Semantic Leaps

Frame-Shifting and Conceptual Blending in Meaning Construction

Seana Coulson

Semantic Leaps

Frame-Shifting and Conceptual Blending in Meaning Construction

Semantic Leaps explores how people combine knowledge from different domains to understand and express new ideas. Concentrating on dynamic aspects of on-line meaning construction, Coulson identifies two related sets of processes: frame-shifting and conceptual blending. Frame-shifting is semantic reanalysis in which existing elements in the contextual representation are reorganized into a new frame. Conceptual blending is a set of cognitive operations for combining partial cognitive models. By addressing linguistic phenomena often ignored in traditional meaning research, Coulson explains how processes of cross-domain mapping, frame-shifting, and conceptual blending enhance the explanatory adequacy of traditional frame-based systems for natural language processing. The focus is on how the constructive processes speakers use to assemble, link, and adapt simple cognitive models underlie a broad range of productive language behavior.

Seana Coulson is Assistant Professor in the Department of Cognitive Science at the University of California, San Diego.

Semantic Leaps

Frame-Shifting and Conceptual Blending in Meaning Construction

SEANA COULSON
University of California, San Diego

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press
The Edinburgh Building, Cambridge CB2 2RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9780521643610

First published 2001
This digitally printed first paperback version 2006

A catalogue record for this publication is available from the British Library

Library of Congress Cataloguing in Publication data

Coulson, Seana, 1966–

Semantic leaps : frame-shifting and conceptual blending in meaning construction / Seana Coulson.

p. cm.

Based on the author's thesis (doctoral – University of California, San Diego, 1997)

ISBN 0-521-64361-9

1. Semantics – Psychological aspects. 2. Frames (Linguistics) 3. Concepts. 4. Cognitive grammar. I. Title.

P325.5.P78 C68 2000
401′.43 – dc21

00-036046

ISBN-13 978-0-521-64361-0 hardback
ISBN-10 0-521-64361-9 hardback

ISBN-13 978-0-521-02783-0 paperback
ISBN-10 0-521-02783-7 paperback

for Steven

Contents

Acknowledgments

This book grew out of my 1997 Ph.D. thesis at the University of California, San Diego (UCSD). An interdisciplinary blend of topics and perspectives, in many ways it is a product of UCSD's Cognitive Science Department, where the hallways are filled with talk about everything from amnesia to aplysia, bumblebees to back-prop, the elegance of X-bar theory to the tractability of XOR. But perhaps more than anything else, it is a product of years of discussion with Gilles Fauconnier. Of course, most of these discussions were rather one-sided as Gilles barreled through the history of pragmatics, opined about the current state of cognitive science, or dove head-first into a new example. During these conversations, I did my best to scribble notes, in the hopes that I might later be able to re-create the feverish train of thought, and perhaps even decipher what Gilles meant. Over the course of time, I came to think about meaning in a whole new way.

Like many cognitive scientists, I had assumed that meaning was something speakers *compile* from linguistic input. The interesting questions in my mind were how contextual knowledge and pragmatic functions affected mechanical parsing operations and the activation of word meanings. However, in looking at actual examples of language use, it soon became clear that the impact of any given utterance could go well beyond its linguistically assembled meaning. In particular, in my 1992 study examining informants' discourse about the morality of abortion, I found that people's utterances draw on a vast set of background information, cultural knowledge, societal attitudes, covert analogies, and mappings between past, present, and future selves. While grammatical knowledge is often necessary to understand utterances in these debates, it is far from sufficient. As a result, I began to see that the apparently

trivial question of how speakers assemble *utterance* meaning was not simple – and, indeed, poses a profound and interesting problem.

My own personal Copernican revolution came with the understanding that utterance meaning is not *in* the speech signal, but actively constructed by speakers in response to linguistic and nonlinguistic cues. Like the paleontologist, the speaker has the task of combining different sorts of information to derive the overall meaning of the discourse event, to exploit her imaginative capacities and derive the life of the organism from its grammatical bones. I grew to realize that, rather than compiling meanings, people were using linguistic information to help them assemble cognitive models of the discourse event. Consequently, I set out to understand how people use contextual and background knowledge in the on-line construction of meaning. This book outlines what I have learned so far.

Traditionally, when approaching this issue, cognitive scientists have appealed to the idea of a frame, a hierarchically organized data structure for representing stereotypical knowledge. Unfortunately, the static nature of traditional frame-based systems stands in marked contrast to speakers' flexible use of language. To reconcile this tension, I point to two processes, *frame-shifting* and *conceptual blending*, intended to help solve problems faced by frame-based systems for reasoning and natural language processing. Frame-shifting is semantic reorganization that occurs when incoming information is inconsistent with an initial interpretation, and conceptual blending is a set of cognitive operations for combining frames from different domains.

I can lay claim to neither process. Though the term *frame-shifting* is mine, the phenomenon has been noted by many before, including Arthur Koestler, Marvin Minsky, and David Rumelhart. Conceptual blending theory, also known as *conceptual integration* theory, owes its origin to Mark Turner and Gilles Fauconnier. I have had the great fortune of being in contact with this dynamic duo as they developed the theory, and my own work has benefited greatly from our conversations and email correspondence. Similarly, my ideas about blending have profited from many discussions with Nili Mandelblit about the ubiquity of conceptual blending operations in the context of machine translation. I would like to thank them all.

One of my goals in *Semantic Leaps* is to supplement classic concepts from artificial intelligence with new ideas from cognitive semantics. Moreover, I want to use these ideas to develop models germane to psychologists and cognitive neuroscientists interested in language

processing. In my attempts to translate the abstract concerns of cognitive semantics into psychologically and neurologically plausible models, I received a great deal of assistance from my colleagues in and around the Kutas Lab at UCSD and the Van Petten ERP Laboratory at the University of Arizona. I would like to thank Marta Kutas, Robert Kluender, Jonathan King, Kara Federmeier, and Cyma Van Petten for sharing their knowledge of cognitive neuroscience. I have found our many meetings, arguments, and hours of speculative theorizing to be immensely stimulating. Special thanks are due to Marta and Cyma for their attempts to make an experimentalist out of me.

Thanks also to John Batali, Per Aage Brandt, Aaron Cicourel, Roy D'Andrade, Jonathan Folstein, Deborah Forster, Joe Grady, George Lakoff, Paul Maglio, Teenie Matlock, Todd Oakley, Sven Ostergaard, Tim Rohrer, and Eve Sweetser for their encouragement and advice. My original research was supported by a predoctoral fellowship from the McDonnell-Pew Center for Cognitive Neuroscience in San Diego. Subsequent work was supported by postdoctoral awards from the McDonnell-Pew Center for Cognitive Neuroscience in Tucson and an NIH National Research Service Award 1 F32 DC00355-01.

Lastly, I would like to thank my family for their love, support, and good humor.

CHAPTER ONE

Semantic Leaps

"One small step for [a] man, one giant leap for mankind."

– Neil Armstrong

Uttered by Neil Armstrong just before he set foot on the moon, this quote demonstrates how the same action can be interpreted in radically different ways. In this case, Armstrong alludes to a literal construal of his action and a more figurative one. Although it was a short distance from the space probe to the moon's surface, Armstrong's step would not have been possible without the combined efforts of those on the Apollo project to conceive and implement the mission. As such, Armstrong's journey to the moon was a sign of the scientific progress made in the twentieth century. The phrase "one giant leap" is not meant to be understood as a physical action, but rather as a metaphor for the scientific progress that had made the Apollo project possible.

Armstrong's action can also be interpreted metonymically. *Metonymy* is a figure of speech in which an object is referred to by one of its attributes, or by something with which it is commonly associated. In this case, Armstrong employs part-whole metonymy in which the whole species ("mankind") is evoked by one of its members. Because the journey represents the results of centuries of general scientific progress, and more specifically the efforts of the thousands of engineers, scientists, and bureaucrats involved in the Apollo program, Armstrong's voyage has been conceptualized as a public voyage on behalf of all earth-bound humanity. In the metonymic mapping, the leap corresponds to the entire journey from earth to the moon, which culminates in the final step.

The different possible interpretations of Armstrong's step provide a good illustration of how the significance of an action is dramatically

altered by a change in background assumptions. Like the interpretation of an action such as stepping onto the moon, interpreting the meaning and significance of natural language utterances depends crucially on contextual factors and background knowledge. Armstrong's utterance, for example, would be interpreted differently if he said it just before stepping into the shower. We shall see that for action and language alike, interpretation depends on the nature and scope of background assumptions.

In an effort to concentrate on tractable problems, researchers in semantics have traditionally focused on how to specify the literal meaning of phrases such as "one small step," thereby avoiding the more difficult problem of figurative meaning (see, e.g., Barwise & Etchemendy, 1989). However, more recent developments suggest the emphasis on these sorts of expressions may have been somewhat misplaced. For example, previous investigators have sought the locus of novelty in language and thought in combinatorial processes. And, while combinatorial processes are indispensable for explaining speakers' productive capacities, they do not provide a full account of novel language use. Moreover, the focus on expressions whose truth values can easily be specified has caused many investigators to overlook the constitutive role of context in meaning construction.

In this book, I locate speaker productivity in the comprehension mechanisms underlying semantic leaps – natural language constructions that yield nonobvious meanings. Characterizing the human ability to make semantic leaps will turn out to have implications for how speakers represent, organize, and use their knowledge in the production and interpretation of language. Just as Armstrong's small step relied on the heroic efforts of the members of the Apollo project, we will see how the interpretation of natural language recruits an elaborate set of meaning construction processes.

Semantic leaps is not a technical term, but, rather, a family of interesting natural language phenomena. It includes all sorts of nonstandard meanings absent from dictionaries and, typically, not computable by traditional parsers. Leaps include things such as metaphoric and metonymic expressions, hyperbole, understatement, and sarcastic quips. They also include things such as innuendo, subtle accusations, and the private meanings that can arise when people live or work closely together. Many leaps are necessary because of the way we deploy background knowledge in meaning construction. For example, until given information to the contrary, speakers tend to assume the pool in (1) is filled with water.

(1) Everyone had so much fun diving from the tree into the swimming pool, we decided to put in a little water.

Upon learning there *was* no water in the pool we do a double take, imagining the scenario in which everyone has fun diving into an empty pool. This reanalysis process is called *frame-shifting* and is the topic of Part I.

Part II deals with the sorts of leaps needed to combine information and create new concepts. In particular, Chapters 5 through 7 argue that information integration recruits *conceptual blending*, a set of noncompositional processes in which the imaginative capacities of meaning construction are invoked to produce emergent structure (Fauconnier & Turner, 1994; 1998). For instance, Chapter 5 addresses conceptual blending coded by modified noun phrases. While some examples, like "brown cow," and "red apple," don't involve much of a leap, others, like "caffeine headache," and "topless bar," rely extensively on background knowledge about the domains in question. Moreover, understanding the difference between "fake guns" and "real guns," or "stone lions" and "real lions," involves recruitment of knowledge about human practices such as deception and representation.

Some leaps are made possible by the human ability to conceptualize one thing in terms of another. For example, in (2), success in life has been conceptualized analogically in terms of success in baseball.

(2) He's a guy who was born on third base and thinks he hit a triple.

However, as if the leap from baseball to success weren't enough, note that the game in (2) has a woman giving birth in the ballpark. Chapter 6 considers this and other instances of metaphor and analogy that seem to involve a blend of elements and structure from the two domains that are being conceptualized.

Moreover, a similar blending process will be seen to operate in examples like (3).

(3) I wouldn't go out with you if you were the last man on earth.

Surely, "I won't go out with you," would suffice to make the speaker's point. However, by positing this armageddonlike scenario, she manages to drive that point home. Chapter 7 considers this and other counterfactual leaps in which speakers combine knowledge about the current state of affairs with knowledge necessary to understand an alternative reality. Overall, examples point to an important source of productive

language behavior that, until recently, has largely escaped the notice of researchers interested in meaning.

Part III addresses instances of frame-shifting and conceptual blending in real-life examples of moral discourse. Chapter 8 explores how speakers use terms to evoke culturally shared knowledge structures that can influence the way we evaluate the world around us. For example, the speaker in (4) draws on cultural understandings of action and responsibility in order to frame the rape victim's pregnancy as punishment.

