvinist) "compliment" applicable to women only. Although the adjectives' root (koos) refers to a women's sex organ (originally an Arabic word), its masculine adjectives (koosi, kooson) are now used to refer to sexy men, in spite of the apparent incompatibility. Conventionality and frequency of use thus allow for the word's sense to be divorced from its originally "sensible" meaning and assume a life of its own.

The question, of course, is whether such prominent and independent meanings would be activated even in the presence of highly constraining contexts favoring an alternative meaning. A look at the context in (4), taken from "Miss Furr and Miss Skeene" by Gertrude Stein (1922/1975: 42), might give us an insight into the nature of the question. The context in (4) is heavily weighted in favor of the "happy" meaning of gay. However, it is dubious whether, under whatever circumstances, that meaning would be accessed exclusively. For instance, would a gay person, who is highly familiar with the homosexual sense of gay, activate only the contextually compatible "happy" sense of the word, while shutting out the contextually incompatible homosexual sense?

(4) Helen Furr had quite a pleasant home. Mrs. Furr was quite a pleasant woman. Mr. Furr was quite a pleasant man. Helen Furr had quite a pleasant voice quite worth cultivating. She worked to cultivate her voice. She did not find it gay living in the same place where she had always been living. She went to a place where some were cultivating something, voices and other things needing cultivating. She met Georgine Skeene there who was cultivating her voice which some thought was quite a pleasant one. Helen Furr and Georgine Skeene lived together then. Georgine Skeene liked travelling. Helen Furr did not care about travelling, she

liked to stay in one place and be gay there. They were together then and travelled to another place and stayed there and were gay there.

They stayed there and were gay there, not very gay there, just gay there. They were both gay there, they were regularly working there both of them cultivating their voices there, they were both gay there. Georgine Skeene was gay there and she was regular, regular in being gay, regular in not being gay, regular in being a gay one who was not being gay longer than was needed to be one being quite a gay one. They were both gay then there and both working there then.

The goal of *On Our Mind*, then, is to explore the effect of accessible meanings on speech production and comprehension. It aims to shed light, primarily empirically, on how, in addition to contextual information, salient meanings and senses of words and fixed expressions shape our linguistic behavior.

### 1. Processing Models

Few major proposals attempting to account for how we process information have captured the attention of researchers for the last two decades or so. While they agree that context affects language comprehension and production, they disagree on the temporal stage at which this information comes into play.

#### 1.1 The Interactionist View

For most contemporary theorists, context effects are primary (Bates, 1999; Bates & MacWhinney, 1989; MacWhinney, 1987). Accordingly, a highly constraining context, strongly biased in favor of a certain interpretation, interacts with lexical processes very early on and enables that interpretation to be tapped directly and, on most versions, exclusively. This constraint-based view, dubbed "the direct access" or "interactionist" model, thus, assumes a single mechanism that is sensitive to both linguistic and nonlinguistic information. Consequently, in a rich ecology, comprehension should proceed smoothly and seamlessly, selectively accessing the appropriate interpretation, without involving a contextually inappropriate stage initially (see also Wilson, 1998).

For instance, Vu, Kellas, and Paul (1998); Vu, Kellas, Metcalf, and Herman (2000); and Martin, Vu, Kellas, and Metcalf (1999) showed that when the sentence context (*The biologist wounded the bat*) was sufficiently constraining (compared to *He located the bat*) and weighted in favor of one meaning ('flying mammal') of an ambiguous word (*bat*), that meaning was activated exclusively, without activating the contextually inappropriate alternative meaning ('wooden racket'). This was true even when the appropriate meaning was the less frequent of the two.

Similarly, in a rich and supportive context, for a nonliteral utterance such as We have come a long way to be understood, the comprehender need not compute the contextually inappropriate literal meaning of the statement (which de-

picts individuals walking up a long path) before deriving its metaphoric meaning. Rather, the nonliteral meaning is accessed directly. Likewise, in an ironically biased context, *Thanks for your help* would be understood ironically directly, without having to analyze the sentence's literal interpretation first (Gibbs, 1986a,b, 1994, 2001, 2002; Sperber & Wilson, 1986/1995; but see Hamblin & Gibbs, 2001).