(4) This bill is drawn so narrowly that it would punitively and without compassion further harm an Idaho woman who may find herself in the horrible, unthinkable position of confronting a pregnancy that resulted from rape or incest.

In contrast, the speaker in (5) presents an alternative framing.

(5) The unwanted pregnancy flows biologically from the sexual act, but not morally from it. ... Even degradation, shame, and emotional disruption are not the moral equivalent of life.

Moreover, (5) represents the use of another sort of culturally based knowledge structure, the *pragmatic scale*. Discussed in more detail in Chapter 9, pragmatic scales consist of objects or events ordered along a relevant semantic dimension. For example, the speaker in (5) evokes a scale of potentially acceptable reasons for abortion. We shall see how rejecting particular pragmatic scales can result in frame-shifting.

1.1 PRODUCTIVE LANGUAGE BEHAVIOR

1.1.1 Cats on Mats

Philosophers and linguists have each been impressed by different things about human language competence, and semantics has been shaped by both sorts of interests. For philosophers, the interesting thing about language is its *intentionality* or aboutness. How is it that an arbitrary set of symbols can represent things in the world? Linguists, on the other hand, are struck by the fact that there are virtually no limits to what competent speakers can say. We can understand sentences we've never heard before and can generate a potentially infinite number of novel ones.

These two issues, though potentially distinct, turn out to be related. The philosopher's problem of how "cat" can represent a cat has led to an

emphasis on truth and reference. The linguist's problem of how a finite brain can process a potentially infinite number of meanings has led to an emphasis on compositionality. A language is compositional if the meaning of a complex expression is systematically related to the meanings of its constituents (Frege, 1970; 1892). The interests of philosophers and linguists are complementary because semantics provides algorithms that show how the truth of a complex expression depends on the truth of its components. A good example is propositional logic, in which the truth or falsity of any well-formed formula can be systematically determined from the truth values of its component propositions and the truth tables for the connectives.

In keeping with the dual tasks of understanding compositionality and intentionality, part of what a semantic theory is supposed to account for is the relationship between the meaning of a sentence and facts about the world that support the truth of the proposition or propositions expressed by that sentence. Semanticists assume that to understand the meaning of a sentence in the indicative mode is to understand the kind of a world in which it could be used to make a true claim. Thus the main goal of semantics is to provide truth conditions for sentences.

On an objectivist account, understanding the meaning of an utterance such as (6) involves the ability to specify the conditions that make the propositions stated by the utterance true.

(6) The cat is on the mat.

Moreover, the truth conditions for a complex expression are determined systematically from those of simpler expressions. The lexical semanticist accounts for the meanings of individual words, and the compositional semanticist provides an account of how the meanings of words are composed into the meanings of their constituents, and how constituents are composed to yield the meaning of the sentence as a whole.

An important component of this account is that sentences have an independent level of representation, syntactic structure, which is used to compute an unambiguous semantic representation. Intuitively, the meaning of (6) would be derived by combining syntactic and semantic information about the components of the sentence. A "cat" is a carnivorous mammal long domesticated by humans and kept as a pet. A "mat" is a piece of coarse, woven, or plaited fabric used as a floor covering or a support. The locative relation "on" indicates that the external argument (in this case, "the cat") is in a position in contact with and supported by the top surface of the object (in this case, "the mat").

Accordingly, understanding the meaning of a sentence such as (6) involves the ability to specify the conditions that make the propositions stated by the utterance true. A semantic account should be able to differentiate the meanings of (6) and (7). Clearly, the state of the world (or worlds) in which (6) is true differs from that in which (7) is true.

(7) The mat is on the cat.

Further, this state differs in predictable ways concerning the spatial relationships between particular cats and particular mats. Formal semantics captures the intuition that anyone who claims to understand the meanings of (6) and (7) ought to be able to discriminate between a world in which (6) is true and a world in which (7) is true.

Of course, a real semanticist would not offer such a crude description as this as a semantic account. To relate words to the world, the objectivist semanticist (actually an *extensional semanticist*) employs set theory to model properties with sets of objects that possess those properties, and relations as ordered pairs (or triples, or quadruples, etc.) of related objects. Semantics provides a set of functions that convert constituents into sets that can be tested for truth in some model.

A semantic account can convert the English representation into an unambiguous semantic representation such as might be provided by predicate calculus. For example, for (8) we might derive the following representation.

(8) A cat is on a mat.

$(\exists x)(\exists y)((Cat(x)\&Mat(y))\&On(x,y))$

The semantic representation is then tested for truth in a set theoretical model such as the following:

$\{cats\} = \{Cat1, Cat2, Cat3\}$
$\{hats\} = \{Hat1, Hat2\}$
$\{mats\} = \{Mat1, Mat2, Mat3\}$
$\{flies\} = \{Fly1, Fly2\}$
$\{on\} = \{<Cat1, Mat3>, <Hat1, Cat1>, <Fly1, Hat1>\}$

In this model, the world consists of cats, hats, mats, and flies. To find the truth value for (8), we first must find the truth value for (9) and (10).

(9) $(\exists x)Cat(x)$ There is a cat.
(10) $(\exists y)Mat(y)$ There is a mat.

To do so, we first check our model to see whether the set of cats, {cats}, contains any members, and second whether the set of mats, {mats}, contains any members. Because both of these return true, we can go on to see whether any ordered pairs consisting of a cat and a mat are in the set {on}. Thus in this model, (8) is true but (11) is not.

(11) A mat is on a cat.

The task of the semanticist is to develop ways of transforming vague and ambiguous natural language statements into one or more unambiguous semantic representations, and to develop procedures for testing those representations for truth in a model. Adequacy of an account depends on the extent to which the output of the formalism satisfies our intuitions about the truth conditions and entailments of the original natural language sentence. Note that here I've developed only a cartoon-like explanation of how the semanticist operates. The reader interested in a thorough introduction to modern formal semantics should consult Dowty, Wall, & Peters (1981).

In any case, the formal approach to meaning has directly affected influential cognitive scientists such as Phillip Johnson-Laird, whose theory of mental models (Johnson-Laird, 1983) attempts to integrate model theoretic semantics with a theory of on-line language comprehension. Clearly, the demands of psychological plausibility preclude any straightforward mapping between formal semantics and language understanding: It is unlikely, for example, that people represent infinite sets. Rather, Johnson-Laird suggests that people construct models with individuals and *finite* sets of individuals. Johnson-Laird's mental models resemble the models of semantics in two ways: First, they represent the world in a structural way; second, one can ask whether a given sentence can be true in one of Johnson-Laird's models (Barwise & Etchemendy, 1989).

And, although he finds his belief problematic, Johnson-Laird (1989: 578) remains committed to the tenet of compositionality:

> A major problem confronting the present theory is to reconcile two important constraints on the process of comprehension. On the one hand, information from an utterance is integrated into the existing model as a function of the referential links, if any, between the utterance and the model; on the other hand, the interpretation of the sense of a sentence almost certainly depends on combining the senses of its constituent words according to the syntactic relations between them.

In fact, most semanticists working today have abandoned the idea that natural language is fully compositional. However, the assumptions

of the paradigm are deeply ingrained in cognitive science. Most notably, objectivist semantics has indirectly affected a huge number of researchers who have assumed that the productive nature of language use reflects the operation of something isomorphic to a compositional system (see Fodor & Pylyshyn, 1988, for a manifestolike review of the arguments for this position). The following sections show that this assumption – that compositionality in some form or other *must* obtain – has resulted in an overemphasis on combinatorial aspects of language comprehension and a corresponding underemphasis on the creative *noncompositional* mechanisms that also contribute to on-line meaning construction.

1.1.2 Mats on Cats

In common parlance, the term "meaning" applies to both linguistic expressions and particular utterances of those expressions. In contrast, semanticists reserve "meaning" for the expressions themselves. An expression such as (6) has a *meaning* that is independent of any particular context in which it might be uttered, while an utterance of this expression has a *content* that depends on things such as the particular cat to which the speaker is referring. One might hope for a semantic theory to relate expression meaning to utterance content. However, I shall argue that we should tackle the opposite problem: how meaning construction processes that generate utterance meanings might also explain why we think of words, phrases, and sentences as having meaning that is independent of any particular context.

Perhaps the main reason for shifting our focus from expressions to utterances is that meaning is never fully context-independent. Searle (1979), for instance, provides a number of examples that argue against the idea that the literal meaning of a sentence is context-independent. Rather, the literal meaning can be applied only relative to particular contextual assumptions. As we have seen, on the traditional account, answering the question of whether the cat is on the mat requires establishing which cat and which mat are being referred to, and determining whether the correct spatial relationship obtains. However, Searle (1979: 122) points to the case in which the cat in (6) is in outer space.

> ... suppose that the cat and the mat are in exactly the relations depicted only they are both floating freely in outer space, perhaps outside the Milky Way galaxy altogether. In such a situation the scene would be just as well depicted if

we turned the paper on edge or upside down since there is no gravitational field relative to which one is above the other. Is the cat still on the mat? And was the earth's gravitational field one of the things depicted in our [original] drawing?

Is the cat on the mat – or is the mat on the cat? Further, consider the case in which the mat is stiff, and has been stuck into the floor at an acute angle. Searle urges us to imagine that the cat is drugged, and lying on the top edge of the mat. Is the cat on the mat? Alternatively, the cat might be in the normal configuration but suspended by wires so that she places no pressure on the mat. Again, is the cat on the mat?

Although we *can* come up with an answer to the question of whether the cat is on the mat in each of these cases, the answers are far from determinate. Searle's point is precisely that we have to *come up with* an answer, rather than consulting the output of a context-invariant semantic function. While a semantic function might help to establish the truth of (6) in normal situations, these examples suggest that the function needs to be augmented in cases where normal conditions don't obtain. Thus Searle argues that giving truth conditions for an unambiguous sentence's literal meaning relies upon the assumption of the correct set of background conditions. Searle's argument thus suggests that truth conditions are not sufficient to specify all possible *sentence* meanings, let alone all *utterance* meanings.

Searle suggests that the establishment of the truth conditions for even the most transparent sentence is not context-independent, but relies on the assumption that conditions are normal. In this way, judgments of truth resemble judgments of category membership. While there is a great deal of disparity in people's truth and category-membership judgments concerning *atypical* cases, people tend to agree in their judgments of more typical cases (see Lakoff, 1987; Taylor, 1995). A similar process underlies both sorts of judgments. Context-independent meaning is an illusion based on the fact that a competent language user will create a context when none is provided.

Any given expression permits an indefinite number of interpretations that depend upon a speaker's understanding of the relevant context. Considering the immense variability in the utterance meaning of (6), Langacker (1987: 155) argues that a compositional account based on context-invariant meanings is unlikely to have much explanatory content.

Consider *The cat is on the mat.* Prototypically it describes a situation where a mat is spread out on the ground and a cat is sitting or lying on it. Already there is

indefinite variability, since the cat can be of any size, coloring, or subspecies; the mat is similarly variable; the cat can assume many different postures; and so on. But this is only the beginning. Possibly the mat is rolled up in a bundle and the cat is sitting or lying (etc.) on top of it. Maybe the operator of a slide show has just managed to project the image of a cat onto a mat being used for a makeshift screen. The sentence is appropriate in a mat factory where a worker has just finished decorating a mat with the outline of a feline. Conceivably a wrestler is holding an exhibition match with a tiger and has just succeeded in pinning its shoulders to the floor of the ring. The possibilities are obviously endless.

Pointing to the plethora of scenarios to which (6) might be applied, Langacker suggests there is a gap between the abstract specification of expression meaning and the rich content of particular utterances. By definition, there is some abstract commonality between the pet cat, the projected cat, and the tiger that licenses the use of the term "cat" in all three cases. However, Langacker seems to question whether there is any more content to this suggestion than the definition itself. Merely positing an unspecified abstract notion of cathood does little to explain it. Thus Langacker suggests the notion of context-independent meaning leaves an important component of productive language behavior unexplained.

After all, a speaker whose entire linguistic repertoire consisted of (6) could potentially use that philosophical chestnut to express an infinite number of *different* things. This property of utterances reflects human productivity just as surely as does center embedding. Moreover, while grammatical regularities often prove to be important for distinguishing between scenarios in which (6) does and does not apply, they do little to explain the immense variability in the configurations of cats on mats. Because the meanings provided by a compositional semantics leave this component of productive language behavior unexplained, they are necessarily incomplete. If we are ever to explain how speakers can relate (6) to all its variegated instantiations, we need to look beyond compositional mechanisms.

1.1.3 The Centrality of Mapping

Indeed, there are a number of problems that semanticists have assumed to be uninteresting, and subsidiary to the development of a theory of meaning, that turn out to be quite interesting for the cognitive scientist interested in a broader range of intelligent behavior. In particular, the way in which speakers construct utterance meaning on-line can potentially reveal general cognitive capacities. This section considers a

number of examples that demonstrate the variety of different uses to which people might put a term. I shall argue that these uses go beyond a straightforward relationship between words and the world. They are too complicated to be captured by objective criteria and will necessarily involve an appeal to this facet of the language user's cognitive capacities. In particular, word use reflects knowledge of typical situations, an ability to recognize abstract commonalities between domains, and an ability to conceptualize one domain in terms of another.

Consider the differences in the interpretation of "ball" in (12)–(14).

(12) Shaq dribbled past his defender and dunked the ball easily.
(13) Joe kicked the ball and smiled as it sailed through the uprights.
(14) Sammy hit the ball right out of the park.

In (12), the "ball" is a small leather sphere, orange or brown, filled with air. In (13), the "ball" is a brown, three-dimensional ellipsoid. In (14), the "ball" is a small white sphere, filled with yarn and covered with stitched cowhide. Clearly (12)–(14) each suggest a different game, and the speaker's understanding of "ball" varies as a function of the game in which it's embedded.

This is not to say that the uses of "ball" in (12)–(14) cannot be subsumed under a single abstract meaning. However, the existence of an abstract meaning does not help very much in explaining why the competent speaker interprets "ball" differently in (12), (13), and (14). It's not just that the reader *can* infer the presence of a basketball in (12), a football in (13), or a baseball in (14). She will automatically do so, and would presumably evidence surprise if (12) turned out to be a football, (13) a baseball, and (14) a basketball. Roth & Shoben (1983) demonstrate that sentence context can affect the time it takes to establish anaphoric reference between an exemplar (such as "football") and a category term (such as "ball"). Their data suggest that people entertain specific expectations about the nature of a vague term such as "ball" that are sensitive to the suitability of a particular exemplar in the context suggested by the sentence.

In answer to the question of how a listener adopts the appropriate understanding, Johnson-Laird (1981, 1983) has argued that it requires making inferences based on the *reference* of expressions rather than the sense or meaning. People's memory for sentences suggests that the product of the comprehension process is a model of the scenario depicted by a given sentence. Moreover, the model may contain details that are not specified by the meanings of the words in the actual sentence. For

example, Anderson & Ortony (1975) found "shark" served as a better recall cue than "fish" for (15).

(15) A fish attacked the swimmer.

Similarly, Garnham (1979) has found that the word "fried" is a better recall cue than "cooked" for (16), but not for (17).

(16) The housewife cooked the chips. [British English]
(17) The housewife cooked the peas.

These data suggest that in processing a sentence such as (18), the word "shark" would serve as a better recall cue than "it."

(18) It frightened the swimmer.

From this, however, we should not conclude that the listener must *always* consider "shark" as a possible meaning for "it." Johnson-Laird (1981) argues that comprehension of (18) requires the listener to imagine a scenario that is more specific than warranted by the meanings of the words. An adequate understanding of the utterance of a particular open-class word is inevitably underdetermined by an abstract characterization of its sense.

These sorts of observations imply that a sense-based explanation of language behavior will necessarily be incomplete. Whether or not a distinct *sense* is retrieved for lexical items, ultimately the language comprehension process involves the recruitment of knowledge about the likely nature of the particular referents of the utterances. Understanding "ball" in (12)–(14), for example, involves recruiting knowledge about the general scenario and the sort of ball that is most likely to be present in that scenario.

Consequently, a speaker who doesn't have the relevant cultural knowledge about baseball might entertain a different understanding of (14) than that of a sports-minded American. In contrast, an Indian speaker of English might be a bit puzzled by (14) and attempt to understand it as a strange variant on a cricket game. Regardless of whether there are shared truth functions between the Indian and the American speaker, there is a need to account for differences in their understanding. Given that the content and extent of a speaker's background knowledge affects her interpretation, it seems reasonable to postulate that integration with background knowledge is part of the interpretive process.

The dependency between a speaker's understanding of "ball" and her general conceptualization of the scenario will become more obvious if we consider how "ball" is understood in some less prototypical situations. Take, for example, the understanding of "ball" in the context of the video game *NBA Jam*. Here, the "ball" is a basketball, but not a small sphere. Rather it's a moving set of orange pixels, arranged in a roughly circular manner, covering the area of a few square centimeters. The ball in *NBA Jam* shares few physical features with a "real" basketball, but functions in a roughly analogous manner – that is, it is "dribbled" electronically, and "passed" and "shot" by pushing buttons.

Moreover, the ball in *NBA Jam* is not the only two-dimensional entity that we call a "ball." We might point to a 4-inch picture of Shaquille O'Neal on the sports page and remark (as in (19)) on how small the "ball" looks in his hand.

(19) That ball looks so small in his hand.

Or, we might point to a 40-foot billboard of Scottie Pippen and note the same thing.

Besides emphasizing the immense amount of contextual variability in the set of objects to which "ball" can be applied, these examples are also intended to highlight the systematicity in the disparate uses. For example, in both readings of (19), it is not the objective size of the "ball" that is being discussed, but the relative size of the ball and the man who is holding it. Similarly, in *NBA Jam*, there is a systematic set of correspondences between the objects and actions in the video game and the objects and actions in its three-dimensional counterpart. However, appreciating these correspondences requires the induction of an abstract schema.

Similarly, we can use the term "ball" to talk about the implications of televised sporting events. For example, during the Superbowl, Americans watching television might debate whether or not the "ball" was carried past the first-down marker. Alternatively, during the World Cup, those same Americans might scan the television screen in hopes of tracking the "ball." Just as knowledge of the game can help understand the significance of the game's events and aid the tracking of the ball, knowledge of the different games in which balls are used allows us to interpret statements about "balls" in ways that are far more specific than would be available from an abstract characterization of balls.

Now let's consider the case of two college students up late studying for an exam. Suddenly one crumples up a piece of paper and heaves it

at the wastepaper basket. The other decides to make a game of it and the two begin to shoot the "ball" at the "basket" from various spots around the room. After a while, one of them might utter (20).

(20) I'm getting good – I've gotten the ball into the basket the last three times.

Or after playing catch with a watermelon, a kid might offer (21) to explain what happened to dessert.

(21) We were playing catch till Jody dropped the ball.

Although neither of these cases involves a "real" game, we use the term "ball" to refer to the entity that functions similarly to a "real" ball in a "real" game.

We can also use the word "ball" to refer to entities that aren't directly involved in games. For example, imagine a bowling trophy with a small sculpture of a man, arm and leg extended, just about to release the "ball." The football coach might point to X's and O's on the chalkboard and say (22).

(22) By this point in the play, the ball is right here.

Recounting a playground basketball game with a much taller opponent, I might point to a spot above my head and say (23).

(23) So I'm going for the ball – which is way up here....

Or, a musician might point to his saxophone and say (24).

(24) I'm Michael Jordan and this is my ball.

While philosophers have assumed a correspondence theory of truth in which the core cases of meaning involve a straightforward relationship between words and objects in the world, these examples highlight the extent to which language use often relates only obliquely to things in the world. It would be quite hard (probably impossible) to account for the occurrence of "ball" on the basis of objective criteria. Besides balls made out of leather and pigskin, we have balls made from fruit, trophy metal, and crumpled-up paper. We also have balls printed on newsprint, painted on wooden billboards, and composed entirely of moving pixels on the CRT. Some balls have physical extension, while others correspond only to a blank spot on the coach's chalkboard, or an empty spot in the air. Finally, it would appear that while some balls are for playing games, others are for melodies.

However, these sorts of arguments are bound to draw a predictable set of replies from the objectivist. For example, she could argue, quite rightly, that the term "ball" is merely vague. Surely, careful analysis will reveal an abstract common meaning to the referential, metaphoric, and represented balls. However, if this is the case, it's not clear what the abstract common meaning buys the language user. To capture the full range of cases, the meaning would have to be so abstract as to overlap with the meaning of other things, such as "saxophone." Further, it leaves unanswered the question of how the listener derives particular understandings of utterances of "ball."

Another retort is to suggest that the three-dimensional objects that are so much fun to throw, catch, and kick are the *core* cases while the others are parasitic extensions. However, this raises the issue of defining exactly which cases are to count as the core uses. Is the ball in *NBA Jam* a "core" instance of a ball or a representational extension? What about a toddler's toy football? My own intuition is that it depends on who's playing with the ball. If two toddlers are playing catch with it, it might count as a "core" instance of a ball. However, if two adults play with it, it becomes more "peripheral." Or what about a toy football used with a G.I. Joe doll? The problem is that there don't seem to be objective criteria that definitively determine the set of utterances that reflect the core.

One possibility is to let statistics adjudicate the core. On such a scheme, speakers will consider the way in which a term is used most frequently to indicate the core meaning. However, defining the core in this way does not help us to identify the essence of the concept. Moreover, there remains the problem of characterizing how the core meaning of "ball" can be extended. Although extension sometimes depends on physical resemblance to the core cases, there's no consistent dimension of similarity in which a "peripheral" type of ball is supposed to be measured. Sometimes color is relevant, sometimes it's not. Sometimes size is relevant, sometimes it's not. Sometimes relative size is relevant, sometimes it's not. As these examples are intended to demonstrate, an object can't be defined in terms of objective criteria. It's not the case, for example, that representational extensions of "ball" are all two-dimensional circular entities. Indeed, it would appear that the only property shared by all the balls discussed here is that they have been construed as balls.

The fact that so many instances of "ball" appear to be motivated by abstract relational features suggests that an adequate account of those relationships necessarily involves an appeal to human understanding. In particular, language use seems to reflect the establishment

of correspondences or *mappings* between elements and relations in different contexts. For example, representations such as the photograph, the billboard, or the television screen afford a systematic mapping of terms for elements in games to apply to splotches of ink, paint, and light. Although the entities in the representation are ontologically quite distinct from those in the actual game, visual similarities exist. Further, the representation contains information about the relationships between entities in the picture that correspond to relationships between the entities they depict, such as the relative size of a basketball and a basketball player's hand.

A crucial component of understanding any usage of "ball" involves knowing what portion of our vast knowledge about balls is to be attended to, and what should be ignored. In the previous examples, I have suggested that the meaning of a word in a particular context is intimately related to the nature of the context. Moreover, capturing the way in which this understanding varies with the context is impossible if we restrict ourselves to objective criteria of objects in the world. However, if we look instead to the cognitive processes that underlie meaning construction, we can understand the contextual variability of utterance meaning as reflecting differential access to information in memory as a function of context.

In the case of the watermelon, for example, its construal as a ball involves recruiting knowledge about how people typically play with balls, while at the same time ignoring what balls are typically made of. The case of the crumpled-up piece of paper is similar, and illustrates how the flexible use of "ball" can be extended to incorporate other entities in the college dorm room, such as the "basket" and indeed the students' "game" itself. Further, the use of terms in this way relies upon the constitutive nature of the participants' construal of their activity. The watermelon can be understood as a "ball" in (21) because the kids treat it *as* a ball; similarly, the students understand their antics *as* a game. This sort of language usage relies on the ability both to appreciate and to create cross-domain mappings.

Although people's behavior does not seem to reflect the use of a context-invariant sense, it does seem to reflect knowledge of typical scenarios in cases that provide no explicit context. For example, the difference between the understanding of the balls in (12)–(14) results from the way in which meaning inherently rests on an implicit set of assumptions about the world in general and (in this case) games in

particular. Further, when a proper context is provided, speakers can accommodate even atypical scenarios. In all these cases, the speaker exploits her knowledge about the context in order to create a scenario that contains the right sort of a ball. So while referential applications of "ball" appear to differ in kind from metaphoric and representational applications of the term, they merely reflect different portions of the spectrum of meaning construction.

1.2 BACKGROUND

While we tend to think of words like "cat," "mat," and "ball" as having simple definitions, the previous examples demonstrate the variability of meanings in concrete utterances – even in artificially constructed scenarios. Moreover, the particular character of utterance content seems to depend on contextual factors in some nonrandom ways. On this view, interdependence of meaning and background knowledge is no accident. Contextual variation in meaning is ubiquitous because context is an inherent component in the meaning construction process. Background knowledge affects utterance meaning precisely because utterances are designed to evoke information from memory in a way that supports action and interaction in the physical and social world.