#### 1.2 The Modular View

The direct access view argues against "the modular view" (Fodor, 1983), which assumes modular architecture of mind alongside nonmodular, central, cognitive mechanisms. According to the modular view, linguistic information (among other types of knowledge) is encapsulated and does not have access to other systems' information when processed initially. Rather, initial input analyses are stimulus driven: they are automatic, rapid, and, on some traditional interpretations, exhaustive: all the meanings of a word are accessed upon encounter. The assumption that retrieval of meanings may be speedy and impervious to contextual and world knowledge effects predicts that contextual fit is accidental: at times, the output of the linguistic module would cohere with contextual information; at times, however, it would not and would require further inferential processes. When the latter holds, comprehension would not be smooth, but would initially go astray, with a later stage of revision and adjustment to contextual information. Integration, then, should not be always seamless.<sup>2</sup>

### 1.3 The Graded Salience Hypothesis

A third option is proposed by "the graded salience hypothesis" (Giora, 1997b, 1999b; Giora & Fein, 1999a,c; Giora, Fein, & Schwartz, 1998; Giora, Peleg, & Fein, 2001: Peleg, Giora, & Fein, 2001, in press; see also chapter 2 in this volume). This hypothesis shares a number of assumptions with the modular view, presuming two distinct mechanisms: one bottom-up, sensitive only to domainspecific (linguistic) information, and another, top-down, sensitive to contextual (both linguistic and extralinguistic) knowledge. However, unlike the traditional modular assumption, the graded salience hypothesis assumes that the modular, lexical access mechanism is ordered: more salient meanings—coded meanings foremost on our mind due to conventionality, frequency, familiarity, or prototypicality—are accessed faster than and reach sufficient levels of activation before less salient ones. According to the graded salience hypothesis, then, coded meanings would be accessed upon encounter, regardless of contextual information or authorial intent. Coded meanings of low salience, however, may not reach sufficient levels of activation to be visible in a context biased toward the more salient meaning of the word (for a different view, see Hillert & Swinney, 2001).

Contextual information may also affect comprehension immediately. A highly predictive context may avail meanings on its own accord very early on.

Nevertheless, it would not penetrate lexical accessing. Although it has a predictive role that may speed up derivation of the appropriate meaning, it would not obstruct inappropriate, coded meanings upon encounter of the lexical stimulus. Indeed, contextual information may be strong and even faster than lexical processes, so much so that it may avail meanings even before the relevant stimulus is 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 effective guessing and inferential processes. However, it does not interact with lexical processes but runs in parallel (Giora, Peleg, & Fein, 2001; Peleg et al., 2001, in press). As will be seen later, assuming a simultaneous operation of the encapsulated, linguistic mechanism and the integrative, central system mechanism 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. Neither does it assume that activation of a whole linguistic unit should be accomplished before contextual information comes into play. Rather, across 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.

Indeed, it is quite often the case that only contextually appropriate meanings are made available for comprehension. However, the claim here is that their exclusive activation is not a consequence of some selective compliance with contextual information but, rather, a function of their salience. For example, although my Hebrew name (Raxel) literally means 'ewe', this meaning most probably escapes anyone who uses my name as a referring expression. This, of course, has nothing to do with context effects. Rather, since the literal meaning of raxel is no longer in use, it is not salient enough, and hence of low accessibility. Similarly, it is quite plausible that the appropriate, computer appliance meaning of mouse would be the first to occur to a computer freak encountering mouse when involved in a discourse about computers. Contrary to appearances, however, directly accessing the appropriate meaning of mouse is not a matter of compliance with context but a matter of the accessibility of that meaning or concept in the mind of that comprehender. A novice learning how to use a software, however, might activate (and benefit from) the literal meanings of window, document, save, and so on, since to her these senses are more accessible than the appropriate ones. Thus, findings showing that the intended meaning of conventional utterances (idioms and indirect requests) has been processed without any analysis of the sentence's literal interpretation (Gibbs, 1994: 89–90) need not attest to context effects and may very well be affected by meaning salience.