This view stems from frame semantics and more generally from the framework of cognitive semantics (Fillmore, 1982; Lakoff, 1988). Rather than viewing the description of the world as the central function of semantics, the cognitive semanticist treats meaning as a cognitive phenomenon invoked to construe conceptual content. The cognitive semanticist is less interested in the truth conditions for sentences than she is in the cognitive operations that realize the production and comprehension of linguistic utterances. Because cognitive activity mediates the relationship between words and the world, the study of meaning is the study of how words arise in the context of human activity, and how they are used to evoke mental representations.

1.2.1 Frame Semantics

Fillmore (1982) suggests that the meaning of many words relies on speakers' experience with the scenarios and social institutions they presuppose. For example, words such as "Tuesday" cannot even be defined without providing a certain amount of background information about

the more general concept of the organization of the week. Similarly the significance of the distinction between "week" and "weekend" arises because its meaning is motivated by the practice of the five-day work week. Thus Fillmore defines a frame as a system of categories whose structure is rooted in some motivating context. Words are defined with respect to a frame and perform a categorization that takes the frame for granted.

Fillmore (1982) emphasizes how meanings grow out of these motivating experiences. He argues that lexical semantics requires an account of how and why people use words as well as a characterization of the scenarios they presuppose. Fillmore construes frame semantics as a far-reaching research program with implications for lexical semantics, changes of word meaning, the creation of novel words, and even the assembly of the overall meaning of a text. In lexical semantics, for example, the aim is to characterize the motivating context for a particular word and explain how the word's meaning relates to that context. Thus words are defined with respect to frames and are used to evoke them.

Fillmore (1977) shows how a number of verbs, including "buy," "sell," and "pay," are related to one another by virtue of how they highlight certain aspects of the same Commercial Event frame. Indeed, a number of verbs can be understood as evoking the same frame, but accentuating (or *profiling*) the perspective, motives, or intentions of particular participants. Examples include "buy" versus "sell" (Fillmore, 1977), "give" versus "take" (Fisher, Hall, Rakowitz, & Gleitman, 1991), and "substitute" versus "replace" (Landau & Gleitman, 1985). Examples such as these accentuate how meaning cues the particular construal of events rather than merely providing speakers with an objective characterization. Moreover, they demonstrate how frames are motivated by human experiences, social institutions, and cultural practices.

Similarly, the meaning of "bachelor," classically defined as an unmarried man, can be shown to depend on the existence of background information grounded in social practice. Questioning whether, for example, the pope, Tarzan, or a gay man in a long-term relationship count as bachelors, Fillmore argues that the definition of "bachelor" as an unmarried man relies on the existence of a frame, or set of propositions that represent common assumptions about the normal course of a man's life in Western society. Much as our understanding of "on" in (6) and (7) involves a tacit assumption of a gravitational field, talk about bachelors involves implicit acceptance of background assumptions about the

course of a man's life in Western society. When these assumptions do not obtain for a particular man, such as Pope John Paul II, we are hesitant to apply the term "bachelor."

Lakoff (1987) emphasizes the idealized character of the background assumptions represented in frames. Because these assumptions involve a large degree of oversimplification, they apply more easily to certain segments of society than others. Thus, Lakoff (1987) shows how grounding lexical semantics in these idealized cognitive models provides an account of prototype effects in categorization. Further, the frames invoked by linguists to understand lexical semantics can be used to explain other cognitive tasks, including reasoning, problem solving, and making judgments about the behavior of others.

Frames as Data Structures

Fillmore's linguistically motivated account of a frame is paralleled by similar suggestions from researchers in other branches of cognitive science. In the field of artificial intelligence, Minsky (1975) proposed the term *frame* for a data structure used to represent commonly encountered, stereotyped situations. Minsky offered a child's birthday party as an example of the sort of thing a frame might be used to represent. A birthday party frame includes *slots*, such as food, games, and presents, that specify general features of the event. Slots are bound to *fillers*, or representations of the particulars of a situation. In a process called *slot-filling*, slots such as food are bound to fillers such as cake and ice cream.

The efficiency of frames as data structures derives from the organization of general slots that can be bound to particular fillers. This provides a means of organizing the similarities as well as the differences that exist between our various experiences of children's birthday parties. Activating a frame creates expectations about important aspects of the context by directing the agent to fill the slots with available information. Moreover, the real power of frames derives from the use of *default values*, that consist of the most typical and/or frequent filler for each slot. If information about the actual slot-filler is unavailable, a slot is assumed to be filled by the default.

The beauty of Minsky's suggestion was the notion that frames contain the sorts of information needed to understand a particular sort of event or scenario, as well as default information about the most probable fillers for any given slot. Moreover, the representational structure in frames proved to be valuable in building computer systems to

understand natural language. In the course of developing a system that could understand simple stories, Schank & Abelson (1977) postulated scripts as analogous to Minsky's frames. Scripts represent stereotyped sequences of events such as going to a restaurant, and contain slots that are either filled by binding the particular fillers manifest in the situation at hand or by instantiating the default value for any particular slot.

While the frame type data structure was largely motivated by considerations of representational utility, cognitive psychologists have found considerable evidence that people utilize frames, or schemata (schemas) as they are called in the psychological literature, in a variety of cognitive tasks. People use frames in perception, planning, and memory for events (Barsalou, 1992). Moreover, frames have been used to explain human ability to make inferences in complex situations, to make default assumptions about unmentioned aspects of situations, and to make predictions about the consequences of actions.*

In cognitive semantics, meaning does not involve mapping from terms to objects, actions, and events in the world. Rather, words designate elements and relations in frames that *may* represent objective aspects of reality, but need not (Fauconnier, 1997). So, instead of positing one set of processes to track correspondences between terms and objects, and another for terms and various abstract, relational properties, we can see the former as following trivially from the latter. That is, words are always understood as setting up frames, regardless of whether those frames apply to actual, representational, or hypothetical referents. Apparent core cases – where frames apply directly to real-world referents – are merely a subset of a more inclusive phenomenon.

Consequently, there is nothing to be gained from treating these so-called core meanings as more fundamental than more exotic-looking utterance meanings. For one thing, doing away with the notion of core meanings dissolves the problem of how to circumscribe the core and parameterize its extensions. Moreover, addressing more exotic cases of meaning construction has led to an important locus of generativity in language production, namely the human capacity to map within and between frames and scenarios.

* Following Fillmore (1982), I will use the term *frame* as a cover term for a whole set of related concepts, including script, schema, scenario, idealized cognitive models, and folk theory. Although differences exist in the scope of these constructs, they are all used to represent structured background knowledge, have important experiential character, and so forth.

1.2.2 Mental Space Theory

Mental space theory (Fauconnier, 1994) is a theory of *referential structure*, a level of conceptual organization between the situation being described and the linguistic structures that describe it (Langacker, 1993). In this framework, words do not refer directly to entities in the world (nor to elements in a set theoretical model). Rather, linguistic cues prompt speakers to set up elements in a referential structure that may or may not refer to objects in the world (see, e.g., Fauconnier, 1997). Created to represent discourse that concerns various factive (and non-factive) contexts, mental spaces can be thought of as temporary containers for relevant information about a particular domain.

A mental space contains a partial representation of the entities and relations of a particular scenario as perceived, imagined, remembered, or otherwise understood by a speaker. This representation typically includes *elements* to represent each of the discourse entities, and simple frames to represent the relationships that exist between them. For example, to represent the information in (25), we would set up a mental space with two elements, structured with a Married Couple frame (see figure 1.1).

(25) Paul is Jaimie's husband.

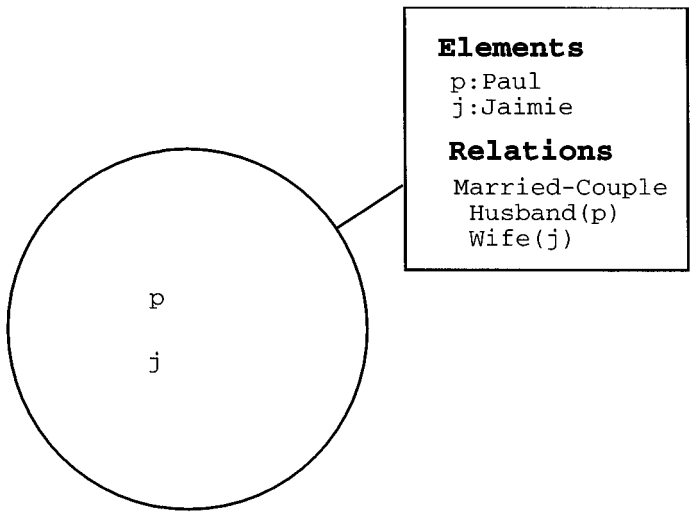

Figure 1.1 A Mental Space. The circle corresponds to the mental space set up to represent "Paul is Jaimie's husband." The attached square contains a representation of the conceptual structure in this space.

Mental space theory also affords a way to represent information about an entity in contexts where its properties might change. For example, in the following story, Arnold has different properties in real life than he does on his TV show.

Arnold is an actor.
He plays a mercenary on TV.
But in real life, he's a pacifist and a vegetarian.

To represent a story like this, spaces are used to partition information about Arnold into those properties that pertain to the actor, and those that pertain to his character on the TV show (see figure 1.2). Each space contains a partial representation of relevant elements and relations structured by a delimited set of assumptions. While different spaces can contain disparate information about the same elements, each individual space contains a logically coherent representation.

Of course, without an adequate understanding of the correspondences between elements in the different spaces, partitioning the information in the discourse has limited utility. To represent these sorts of relationships, speakers exploit abstract structure in each of the spaces to establish *links* (or *connectors*) between *counterparts*, or corresponding elements. For example, in figure 1.2, the element (a) corresponding to Arnold in the base space is linked to his counterpart (a') in the Television space. The link represents the counterpart mapping between the actor slot and the character slot in their respective frames. The combination of frames and links between elements in different spaces considerably simplifies the speaker's representational task. Instead of formulating an elaborate integrated representation, the speaker merely maintains a number of partial structures and simple models of the links between them (Dinsmore, 1991).

Spaces represent such diverse things as hypothetical scenarios, beliefs, quantified domains, thematically defined domains, fictional scenarios, and situations located in time and space. As discourse unfolds, the language user extends existing spaces by adding new elements and relations to the cognitive models already evoked. The need for new spaces can be prompted either explicitly by *space builders* or implicitly by the knowledge of the current importance of an alternative domain of reference. Examples of space builders include adverbials (e.g., *maybe, really, possibly*), prepositional phrases (*in Japan, in 1999, behind the house,*

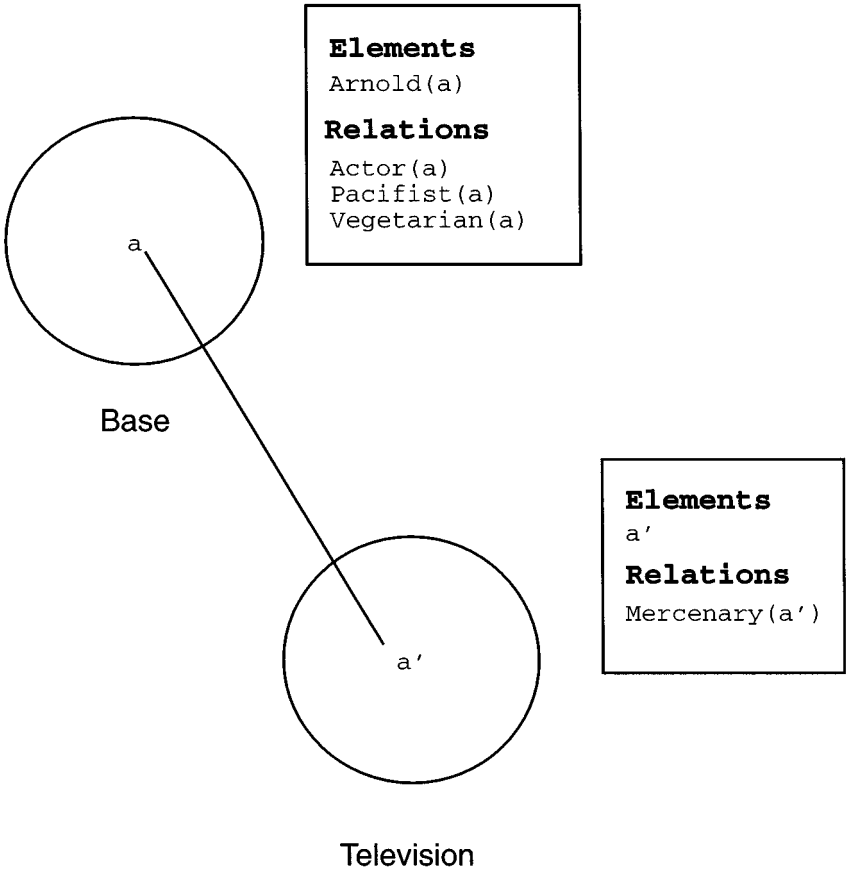

Figure 1.2 A Mental Space Configuration. The configuration of a base space and a television space enables the speaker to represent both the disparate facts about Arnold and the character he plays on TV and the connection between Arnold and his character (the link).

on TV), subject-verb complexes (*Jen believes, Lee thinks, Chuck argues*), and conjunctions (*if, or, because*) (Fauconnier, 1988).