In the next chapters I explain the notion of salience (chapter 2) and look into how salient meanings affect ambiguity resolution (chapter 3), comprehension and production of irony (chapter 4), metaphor and idiom (chapter 5), and

jokes (chapter 6). In chapter 7, I look into the role of salience in aesthetic novelty. Chapter 8 sheds a new light on contemporary research and attempts to reconcile the conflicting findings on literal and figurative language in terms of the graded salience hypothesis. Chapter 9 suggests possible extensions and implications of the various findings to theories of mind.

# Salience and Context

Linguists, cognitive psychologists, and psycholinguists interested in the cognitive processes involved in language comprehension and production need mostly make do with observing the products of such processes—the linguistic entities speakers either emit or attempt to comprehend. But, at times, even these entities may be highly revealing about the nature of the unobservable processes. Consider, for instance, the text in (1). This is a very old riddle. However, even today, when I present it to my students in an introductory course on feminist criticism, it is still as effective as it used to be 15 years ago and is not resolved on the spot:

(1) A young man and his father had a severe car accident. The father died, and the young man was rushed to the hospital. The surgeon at the emergency room refused to operate on him, saying: "I can't. He is my son." How is this possible?

What does that riddle tell us about the mind? What is (the mindset) ridiculed here? It is (among other things) our inability to instantly adjust to contextual information. Despite ample contextual evidence to the contrary (for instance, the riddle is presented in a course whose orientation is feminist; the immediate and explicit context eliminates the possibility that the surgeon is the young man's father, i.e., male<sup>1</sup>), it is the contextually inappropriate feature of surgeon ('male') that comes to mind first (and makes the riddle work). Redressing the "incongruity" or mismatch with context occurs somewhat later, when 'female'—the less frequent/stereotypical/salient feature of surgeon—is ac-

cessed (for empirical evidence, see Carreiras, Garnham, Oakhill, & Cain, 1996; Oakhill, Garnham, & Reynolds, 1999).

Consider yet another example. The other day, when I was shopping in the mall, I saw a shoe shop named *Body and Sole*. Despite ample contextual evidence supporting the shoe sense of *sole* (I was in a shopping mall, the window shop displayed shoes), I could not refrain from accessing 'soul'—the contextually inappropriate meaning of the homophon—alongside *sole*.

Or consider the following anecdote. Brigitte Nerlich, living in cold England, frequently whining about the cold weather, said to her husband one cold evening that she got cold feet. She knew while saying it that he would never in the world reply to her literal meaning and her literal intention to tell him that she had cold feet, but he would start a joke and comfort her and ask her what she was so worried about, which he naturally did (Nerlich & Clarke, 2001). Again, although the context was heavily biased in favor of the literal meaning of the idiom, it did not inhibit the idiomatic meaning.

This is one thing a product such as these anecdotes may tell us about the mind: the mind is rigid and, at times, error prone. Rather than accessing the compatible information exclusively, it opts for the most accessible information first. On most occasions, the accessible information is also correct—that is, contextually appropriate and compatible with the speaker's intention. On some occasions, however, it is not, and it will result in a search for a contextually compatible meaning.<sup>2</sup>

Has this folk theory about the mind (unveiled by jokes) gained any empirical support? Recently, Keysar (1994b, 1998, 2000); Keysar, Barr, Balin, and Paek (1998); and Keysar, Barr, Balin, and Brauner (2000) adduced evidence supporting the hypothesis that available information is accessed initially, regardless of contextual fit or speaker's intent. For instance, they showed that while searching for anaphor antecedent, the comprehender first accessed the candidate made available by the immediately preceding context, in spite of contextual information demonstrating that this referent could not be the one intended by the speaker. As illustration, consider the following example (taken from Keysar, Barr et al., 1998):

(2) It is evening, and Boris's young daughter is playing in the other room. Boris, who lives in Chicago, is thinking of calling his lover in Europe. He decides not to call because she is probably asleep given the transatlantic time difference. At that moment his wife returns home and asks, "Is she asleep?"

Keysar, Barr et al. (1998) showed that retrieving the contextually appropriate referent of *she*—the one that could be intended by the speaker, the daughter—was slowed down when there was a more accessible, albeit contextually inappropriate candidate: the lover. The slower reaction suggests that the accessible, albeit contextually inappropriate referent was accessed and interfered with the retrieval of the contextually compatible referent. Such findings are consistent with the view that contextual information does not affect initial access: it

does not preselect the contextually appropriate candidate (the daughter) while inhibiting the inappropriate one (the lover).

What is there about these meanings that makes them so privileged that even a highly biasing context is ineffective in blocking them? What makes them so salient, so foremost on our mind? The aim of this chapter is to point out some of the factors that make meanings unavoidable; more specifically, it aims to render precise the notion of salience so as to lay the foundation for the imminent discussions (in the following chapters) of the role of salience in language comprehension and production.

#### 1. Salience

According to the graded salience hypothesis (Giora, 1997b, 1999b), for information to be salient—to be foremost on one's mind—it needs to undergo consolidation.<sup>3</sup> that is, to be stored or coded in the mental lexicon (for instance, the 'male' and 'female' features of surgeon in (1); the idiomatic meaning 'fear' of cold feet). Stored information is superior to unstored information such as novel information or information inferable from context; while salient information is highly accessible, nonsalient information requires strongly supportive contextual information to become as accessible as salient information. Information derived on the basis of given/stored knowledge (such as script-based or thematicrelated knowledge), though predictable, is not salient and, at times, slower than when it is mentioned. For example, The beer was warm was read faster following We got some beer out of the trunk than following We checked the picnic supplies (Haviland & Clark, 1974; Gernsbacher, 1990; Sanford & Garrod, 1981). Similarly, meanings associated with a word that was mentioned in the text were shown to be primed instantly (about 300-350 milliseconds after onset of the prime). In contrast, inferences made on the basis of script or thematic information took longer to retrieve (about 500-750 milliseconds after onset of the prime; see Till, Mross, & Kintsch, 1988; Long, Oppy, & Seely, 1994). Specifically, immediately after reading The townspeople were amazed to find that all the buildings had collapsed except the mint, test words related to the contextually appropriate ('money') and contextually inappropriate ('candy') meanings of the target word (mint) were primed. However, a test word related to a thematic inference ('earthquake') took longer to retrieve (1000 milliseconds after onset of the prime).4

Salience is not an either-or notion, however. Rather, it admits degrees. The more frequent (Hogaboam & Perfetti, 1975; Liu, Bates, Powell, & Wulfeck, 1997; Neill, Hilliard, & Cooper, 1988), familiar (Blasko & Connine, 1993; Gentner & Wolff, 1997; Gernsbacher, 1984; Hintzman & Curran, 1994; Wiley & Rayner, 2000), conventional (Dascal, 1987, 1989; Gibbs, 1980, 1982, 1983; Ortony, Vondruska, Foss, & Jones, 1985: 587), or prototypical/stereotypical (Armstrong, Gleitman, & Gleitman, 1983; Rosch, 1973; Rosch & Mervis, 1975; Talmy, 2000: chap. 5) the information in the mind of the individual or in a

certain linguistic community, the more salient it is in that mind or among the community members. Let us consider the relation of these factors to salience in more detail.

#### 1.1 Frequency

Both (noun) meanings of bank—the 'financial institution' and the 'river-edge' meanings—are coded in the mental lexicon. However, for those of us who interact with commercial banks more often than with riverbanks, the institutional sense of bank is more frequently used and hence foremost on our minds—that is, salient. The more frequent the meaning, the quicker it is to retrieve (Galbraith & Taschman, 1969; Howes & Soloman, 1951; Just & Carpenter, 1980; Kawamoto, 1993; Rubenstein, Garfield, & Millikan, 1970; Seidenberg, Waters, Barnes, & Tanenhaus, 1984; Whaley, 1978; Zwitserlood, 1989).