A new space is also set up when utterances concern objects or events that require different background assumptions from those of the current space. For example, if extending a space would result in the representation of contradictory information, a new space is set up. As much as possible, the new space will maintain aspects of the background structure in the previous space. Although numerous grammatical clues aid the speaker to construct mental space configurations, the meaning

construction process is underdetermined by the linguistic information. For a given utterance, there are often multiple configurations of spaces that are consistent with the linguistic information (see Fauconnier, 1997, for examples). However, in practice, speakers are usually able to exploit background knowledge and contextual information to set up the correct configuration.

Speakers thus create a two-dimensional lattice of spaces made up of parent and child spaces. Except for the base space, all spaces have parents. Further, a parent space can have a number of child spaces. It's possible to move from higher to lower spaces in the lattice, and also from lower to higher spaces. The *focus* is the space in which meaning is currently being constructed, and the *viewpoint* is the space from which other spaces can be accessed. Meaning construction thus consists of mapping cognitive models from space to space while keeping track of the links between spaces and between elements and their counterparts. In contrast to traditional approaches to meaning construction, mental space theory assigns the bulk of the cognitive work to the mappings and links between domains rather than the derivation of a logical representation of sentence meaning.

Although motivated by linguistic data, mental spaces are not specifically linguistic in nature. Rather, they reflect the operation of more general cognitive processes. Kids' ability to play catch with a watermelon or to play basketball with crumpled-up paper invokes the same sort of linking and partitioning ability involved in language comprehension. These sorts of activities are called symbolic, pretend, and imaginative play by developmental psychologists and are considered an important aspect in normal socialization (see Flavell, 1999, or Singer, 1995, for review). Leslie (1987) has proposed that imaginative play is critical for the development of a child's understanding of theory of mind (that is, the fact that others can have different beliefs than oneself) and contributes to her general metarepresentational capacity.

Similarly, the idea of a slot that can be filled with a variety of different individuals has a parallel in social life: a *role* such as mother or boss. Roles in a frame can frequently be expanded into their own frames, depending on the granularity of the focus. Just as individuals have properties, roles have properties. Moreover, the properties of roles are to a certain extent distinct from those of the individuals who fill them. To illustrate, Sweetser & Fauconnier (1996) point to (26):

(26) The president has been commander in chief since 1776.

This is a property of the role of president of the United States, but clearly not of its current filler. Just as actors have different properties from the roles they play on stage, people have different properties from the roles they play in various frames. For example, George Washington, besides being the president, was a husband, a father, and a grandfather. Examples of referentially ambiguous statements like these accentuate the extent to which meaning depends on our ability to delimit the orbit of reference.

Links and Connectors

If we understand meaning construction as the process of structuring and linking mental spaces, it suggests that meaning always emerges from understanding in a particular context. In the absence of an explicit context, speakers will create one for themselves based on their knowledge of typical situations and their default values. For example, in the absence of a particular context, we might say that "Foucault" refers to a particular dead white male philosopher. However, in many contexts, "Foucault" refers to something entirely different. For instance, if the book on top of my television set is *The Order of Things*, we might refer to the book by the name of its author, Foucault.

In objectivist semantics, the problem is to provide a meaning that is abstract enough to account for the contextual variability in its application. How can "Foucault" refer both to a philosopher and to a book? What sort of a function could relate "Foucault" to both sorts of entities without creating a function that overgenerates? The cognitive semanticist's answer is to skirt the problem of relating words to the world, looking instead to the sorts of cognitive operations that would allow people to evoke different representations in response to hearing "Foucault" in different contexts.

Nunberg (1978) provides an account of these *pragmatic functions* that allows us to use a term from one domain (the *trigger*) to refer to an entity in another domain (the *target*). Here the connection between authors and their books allows us to use a term that customarily refers to a dead white philosopher to refer to the book on top of my television set. Although the default reference of the trigger (a dead white philosopher) is very different from the default reference of the target (the book on top of the TV), there is a nonarbitrary relationship between them that allows us to map between domains. Moreover, this relationship is mediated by both general and situation-specific background knowledge. In mental space theory, the possibility of using a term from one space to refer to a linked

element in another domain is known as the *Access Principle*. Because elements in different spaces can be linked by identity, similarity, analogy, metonymy, and other pragmatic functions, this principle accounts for much of the contextual variability in word use.

Objectivist semantics predicts objects to be named on the basis of their properties. Since Joe Schmoe shares few properties with the ham sandwich he's currently consuming, examples such as (27) are hard to account for.

(27) The ham sandwich wants another Coke.

In contrast, cognitive semantics predicts speakers will exploit properties of conceptual structure, regardless of its relationship to objective properties. Sweetser & Fauconnier (1996) pose the term *frame metonymy* to account for the relationship between metonymic triggers and targets. We can refer to customers in a restaurant by the food they order because both customers and food orders occur in the restaurant frame. Similarly, because authors and their books occur together in a frame, we can refer to *The Order of Things* by the filler of the author slot, Foucault.

Rather than viewing context as something that modulates a context-invariant meaning, we can view what *appears to be* the context-invariant meaning as a product of the default context. While the use of defaults might produce the appearance of a context-invariant meaning, this meaning is no different in kind from the meaning that emerges from the use of less frequent background assumptions. The general dependency between meaning and background assumptions obtains, irrespective of the content of the background assumptions that operate at a given time.

What Motivates Mental Space Theory?

Mental space theory was set up to answer questions about indirect reference and referential opacity. Imagine that Agent Mulder thinks his partner Agent Scully is the American ecoterrorist known as the Unabomber. The astute reader, however, will know that the real Unabomber is not Scully, but a man named Ted. On traditional accounts, sentences like these, with a clause embedded in a psychological predicate, are known as referentially opaque contexts. The problem for extensional accounts is that while (28) and (30) have the same referents, it is (28) and (29) that have the same truth value.

(28) Mulder thinks he had lunch with the Unabomber today.
(29) Mulder thinks he had lunch with Scully today.
(30) Mulder thinks he had lunch with Ted today.

Because people's conceptions of reality are not always accurate, contexts that contain propositional attitudes (sentences about beliefs, hopes, desires, and so forth) create referential ambiguities.

Jackendoff (1975) notes similarities between ambiguities present in propositional attitude sentences and those in sentences about pictorial representations as in (31).

(31) In the picture, the girl with green eyes has brown eyes.

Fauconnier (1994) argues that these ambiguities are actually quite general and can be found in any case where domain connections apply. Moreover, both sorts of examples (propositional attitudes and representations) can be handled by partitioning the disparate referential information into different spaces.

In example (28), the information is partitioned into two spaces: one for Mulder's beliefs, and one for the base (reality) space (see figure 1.3). In the belief space, Mulder had lunch with the Unabomber; in reality space, he had lunch with Agent Scully. From the viewpoint of Mulder's belief space, we can talk about "the Unabomber" by using the term

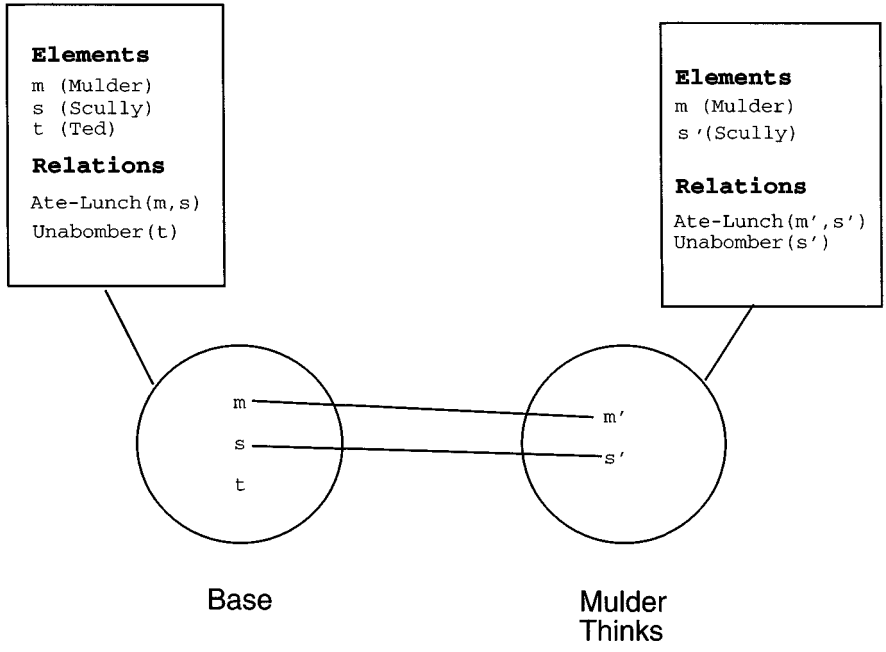

Figure 1.3 Mulder's Erroneous Beliefs. This mental space configuration has partitioned information pertaining to reality (Base Space) and information pertaining to Mulder's beliefs (Mulder Thinks).

"Scully": We can *access* the Unabomber from "Scully" (and vice versa) because Mulder's model that Scully = Unabomber affords a link between the two entities. From the viewpoint of reality, knowledge that Ted = Unabomber affords a link between Ted and the Unabomber, which makes either term accessible from the other. However, from the viewpoint of Mulder's belief space, there is no link between Ted and the Unabomber, so "Ted" is not cognitively accessible from "the Unabomber."

Similarly, for pictures or other representations, sentences such as (31) are partitioned into representation space and reality space. In one space the girl has green eyes, and in the other her counterpart has brown eyes. Knowledge of the identity of the two makes either cognitively accessible from the other space. On the default reading of (31), the girl has *green* eyes in reality space, and *brown* eyes in the picture. Because of the link between counterparts, a description of the real girl can be used to refer to the girl in the picture.

The basic insight of mental space theory is that apparently anomalous phenomena such as referential opacity are not actually different in kind from the way reference normally occurs. Words do not refer directly to entities in the world. Rather, linguistic cues allow the language user to set up elements in mental spaces that may or may not refer to objects in the world. Moreover, because discourse demands frequently require repeated reference to the same entity, elements are often accessed from spaces that have been set up previously. The sorts of ambiguities that result from the different properties that elements can have in belief and reality spaces can also occur when opening time spaces, hypothetical spaces, and other sorts of representational spaces.

1.3 CONNECTIONS

I have suggested that meaning construction crucially involves constructing frames whose elements and relations are linked across mental spaces. However, viewing frames as central components of language and reasoning presents certain problems. First, there is a need to reconcile the simplified nature of frame-type representations with the complex nature of the cognitive tasks for which they are employed. Similarly, there is a tension between the static nature of traditional frames and the flexibility people demonstrate in their use of language. For all the representational efficiency of traditional frame-based systems in artificial intelligence, their capacity pales in comparison to the fluid, locally creative character of human discourse.

Brachman (1985) has pointed to a number of inadequacies of frames as knowledge representation structures. For example, one of the supposed advantages of frames is that a single frame can represent a potentially infinite number of related objects. However, the need to accommodate all possible variations from the stereotype may ultimately undermine the frame's utility for understanding the general case. To illustrate, Firebaugh (1988) questions how to best represent Clyde, an elephant who has lost his trunk. To be sure, we can add a slot to the existing elephant frame to represent the number of trunks. However, if we pursue such a strategy for all possible elephant variations, it will quickly result in frames so unwieldy as to contravene their initial representational motivation.

The first step in solving this problem is to distinguish between the representation of knowledge in long-term memory from the representations employed in on-line comprehension and reasoning tasks. While people seem to have an immense amount of knowledge about particular topics, only a small portion of that knowledge is available at any one time. Moreover, as I argued in section 1.2, on mental space theory, the demands of discourse comprehension involve organizing knowledge about a single object into discrete packets of related, yet often distinct, information. While our long-term memory might have a slot for the number of trunks in our representation of elephants, this facet of our knowledge base would be recruited only in relevant contexts concerning Clyde.

Another shortcoming of frames, Brachman (1985) argues, is that they preclude the representation of composite descriptions that inherit structure from each of their components. However, in a discussion of the analogical nature of many jokes, Hofstadter & Gabora (1989) pose the term *frame blend* for a frame whose elements and relations are constructed from a combination of two frames that share some abstract structure. Fauconnier & Turner (1994;1998) have shown how frame blends occur in a wide variety of cognitive phenomena, and have developed an elaborate theory of conceptual blending (also known as *conceptual integration*) to explain the representation of composite descriptions.

The central tenet of the approach developed in the succeeding chapters is that the on-line construction of meaning is not a matter of manipulating preexisting representations from memory, but rather an active process in which the speaker integrates ongoing aspects of perceptual and conceptual information with more abstract information available from long-term memory. So, to understand the plight of Clyde the

elephant, we do not merely retrieve an elephant frame and bind a fixed set of slots. Rather, we construct a partial representation of Clyde by integrating contextually available information with selected aspects of our knowledge of elephants, and potentially aspects of knowledge about a whole range of other things. The dynamic character of on-line meaning construction thus arises because language users are continuously and creatively building and blending frames rather than merely retrieving and instantiating them.

In an attempt to reconcile the tension between the static nature of frames with the flexibility people manifest in language use, I hope to show that the analysis of semantic leaps can reveal the nature of information in frames recruited from memory, and how this information is integrated with local contextual information in the course of meaning construction. In particular, processes for cross-domain mapping, conceptual blending, and frame-shifting are proposed to enhance the explanatory adequacy of traditional frame-based systems. The diversity of topics in the following chapters is a testament to the generality of the processes of conceptual integration, and a demonstration of the many levels of cognition at which they operate.

Part I addresses Brachman's first complaint about the representational capacity of frames by supplementing them with ideas from mental space theory. The chapters in Part I explore the tenebility of a frame-based approach to meaning construction from the perspective of the linguist, the psycholinguist, the cognitive modeler, and the cognitive neuroscientist. Further, to address Brachman's second complaint about how to combine frames, Part II explains how conceptual blending operates in noun modification, metaphor, analogy, conditionals, and counterfactuals. In Part III, I show how frame-shifting and conceptual blending are evident in people's discourse about social issues. We will see how a dynamic theory of meaning can help us understand the semantic leaps that underlie our ability to comprehend simple jokes, to do the moral reasoning necessary for social discourse, to combine concepts, to understand analogy, and to reason about things that might have been.

PART ONE

FRAME-SHIFTING

On standard views of sentence processing, meaning construction results largely from *parsing*. A string is decomposed into representations of syntactic or structural information, and of lexical semantics, or word meanings. The meaning construction system on such a view consists of two discrete components, one devoted to word meaning and one devoted to syntactic structure. Parsing is accomplished by combining these two sorts of information to assemble a context-invariant meaning. On such approaches, pragmatics is knowledge that enables speakers to adapt context-invariant meanings to the situation at hand. In contrast, Part I outlines a constructivist comprehension system in which contextual and background knowledge do more than merely *clarify* the application of context-invariant meanings. Rather, meaning emerges from the integration of linguistic and nonlinguistic knowledge, as meaning and background are intimately intertwined.

However, the interdependence of meaning and background presents the language user with a profound challenge. She must determine which background assumptions are relevant at a given time, and which should be ignored. In fact, the challenge becomes quite poignant when we realize that background assumptions differ from context to context, and can even conflict. For example, note the following exchange between an interviewer and a famous Shakespearean actor.

Interviewer: Were Romeo and Juliet sleeping together?
Barrymore: They certainly were in the Chicago company.

Here the interviewer has built a space structured by background assumptions of the play *Romeo and Juliet*. In the play space, Romeo and Juliet are Italian teenagers whose families had a longstanding feud. The interviewer seems to want Barrymore to speculate on whether the

information in the play licenses the inference that the Italians are sleeping together. However, Barrymore has built another space structured by background assumptions of the *production* of the play in which Romeo and Juliet are played by American actors. Although the actor is not usually called Romeo when he's off the stage, the pragmatic function connecting actors to their characters allows us to access the actor via his character's name. Thus partitioning the information into two discrete but related spaces allows us to understand that while the Americans are sleeping together, the Italians need not be.

Understanding which space is being built, and consequently which background frames are relevant, often proves crucial for understanding the overall significance of a linguistic utterance. In the preceding example, Barrymore's remark has changed the interviewer's request for literary speculation into an opportunity for gossip. Interpretation of Barrymore's reply thus requires semantic reanalysis that results in a shift in its implications. This process, called *frame-shifting*, is discussed in detail in Chapters 2–4. In fact, some of the most obvious examples of frame-shifting can be found in jokes. For example, Minsky (1980: 10) notes the knowledge-driven character of humor:

> The element that seems to me most common to all the different kinds of humor is that of unexpected frame-substitution, in which a scene is first described from one viewpoint and then suddenly – typically by a single word – one is made to view all the scene-elements in another, quite different way.

The ubiquity of frame-based meaning construction is thus supported by the suggestion that people play with each other's ability to update their representations adaptively when they tell each other jokes. Moreover, because it highlights the need for dynamic inferencing mechanisms in natural language processing, frame-shifting is seen as a test case for any model of meaning construction. Chapter 2 examines frame-shifting in a number of one-line jokes, addressing what sorts of expectations are set up and how they are revised. Chapters 3 and 4 further develop the constructivist approach, and examine whether it is supported by data from on-line language comprehension tasks.

CHAPTER TWO

Frame-Shifting and Models of Language Processing

As I scan the kitchen, I see a table, chairs, cupboards, a toaster, and a sink piled high with dishes. On the counter is a jar of baby food, half-empty, with a spoon still sticking out of it. While for most of us scanning a room for its contents is a trivial, seemingly effortless, activity, it is in fact an immense computational achievement. The brain is faced with the formidable task of transforming the information in the light that hits our two-dimensional retina into perceptual information that allows us to navigate the three-dimensional world. Moreover, while visual perception requires extracting shape properties and spatial relations, what we perceive are not just shapes, but tables and chairs. Similarly, on the baby food jar we see a drawing of a baby, and not just flecks of ink on paper. Moreover, in spite of the many objective differences between them, we understand the baby in the drawing to be similar in some respects to the real baby sleeping in the other room.

Interpretation always involves the integration of physical information with knowledge at multiple levels of representation. Moreover, the approach the cognitive semanticist adopts toward meaning in natural language suggests commonalities in the interpretation of language and general processes involved in the comprehension of objects, actions, and events in the world around us. Consider again our interpretation of the picture on a jar of baby food. Understanding the label involves knowledge about babies, baby food, and perhaps even the Gerber company. Normally, we interpret the smiling baby as being happy and healthy as a result of eating the baby food. However, the possibility of alternative interpretations exists, and is accentuated by a tale about the Gerber company's move into overseas markets.

Advertising executives claim that Gerber initially used the same packaging overseas that it uses domestically in the United States, with

a smiling baby on the label. However, when Gerber found out about local labeling practices, it decided to change the label. Apparently, in areas where illiteracy was not uncommon, most companies depicted their contents on the label. For most people, hearing about these alternative labeling practices causes a dramatic revision in the understanding of the Gerber label. Knowing that companies often depict their product on the label, we reinterpret the label as suggesting a new meaning for the phrase "baby food."

I call this sort of conceptual revision *frame-shifting*, and suggest it reflects the operation of a semantic reanalysis process that reorganizes existing information into a new frame. Examples like this one violate normal expectations, and thus highlight the tacit role background information plays in the interpretation of all activity. When interpreting objects, actions, and events, we constantly invoke frames from background knowledge to structure expectations and make inferences that go beyond what's immediately present. The sudden shift in the significance of the Gerber baby food label, for example, results in part from its connection to evaluative frames for cannibalism.

Similarly, in language interpretation, linguistic data suggest people access frames from long-term memory to construct meaning for utterances. To comprehend language in the way that people do, it is often necessary to connect actions and events that are not prima facie related. For example, contrast the meaning of "bouncing" in the following passages (adapted from Sanford & Garrod, 1981):

(1a) Jaimie came bouncing down the stairs.
(1b) Paul ran over to kiss her.
(2a) Jaimie came bouncing down the stairs.
(2b) Paul rushed off to get the doctor.

While in (2a) Jaimie tumbles dangerously down the stairs, in (1a) she walks down with excited, springing steps. To understand the difference in these two bounces, it's necessary to invoke background knowledge to connect Paul's actions in the (b) sentences to Jaimie's actions in the (a) sentences. Moreover, when the events involve people, interpretation requires the development of an understanding of actors' goals and motives. The ease with which people seem to make these sorts of connections suggests processes of meaning construction must frequently appeal to this sort of causal and relational information (Rumelhart, 1980; Sanford & Garrod, 1981).

Lange & Dyer (1989) argue that natural language processing requires a capacity for dynamic inferencing based on general knowledge rules represented as *frames*. The general idea of a frame dates back to work by Bateson ([1954] 1972), who argued that communicative acts require interpretive frames that convey the meta-level description of what is going on (e.g., "This is play"). Following in Bateson's footsteps, sociologist Goffman (1974) argues frames can be used to explain the organization of events and our subjective involvement in them. We use frames as representational structures to structure our experience of ongoing activity. Moreover, frames serve as interpretive resources for socially defined activities, giving meaning to things that would otherwise be meaningless. Although individuals constantly interpret their own and each other's activities, they do not singly define situations. Rather, culture provides us with the resources we need to interpret ongoing activity and to define what is real. Goffman thus demonstrates the interdependence of social organization and the use of culturally defined frames to anticipate and formulate expectations about events.

Frames are representations with slot/filler structure, default values, weak constraints on the type of fillers for a given slot, and a hierarchical organization that allows recursive embedding of frames within frames (see Rumelhart & Ortony, 1977, for review). Frames, as representational structures, can be used to represent knowledge about a wide variety of objects, actions, and events. Object frames represent the attributes of objects such as cubes (Minsky, 1975); verb case frames (inspired by Fillmore, 1968) contain abstract slots for agents, patients, and so forth that are bound to the particulars of a situation; temporal frames, or scripts, represent sequences that extend in time (Schank & Abelson, 1977); and situation frames contain both descriptive elements and event sequences (Sanford & Garrod, 1980).

One of the difficulties of natural language processing is that there is often more than one way to interpret events and the various connections between them. Because language (and events in general) can be ambiguous, the choice of which frame to instantiate is often nontrivial. Moreover, as we have seen, additional context can necessitate reinterpretation of the existing message-level representation. Pointing to (3), Lange & Dyer argue that it is consistent both with the goals of washing a cooking pot and hiding marijuana.

(3) John put the pot inside the dishwasher

Lange & Dyer suggest that (3), given alone, will be interpreted as a scenario in which John is attempting to clean a pot he has used for cooking. However, if (3) is followed by (4), new information suggests the need to reanalyze (3) as instantiating the Hiding Pot scenario.

(4) because the police were coming.

Note that this reanalysis requires a complex set of inferences.*

Because frame-shifting highlights the need for dynamic inferencing mechanisms in language processing, it is seen as a test case for models of language comprehension. That is, because the meaning construction process is complicated by the presence of various sorts of lexical and conceptual ambiguities, models must include procedures for deciding between possible interpretations. This chapter examines some possibilities, including the assumption of traditional parsers. Section 1 considers the computational demands of interpreting a simple phrase and suggests shortcomings of a fully compositional account of meaning construction. Section 2 examines frame-shifting in a number of one-line jokes. The focus is on what sorts of expectations are set up in these jokes, and how these expectations are revised. Finally, in section 3, I summarize the relative advantages of a constructivist approach to meaning construction.