The notion of frequency assumed here may be related to probability of occurrence recently investigated by scholars such as Burgess and Lund (1997); Carroll and Rooth (1998); Hindle and Rooth (1993); Jurafsky (1996); Landauer (1999); Landauer and Dumais (1997); Landauer, Foltz, & Laham (1998); and Lund, Burgess, and Atchley (1995). It may also be correlated with frequency of cooccurrence of meanings in the mental lexicon, which may but need not be correlated with frequency of occurrence at the referential or real world level. For example, the Palestinians' neologism, expressing their suspicion that "the Wye agreement" (occurred in Baltimore in October 1998) would turn into a peace of paper, hinges on the frequency of cooccurrence of piece and paper in the collocation a piece of paper. Although the discourse in which a peace of paper occurs abounds in repetitions of its components (peace, paper), the new coinage owes its meaningfulness to the frequency of cooccurrence of the linguistic components of a piece of paper that make up the lexicalized unit. This cooccurrence is frequent, regardless of the frequency of occurrence of concrete. real-life 'pieces' and 'paper', and it establishes the strength of association between the components of the collocation that accounts for its salience.<sup>5</sup>

#### 1.2 Familiarity

The meanings of 'plant' and 'structure' of tree might both be coded in the mental lexicon. However, while both are familiar among linguists, the (syntactic) 'structure' meaning of tree may be entirely unfamiliar to lay persons (although they might have the notion of family tree). Similarly, since I am a female feminist linguist, I am more experientially familiar with the generic sense of the feminine gender (e.g., her/she in The speaker must be relevant to her addressee. She should take his background knowledge into consideration) than most of my students. Meanings may be frequent yet less experientially familiar to an individual, and vice versa (Gernsbacher, 1984). Or they can be familiar though less frequent (Wiley & Rayner, 2000). The more familiar the meaning, the quicker it is to retrieve (Blasko & Connine, 1993; Gernsbacher, 1984; Hintzman & Curran, 1994).

### 1.3 Conventionality

I need four fifteen-cent stamps is a more conventional way of requesting stamps at the post office than I need the time when asking someone to tell you the time; A fine friend you are is more conventionally sarcastic than You are a fine/excellent friend (Gibbs, 1982); I'm afraid is a more conventional way of hedging a piece of bad news than reporting on one's feelings ('I am frightened'; see Erman & Warren, 2001). Conventionality may be viewed as "a relation among a linguistic regularity, a situation of use, and a population that has implicitly agreed to conform to that regularity in that situation out of preference for general uniformity, rather than because there is some obvious and compelling reason to conform to that regularity instead of some other" (Nunberg, Sag, & Wasow, 1994: 492n). Whether conventionality is motivated or arbitrary makes no difference at this point (for a discussion of motivated and arbitrary conventions, see Gibbs, 1986a, 1994: 319–358; Lewis, 1969). The more conventional the meaning, the quicker it is to retrieve.

#### 1.4 Prototypicality/Stereotypicality

On encountering *bird*, comprehenders are more likely to access a prototypical (robin, sparrow) than a marginal (chicken, penguin) member of the category of birds, given similar frequency of occurrence and sometimes regardless of frequency of occurrence of real-life members (Rosch, 1973; Rosch & Mervis, 1975). Indeed, participants were faster to verify that a stereotypical member (e.g., secretary, nurse) is a *woman* than a marginal member of the category (boss, surgeon, comedian; see Armstrong et al., 1983). The more prototypical the meaning, the quicker it is to retrieve.