2.1 APPROACHES TO COMPREHENSION

As we saw in Chapter 1, compositional accounts of language comprehension are motivated by the productive character of speakers' language behavior. Speakers know a finite number of words and yet are able to combine them to understand and produce a potentially infinite set of

* I1: If the police see John's marijuana, then they will know that he possesses an illegal object (since marijuana is an illegal substance).
I2: If the police know that John is in possession of an illegal object, then they will arrest him, since possessing an illegal object is a crime.
I3: John does not want to get arrested.
I4: John has the goal of stopping the police from seeing his marijuana.
I5: The police coming results in them being in the proximity of John and his marijuana.
I6: The police being in the proximity of John's marijuana enables them to see it.
I7: John's putting the marijuana inside the dishwasher results in the marijuana being inside the dishwasher.
I8: The marijuana is inside an opaque object (the dishwasher).
I9: Since the marijuana is inside an opaque object, the police cannot see it, thus satisfying John's goal.
(Lange & Dyer, 1989: 182).

sentences. In explaining this productivity, the critical issue is how the meaning of each word contributes to the overall meaning of the utterance. Proponents of compositional approaches have considered mechanisms of lexical access and parsing as essentially responsible for this productivity. On such accounts, the processor's task is to assign a syntactic representation (or representations), access semantic information contained in the lexicon, and compose sentence meaning from these two sorts of information.

In contrast, in constructivist accounts, language prompts the language user to construct a frame-based representation of sentence meaning. Language input triggers a search through long-term memory for a frame that will anchor the message-level representation. The slots in the activated frame structure the language user's interpretation and expectations in two ways. First, it allows her to understand the relationship between elements by binding them to slots in the frame. Second, it allows speakers to infer the existence of unstated elements.

Background knowledge and contextual information play an important role in both compositional and constructivist accounts. In compositional accounts, background knowledge and contextual information are crucial for resolving lexical and syntactic ambiguities. Moreover, this information is combined with literal meaning in order to compute utterance meaning. In constructivist accounts, background knowledge and contextual information are the raw material from which the message-level representation is constructed. Thus the former class of models relies on a firm distinction between the computation of literal and nonliteral meaning (e.g., Sperber & Wilson, 1987; van Dijk & Kintsch, 1983), while constructivist models (e.g., Sanford & Garrod, 1981; Schank & Abelson, 1977; Schank, 1983) focus instead on how different sorts of linguistic and nonlinguistic information are combined to form the message-level representation.

This section considers the computational demands of sentential integration in two simple sentences, one which requires frame-shifting, and one which does not. I suggest inadequacies in the compositional account of meaning construction, arguing that it provides an incomplete account of the interpretive process. Because the message-level representation is underspecified by the output of a compositional parser, it requires supplementation by some sort of "postlexical" processes. My argument is that any model of a postlexical process complex enough to bridge the gap between the output of a compositional parser and the resultant message-level interpretation would render compositional processes redundant.

Moreover, we will see that some aspects of message-level representations aren't attainable from operations on the output of a compositional parser.

2.1.1 Composition: A Thoughtful Wife

Consider the following phrase:

(5) A thoughtful wife

On the compositional account, interpretation of (5) would proceed by accessing the meanings of each word from entries in the lexicon, and using syntactic information to combine them into the meaning of the noun phrase "a thoughtful wife" (see, e.g., Fodor & Pylyshyn, 1988). For example, upon hearing "a thoughtful," the processor would start by accessing a set of definitions for "thoughtful" such as the following (taken from an on-line dictionary):

thought-ful adj 1a: absorbed in thought: MEDITATIVE 1b: characterized by careful reasoned thinking 2a: having thoughts: HEEDFUL [became thoughtful about his parents] 2b: given to heedful anticipation of the needs and wants of others: SOLICITOUS

The parser's first problem is to determine which of the four meanings of "thoughtful" is being invoked. Hardly unique to this example, the problem of ambiguity is a general obstacle for natural language parsers, and can be dealt with by anticipating the correct meaning of an ambiguous word or phrase, selecting the correct meaning after the fact, or using a combination of these two techniques. In fact, a large literature in psycholinguistics concerns how good the parser is at anticipating the correct meaning, and how soon after the initial parse the correct meaning has been selected (see Simpson, 1994, for review). The consensus in this literature is that although all four meanings are initially activated, the inappropriate ones are quickly discarded (Swinney, 1979).

On this sort of an account, the lexical processor activates the whole set of "thoughtful" entries in the lexicon, and the parser combines them with syntactic information. The central processor evaluates the resultant interpretations and picks the one that makes the most sense in the current context. Surprisingly, there have been very few explicit proposals about how this latter (clearly important) process works. However, one thing that *has* been established is that in the absence of a constraining context, the processor uses information about the frequency of various

meanings (Simpson, 1994). Default mechanisms are important because the demands of speech comprehension rarely afford the luxury of waiting until there is enough contextual information to disambiguate all terms.

For example, upon hearing "a thoughtful," the processor might start by picking 1a (assuming it is the most frequent), e.g., $(\exists x)\,Meditative(x)$. Upon encountering "wife," the processor could make another frequency-based decision and choose meaning 2 from the lexical entry for "wife":

wife, pl wives 1a dial: WOMAN 1b: a woman acting in a specified capacity – used in combination [fishwife] 2: a married woman

Because "thoughtful" is an adjective and "wife" is a noun, "thoughtful wife" has the semantic force of a conjunction of the two predicates, is-thoughtful and is-a-wife (as in Fodor & Pylyshyn, 1988). Consequently, the parser could build a semantic representation like the following:

$$(\exists x)(Woman(x)\&Married(x)\&Meditative(x))$$

However, one consequence of a frequency-based strategy is that it occasionally results in the need to reassign a lexical entry. In (1), for example, we might reinterpret "thoughtful" once we know that it is to be conjoined with "wife." Although there is nothing contradictory (or even improbable) about the idea of a meditative married woman, intuitively meanings 2a and 2b seem more a propos. Given societal expectations about marriage, our thoughtful wife is more likely to be heedful or solicitous than she is to be meditative. Consequently, the parser would recompute the semantic representation based on the new lexical information for "thoughtful":

$$(\exists x)(Woman(x)\&Married(x)\&Solicitous(x))$$

This model is compositional because regardless of their actual content, the adjective "thoughtful" and the noun "wife" make the same *sort* of contribution to the semantic representation of the sentence. That is, regardless of whether "thoughtful" results in the activation of *Meditative(x)* or *Solicitous(x)*, it contributes a predicate for conjunction. A key assumption of compositional models is that each word has a discrete set of fixed meanings so that a given word makes approximately the same semantic contribution to each constituent in which it occurs (Fodor & Pylyshyn, 1988). However, even in this simple example, our compositional parser

had to appeal to an unspecified noncompositional mechanism to decide between two perfectly congruous interpretations of "a thoughtful wife."

Further, one might argue that, in fact, *Solicitous* is a two-place predicate because one is solicitous *to* someone or something.

$$(\exists x)(\exists y)((Woman(x)\&Married(x))\&Solicitous(x,y))$$

Of course, this raises the question of to whom she is solicitous. The natural answer is her husband. One possibility is to change the representation of "wife" to a relational predicate that entails the existence of a husband. At first blush, this seems to capture a generalization that extends to other relational predicates. For example, if being a mother entails having a child, we might expect a thoughtful mother to be solicitous to her child. Similarly, we might expect a thoughtful teacher to be solicitous to her students.

However, what about a thoughtful criminal? If being a teacher entails the existence of students, being a criminal presumably entails the existence of victims. Do we really want to say a thoughtful criminal anticipates the needs of his victims? Or what about a nonrelational predicate such as *Woman*? A "thoughtful woman" is thoughtful of people in general, but we wouldn't want to say that "woman" entails the existence of other people.

While the idea that "wife" suggests the existence of a husband and that "teacher" suggests the existence of students has some merit, I don't believe we can solve this problem by putting husbands and students into a storehouse of context-invariant meanings. These observations suggest that the recipient of thoughtfulness comes from "wife" in "a thoughtful wife," but from "thoughtful" in "a thoughtful woman." Further, there is a difference in the character of thoughtfulness in "a thoughtful wife," "a thoughtful mother," and a "thoughtful woman" that cannot be adequately explained by a semantic representation that captures the meaning of all three.

In general, a continuum of nouns spans from highly relational descriptions, such as "wife," which suggests the existence of a husband, to the more general "woman." Although the meaning of "thoughtful" is underspecified, it accommodates relational descriptions like "wife" in a nonrandom way such that the second argument of the relational description is a candidate to be the argument of "thoughtful" as well. However, as the "thoughtful criminal" demonstrates, the coordination of information is not strictly compositional. Rather, the meaning of the

resultant noun phrase seems to reflect specific knowledge about wives and criminals.

2.1.2 Semantic Indeterminacy

Moreover, consider (6), a continuation of the phrase in (5):

(6) A thoughtful wife has pork chops ready when her husband comes home from work.

Given the additional context provided by "has pork chops ready," it seems clear that the processor should pick meaning 2b for "thoughtful." At this point in the sentence, the processor might have built a representation that may be glossed as follows:

$$(\forall x)((Wife(x)\&Solicitous(x)) \rightarrow (\exists t)Makes-Ready(x, PorkChops, Timet)))$$

Unfortunately, this representation does not capture all the information that would be available to a competent speaker. For example, what exactly does it mean for pork chops to be ready? Assume that we can eliminate all but the following from the seven listed dictionary entries for "ready":

1b: prepared for immediate use [dinner is ready]

There remains a question of how to specify what it is to be "prepared for immediate use." In this context, the immediate use for the pork chops involves ingesting them; consequently, the pork chops in our example have been cooked. However, if the person preparing the pork chops had been a butcher instead of a thoughtful wife, we would expect the pork chops to be raw. If the person preparing the pork chops had been a supermarket clerk, the state of being "ready" might consist of being packaged in a Styrofoam plate and covered with saran wrap.

The problem the parser has with "ready" is slightly different from the problem it has with "thoughtful." As characterized previously, "thoughtful" is ambiguous; "ready," however, is both ambiguous and semantically indeterminate. While ambiguity forces the listener to decide between a set of meanings, semantic indeterminacy poses a harder problem. Because there is no set of candidate meanings from which to choose, the listener must *create* an interpretation consistent with both the semantics of the term and the contextually available information.

Although ambiguity and semantic indeterminacy are traditionally understood as different problems, the solution to both is to draw on background knowledge to formulate a contextually appropriate representation at the message level.

Our thoughtful wife, here, highlights the extent to which the message-level representation frequently contains information absent or underspecified in the semantic representation in the lexicon. Further, these emergent properties do not reflect the output of the parser, but subsequent noncompositional operations on knowledge associated with each of the lexical items. The choice of which meaning of "thoughtful" should be conjoined with the meaning of "wife," for example, seems to depend on a somewhat dated cultural model of marriage rather than on content in the lexicon itself. Moreover, the contextual modulation of "ready" is not entailed by the definitions of "butcher" or "supermarket clerk." Rather, it is a function of background knowledge about butchers, clerks, and wives.

The interpretation of "ready" in different contexts thus displays systematicity (syncategorematicity), but not truth-functional compositionality. The abstract meaning that the parser generates for "ready" requires substantial refinement before it can contribute to the meaning of the sentence as a whole. Understanding what it is to be "ready" in any particular context depends both on what is being prepared, and what it is being prepared for. This process requires use of background knowledge in a creative manner that suggests a property Clark (1983) has called *contextuality*. Contextuality obtains in cases where arriving at the correct understanding depends crucially on information about the coordination among the speaker, the hearer, and the context of utterance.

2.1.3 Sense and Nonce Sense

Clark (1983) has argued quite persuasively that the existence of semantically indeterminate expressions, especially *nonce senses*, violates a core assumption of the compositional approach and ultimately undermines the utility of the compositional parser. Nonce senses are meanings speakers construct for terms used in a nonconventional way. Interpretation of a given nonce sense requires creating a contextually appropriate meaning. This process often appeals to background information associated with more conventional uses for a term. For example, in (7) the expression "pulled a Sampras" invokes a nonce sense.

(7) Yesterday on the court she pulled a Sampras.