Are these factors (familiarity, frequency, conventionality, and prototypicality) correlated? Is any of these factors more weighted? Could these factors be subsumed under one category? Familiarity seems the more crucial factor (see Gernsbacher, 1984; Wiley & Rayner, 2000). Indeed, it is plausible to assume that a meaning that is either frequent, conventional, or prototypical/stereotypical is likely to be more familiar to an individual than is an alternative meaning low on these variables. Yet, this is not necessarily the case. For instance, a meaning may be frequent but experientially unfamiliar to an individual (see Gernsbacher, 1984). Similarly, a meaning may be infrequent (the 'structure' meaning of tree) but familiar to an individual (Wiley & Rayner, 2000). Such are meanings of "forbidden" words (four-letter words, taboo words related to, for example, sex), which must be foremost on our mind in spite of their infrequent use. Further, a certain word or utterance may be a conventional manner of speech under certain circumstances without being used frequently (for example, I may know the conventional way of asking someone to tell the time, although this may be a less frequent or familiar experience since I spend most of my time with intimates among whom politeness is rare; see Brown & Levinson, 1987). Similarly, two members (chickens; pigeons) of the same category ('birds') may be equally frequent or familiar (such as among farmers who grow both birds).

However one (pigeon) may be a more salient member in the category (of birds) on account of its prototypicality.

#### 1.5 Gradability and Its Dynamics

Given that a salient sense of a word or an expression is the one directly computable from the mental lexicon irrespective of inferences drawn on the basis of contextual information, to what extent is salience a graded notion? It is graded in that it conceives of the lexicon as hierarchically structured. Although the various meanings of a word (bank) may be listed, one (the 'institutional' sense of bank) may be more salient, while the other (the 'riverside' sense of bank) may be less salient.

In addition, it is graded in that it assumes that the internal hierarchical structure of the lexicon is not fixed, but dynamic. For example, when the term the West Bank was first used by a metonymic extension to refer to the Palestinian territories occupied by Israel, it had a novel, nonsalient (though inferable) meaning. (The term probably started its career as 'the west bank of the Jordan River'). As of now, in a certain community of speakers, its salient meaning is its political sense (so much so that my spell-check won't accept the phrase in lowercase letters).

Or, take Swinney's (1979) notorious bug as another example (see also section 2.1). Possibly, it is undergoing a linguistic change—a reshuffling of salience—as might be deduced from the following title of an article in Newsweek: "Bugged by Y2K? Investors wonder what to do about their stocks, just in case the millennium bites" (August 1998: 48B). Comprehending bugged activates instantly both the 'concern' and the 'computer bug' meanings. As for the 'insect' meaning of bug, it seems quite plausible to assume that that meaning is now less salient, given that it catches us by surprise at the end of the sentence. Or take related examples. For people frequently interacting with computers, the salient meanings of desktop, windows, folder, document, paste, inbox, mail, download, and other such terms are metaphorical. Their literal meanings, once salient, have lost their priority. No wonder mail now denotes 'electronic' mail, while the 'post' system meaning is now a marked case, entitled snail mail.

#### 1.6 Can a Constituent Longer Than a Word Be Salient?

Though salience relates primarily to lexical meanings, it is also concerned with the meaning of an utterance or message. Could phrases or sentences have salient/stored meanings over and above those of their individual words? Could meanings of whole phrases (*spill the beans; get cold feet*) be salient, too? Are they accessed alone? Are they accessed alongside the meanings of the individual words (or morphemes) that make them up? Would that meaning override the salient meaning of the individual words? According to the graded salience hypothesis, they would reach similar levels of activation if they are equally salient (see also Clark, 1979; McGlone, Glucksberg, & Cacciari, 1994). They would not, in case they are unbalanced in terms of salience, as is the idiomatic meaning

of a familiar idiom relative to the literal interpretation of its individual words (Van de Voort & Vonk, 1995).

For illustration, consider the following example (originally in Hebrew, dated February 1998), in which both the intended, contextually compatible, literal meaning and the unintended, contextually incompatible, idiomatic meaning of the whole sentence were accessed:

Iddo and Omri (both aged 7 years, 8 months and both native speakers of Hebrew) are eating supper together. Iddo fetches himself a glass of juice out of the refrigerator.

(3) Omri: I want to drink, too.

Iddo's mother: Iddo, totci lo et ha-mic ('take the juice out [of the refrigerator] for him').

Iddo (laughingly): ha... le-hotci lo et ha-mic ('to take/squeeze the juice out of him'—a Hebrew idiom meaning 'drive one crazy'). (Mira Ariel, personal communication)

Iddo's laughter suggests that he has accessed the idiomatic meaning of the whole idiomatic phrase, as well as the compositional (literal) meaning. His amusement seems to disclose the computed discrepancy between the two (suggested, further, by the need he felt to repeat the idiom). Although the idiomatic meaning of the speaker's utterance was not the intended meaning, nor was it compatible with the context, it was not ignored; contextual information did not inhibit its activation.

This example suggests that a phrase or a sentence as a whole may have salient meaning(s) that will be activated automatically, irrespective of contextual information. Indeed, Giora and Fein (1999c) provide evidence consistent with the view that familiar idioms always avail their salient idiomatic meanings, even when contextually incompatible, as in contexts biasing their interpretation toward the literal meaning. Thus, subjects completed as many fragmented words related to the idiomatic meaning of familiar idioms (on one leg, meaning in Hebrew 'briefly') embedded in literally biasing contexts (In the zoo, I saw a stork standing on one leg) as they did words related to the literal (compositional) meaning. Further, they completed more fragmented words related to the idiomatic meaning of the idiom in the literally biasing context than they did words related to the literal meaning of the idiom in the idiomatically biased context. These findings suggest that the salient (idiomatic) meaning of the message was inescapable, despite the contextual bias towards the literal interpretation (see chapter 5 in this volume).

It is also possible, however, that highly salient idioms would activate their idiomatic meanings faster. Recall how Nerlich's complaint about her cold feet was entirely lost on her husband who took her only idiomatically. Empirical-findings also suggest that salient meanings of fixed expressions may override those of their individual words. In Gibbs (1980), idiomatic meanings were activated rapidly: they were faster to process when the context was biased toward the idiomatic than toward the literal interpretation. Apparently the frequency of kick the bucket as a fixed, idiomatic unit exceeds that of its literal (composi-

tional) interpretation, which relies on the frequency of the individual words, and thus makes the former more salient than the latter (see also Popiel & McRae, 1988; Van de Voort & Vonk, 1995). In the same vein, the degree of frozenness of idioms may index their salience—their storage as a lexicalized unit. In light of these findings, it seems plausible to assume that some phrases may have salient (lexicalized) meanings processed first or in parallel with the meanings of their individual words (McGlone, Glucksberg, & Cacciari, 1994; Nunberg, Sag, & Wasow, 1994; for a review, see Gibbs, 1994 and Gibbs & Gonzales, 1985).

Evidence from other familiar, conventionalized phrases/sentences seems to be consistent with the assumption that meanings of phrases or sentences may be stored as discrete units. For instance, familiar metaphorical phrases/sentences have been shown to be accessed initially both literally and metaphorically (Williams, 1992)<sup>8</sup> and avail both meanings in metaphorically as well as in literally biasing contexts (Giora & Fein, 1999c). Similarly, familiar ironic utterances (Tell me about it) were shown to be accessed initially both literally and ironically in the ironically as well as in the literally biasing contexts (Giora & Fein, 1999a), suggesting that their phrasal salient meanings were activated regardless of contextual information. In addition, familiar proverbs (The grass is greener on the other side of the fence) were shown to take equally long to read as their literal equivalents and to be processed literally and figuratively, regardless of contextual bias (Turner & Katz, 1997; see also Katz & Ferretti, 2000, 2001).

Would the compositional literal meaning of a sentence be salient? The compositional meaning of a sentence would not be salient, since it has not been lexicalized. However, to the extent that the intended, compositional meaning of a sentence relies on the salient meanings of its components, it would be easier to process than when it is not. Although the sum of the components' meanings is not listed in the mental lexicon as a discrete unit, when integrated into a sentence, this compositional meaning will be more accessible than is the compositional meaning of an equivalent sentence whose intended meaning hinges on the less salient meanings of its lexical entries. For instance, the meaning of the following New York Times headline (cited in Glucksberg, 2001: 16)—Price Soars for Eggs, Setting off a Debate on a Clinic's Ethics—is hard to get because its intended meaning relies on the less salient meanings of its components (the less salient 'human eggs' or 'ova' sense of eggs). Had eggs been replaced by 'human eggs', the difficulty of making up the compositional sense of the sentence would be reduced significantly.

Note, further, that the compositional meaning of a sentence (relying on the salient meanings of its component) is also more accessible than an inference or a conversational implicature drawn on its basis. Thus, the meaning of What a lovely day for a picnic in the context of a stormy day involves the implicature that the day is far from being lovely. This implicature is less accessible than the compositional meaning of the sentence, suggesting that the day is lovely (Giora, Fein, & Schwartz, 1998; Giora & Fein, 1999a). Recall, however, that some implicatures have become conventionalized, and are, therefore, salient (kick the bucket; Tell me about it).

On the basis of existing evidence, it is plausible to assume that some strings have stored meanings above and beyond the individual meanings of the words that make them up, but the measures prevalent in the literature have not been specialized to warrant a distinction between word and phrase/sentence level. For one, reading times of whole sentences might be particularly misleading in that they may mask differences in underlying cognitive process (Gibbs & Gerrig, 1989). Indeed, findings of equal reading times for metaphoric and literal target sentences which were taken to support an assumption of equivalent processes (Inhoff, Lima, & Carroll, 1984; Ortony, Schallert, Reynolds, & Antos, 1978) were shown to mask processing difficulty at the critical word level (Brisard, Frisson, & Sandra, 2001; Janus & Bever, 1985; chapter 8 in this volume). Similarly, attesting to the availability of a phrase or a sentence meaning by tapping word meaning (i.e., using words as probes) may not be sufficiently revealing about the meaning of the phrase or the sentence as a whole. As will be seen later, many of the probes used were related to a specific word in the sentence rather than to the meaning of the sentence as a whole (see Gibbs, 2002, for such a critique). In some cases this insensitivity is crucial. For example, if, when testing whether kick the bucket is interpreted either literally or idiomatically (or both), the probes used are 'die' (related to the idiomatic, message level sense) and 'jug' (literally associated with the word 'bucket'), then these probes measure different levels of meanings.

In sum, the graded salience hypothesis does not subscribe to any unified view of the mental lexicon as containing just one type of entry. It is consistent with the view that the lexicon accommodates specified, discrete entries for individual word meanings alongside underspecified meanings and senses. It is consistent with the view that word meanings may be distributed across layers of units representing semantic features in a network. It has no difficulty accepting that "each word is a constellation of semantic features so that subsets of these features participate in the meanings of other related words" (Vu, personal communication, 25 August 2000). Similarly, it has no difficulty accepting that the meanings of segments longer than a word (idioms, fixed expressions) are coded alongside the meanings of their components. It only postulates that coded meanings/senses are always activated upon encounter of the relevant stimulus, with less salient meanings not always reaching a threshold.

#### 1.7 Measures of Salience

How can meaning salience be measured, then? What norming studies can be used to test the salience—that is, the accessibility—of meanings of words or collocations out of context? Measures of word frequency and probability are already available (Burgess & Lund, 1997; Carroll, Davis, & Richman, 1976; Carroll & Rooth, 1998; Hindle & Rooth, 1993; Jurafsky, 1996; Kucera & Francis, 1967; Landauer & Dumais, 1997; Lund, Burgess, & Atchley, 1995). Alternatively, frequency or familiarity ratings of both meanings of words and phrases can be collected from native speakers. This seems to be the most frequent measure used in the literature.