Utterance of (7) can evoke an indefinite number of interpretations depending on the particulars of the context and the various speakers' knowledge of tennis pro Pete Sampras. For example, if it were uttered shortly after the 1995 Australian Open, it might be used to describe losing in the first round of a tournament in which she was a top seed. Uttered after the semifinals of the 1996 U.S. Open, it might imply she got sick to her stomach. Or, it might simply be used to imply she has a very powerful serve.

Such instances of nonce sense clearly violate the requirement of compositional parsers that a term make approximately the same semantic contribution in every constituent in which it appears. Further, the contextuality of nonce sense suggests that the number of different interpretations for "pulling a Sampras" is limited only by our ability to imagine parallels between a given scenario and salient (for whatever reason) aspects of Pete Sampras' existence. The indefinite number of meanings that can be given for any contextual expression poses a problem for the *sense selection assumption*. This is the assumption that each constituent has a finite number of discrete meanings, and that the comprehension task is to select a meaning from this set (Clark, 1983). However, the parser cannot possibly select between the meanings for a semantically indeterminate expression since they are indefinite in number.

Clark argues nonce senses pose two problems for parsers that incorporate the sense selection assumption: the *nonparsing* problem and the *misparsing* problem. Nonparsing occurs when the intended meaning cannot be computed from the senses stored in the lexicon. To illustrate, Clark uses a quote from an Erma Bombeck column in which the columnist describes her daughter's search for a roommate.* Citing (8), Clark suggests a typical parser would deem this utterance "uninterpretable nonsense rather than interpretable nonce sense" (Clark, 1983: 299).

(8) Our electric typewriter got married and split.

* We thought we were onto a steam iron yesterday, but we were too late. Steam irons never have any trouble finding roommates. She could pick her own pad and not even have to share a bathroom. Stereos are a dime a dozen. Everyone's got their own systems. We've just had a streak of bad luck. First, our Mr. Coffee flunked out of school and went back home. When we replaced her, our electric typewriter got married and split, and we got stuck with a girl who said she was getting a leather coat, but she just said that to get the room.

– Erma Bombeck, quoted in Clark, 1983: 298.

Rather than interpreting the typewriter in (8) metonymically, the typical parser would balk at the implication that an inanimate object might get married. The nonparsing problem thus arises when the lexicon does not include the contextually relevant meaning, and the compositional meaning is semantically anomalous.

In contrast, the misparsing problem arises when the compositional meaning is congruous, but not the one intended by the speaker. Going back to Bombeck's column, Clark points to (9), noting the columnist's intended meaning of roommates who own stereos.

(9) Stereos are a dime a dozen.

The problem is that "stereos" has an existing lexical entry consistent with the standard sound system interpretation. Because (9) is perfectly congruous within the sound system interpretation, the parser will have no reason to compute the intended nonce sense.

2.1.4 Frame-Shifting: Another Thoughtful Wife

Another dimension of the contextuality of language interpretations can be seen in sentences that trigger frame-shifting. For example, contrast the meaning of (6) with that of (10):

(10) A thoughtful wife has pork chops ready when her husband comes home from fishing.

The word "fishing" here cues the reader to make dramatic changes in her representation of the scenario. The motives of the wife in (10) differ markedly from those in (6), as does the status of the stereotypical husband in (6) compared to the unsuccessful fisherman in (10). To form the intended (joke) interpretation of (10), the reader needs to reinterpret the noun phrase "a thoughtful wife" as one who anticipates her husband will come home without any fish. However, this poses a number of problems for a compositional parser.

First, it's unclear what would prompt reanalysis in a compositional model. Recall that on a compositional approach, the parser invokes background knowledge primarily to resolve ambiguities and to bridge informational gaps. This suggests that any differences in the interpretation of (6) and (10) must either be due to a consequence of the lexical semantics of "work" versus "fishing," or pragmatic processing prompted by deficiencies in the output of the parser. However, (6) and (10) do not noticeably differ in the amount of lexical ambiguity that they present to the

parser. If anything, my on-line dictionary suggests (10) is less ambiguous, listing only two definitions for "fishing," compared to twenty-three for the noun "work."

Moreover, a compositional interpretation of (10) presumably does not present the reader with any more informational gaps than does (6). Consequently, a parser interpreting (10) will fall prey to Clark's misparsing problem. Although it can trivially derive an interpretation for (10), it will not arrive at the desired inference that the husband does not catch any fish. Because the compositional interpretation of (10) is perfectly congruous, there is no need to appeal to background knowledge necessary to construct the intended joke interpretation.

Further, it is quite possible that the preceding lexicon cannot handle the reader's new interpretation. Although our thoughtful wife is still "given to heedful anticipation of the needs and wants of others," we might not consider her to be solicitous, as listed in definition 2b. In any case, the particular understanding at the message level will be far richer than that available from whatever lexical entry we can decide on. As argued in Chapter 1, it is this rich, specific interpretation that should concern the cognitive scientist; however, it is this rich, specific interpretation that is left unexplained by the compositional parser.

2.1.5 Constructivist Accounts

The question then arises as to what sort of a model *can* explain the specific interpretations readers seem to be able to construct. Let's begin by considering the shortcomings of the compositional approach. First, the language user seems to know more about the scenario in (10) than does the compositional parser. The need to reorganize is not prompted by the mechanisms of the compositional parser, but rather by expectations about how husbands and wives interact. Further, the nature of the resultant interpretation also suggests the use of structure that does not originate in a representation such as that built by a compositional parser. For instance, understanding the thoughtful wife's action as an insult to her husband's fishing ability requires the reader to infer both that the wife has predicted her husband will not catch any fish, and that he *should* be expected to catch some fish. Thus the difference in the meanings of (6) and (10) is less a consequence of the meaning of "fishing" than the way in which background knowledge about fishing interacts with the meaning evoked by the rest of the sentence.

A constructivist account easily explains the "extra" material that readers derive from such sentences. Because background information is

crucial to the construction of sentence meaning, readers form the message-level representation by accessing frames from long-term memory to aid the construction of cognitive models in a mental space lattice. Fillmore's (1982) observation that words have meaning relative to a set of background assumptions reflects the fact that a word serves to activate knowledge structures pertinent to the current context. Moreover, Langacker (1987) suggests morphosyntactic marking, word order, and other features of language that have traditionally been considered purely structural can be understood as *profiling*, or attentionally focusing, various aspects of evoked knowledge structures, as well as prompting the retrieval of abstract frames to help coordinate newly profiled structure with established structure.

On the constructivist account of (10), the reader begins by activating the Good Wife frame to guide his construction of the meaning. However, at "fishing," the Good Wife frame is no longer adequate to represent the current scenario, and the existing representation must be reorganized into a new frame. Frame-shifting in (10) involves activating the fishing event frame (getting up early, finding bait, finding a spot, catching fish, going home, cleaning fish, cooking fish) and integrating it with existing structure. Moreover, the processor seems to be sensitive enough to the buildup of semantic context to fill the catching-fish slot with "no fish," and the going-home slot with "empty-handed."

Semantic indeterminacy is not a problem for a constructivist account, as words do not retrieve lexical entries but rather activate abstract structures and processes for integration with contextually available information. Obviously the processor relies upon regularity and systematicity in the sorts of structures evoked by a given word. However, there is no requirement that a context-invariant meaning be evoked at any stage in the meaning construction process. In the absence of explicit context, speakers activate generic frames filled with default values. But because language use is usually situated in a shared physical, social, and informational context, the frames that speakers construct in normal discourse can diverge substantially from the prototype.

Consequently, the semantic indeterminacy of "ready" occurs because it tends to evoke an abstract frame whose specific content is dependent on how slots for User and Use are filled. Moreover, constructing a meaning for "ready" in a given context will require the speaker to employ variable degrees of creativity in the meaning construction process. In a familiar scenario such as preparing dinner, a cooking frame should be readily available. However, in an unfamiliar scenario, such

as determining how "ready" might pertain to a member of an alien culture preparing for a ritual, the meaning construction task requires more imagination. Consequently, meaning construction is best construed as occurring on a continuum that has the retrieval of standard frames and default values at one end, and the creation of nonce senses at the other.

Nonce senses, then, reflect the more creative end of the meaning construction process. An expression such as "pulled a Sampras" in (7) directs the reader's attention to knowledge of Pete Sampras and prompts the creation of an abstract frame that can be mapped (analogically) onto contextually evoked information about a woman on an unspecified tennis court. Frame-shifting, too, reflects this more creative aspect of meaning construction. In (10), for example, the term "fishing" invites the reader to create a new frame that relates a wife cooking pork chops to her husband's return from a fishing trip.

Besides providing a way of including the processes that subserve the full spectrum of word and expression meanings, constructivist accounts can also explain more compositional-looking phenomena. Traditional compositional behavior results from typical patterns of conceptualization being imposed upon prototypical representations retrieved from memory. Because frames with default values are primarily used in cases where no contextual information is available, interpretation of novel items in a null context will often look compositional.

The data that have traditionally supported the idea of compositional meaning construction are sentences understood with little or no supporting contextual information. However, this is exactly the situation in which the language user is most likely to recruit frames whose slots are filled with default values. Moreover, in this sort of situation, the context is barren enough to promote typical patterns of conceptualization. Consequently, the traditional treatment of composition as a paradigm for meaning construction is probably an artifact of the context in which linguists obtain their data.

One might object that the old compositional approach can be saved by frames with slots filled by default values. That is, perhaps lexical content as traditionally construed is equivalent to frames with defaults, and there are other postlexical processes that produce our continuum from sense to nonce sense. However, to attempt to save compositionality is to forfeit the parsimony of the constructivist approach. The constructivist's parsimony results from positing a single constructive process to explain cases that look compositional as well as those that do not. Moreover, regardless of whether we think of meaning construction as

fundamentally constructive, noncompositional phenomena require some sort of explanation.

Langacker's (1987) theory of cognitive grammar is a good example of a constructivist account that can encompass both compositional and noncompositional phenomena. As noted earlier, on this theory grammar is understood as systematic pairings between phonological and semantic structure. However, rather than positing an independent level of syntactic combination and associated rules of semantic interpretation, Langacker suggests that the lexicon, morphology, and syntax form a continuum in which traditional syntactic structures are the most abstract. Regularities in form-meaning pairings arise because of abstract frames that specify how two input frames can be integrated into a composite frame. However, even these abstract frames are construed as fundamentally semantic.

In cognitive grammar, compositional patterns arise from abstract frames that specify how two input frames can be integrated into a composite frame. When these abstract frames (which Langacker, 1987, calls "schematic constructions") specify the contribution that each input frame makes to the resultant meaning, the resulting integration appears (and, indeed, is) compositional. For example, one can combine default frames for "black" and "bird" to yield a composite frame for a bird with a black exterior. The composite structure "black bird" is computable from the information in default frames for "black" and "bird" and the structure that results from coordinating them.* Langacker (1987) requires that coordinated structures fulfill a criterion of *analyzability* such that composite expressions can derive from representational structure built from component expressions, contextual information, and background knowledge. An analyzable expression includes (i), a salient frame that represents the relationship between each input frame and the composite frame and (ii), a salient frame that represents the relationship between the two input frames. Analyzable expressions are not built from their component expressions. Rather, they are *motivated* by these expressions. Consequently, composite structures can contain information that is different or more specific than that available from either component structure. The criterion of analyzability thus subsumes compositionality as a special, simple case.

* See Travis, 1981, for discussion of how the inclusion of nonprototypical information affects the meaning of similar constructions such as "black kettle."

2.2 FRAME-SHIFTING IN ONE-LINE JOKES

The flexibility of the interpretive process is perhaps nowhere more evident than in jokes. Deliberately constructed to violate the listener's expectations, jokes highlight the various sorts of expectancies we make as a matter of course. Consequently, this section explores frame-based inference by analyzing the demands of meaning construction in a number of simple jokes. Analyses demonstrate the extent to which comprehension relies extensively on knowledge of typical scenarios. However, these jokes suggest people are quite capable of going beyond knowledge of the typical to construct creative nonconventional meanings.

2.2.1 Typical Situations

(11) By the time Mary had had her fourteenth child, she'd finally run out of names to call her husband.

In interpreting (11), the reader invokes knowledge of the events that surround childbirth. Because parents typically name their children soon after they are born, the reader initially interprets "names" as baby names. This initial interpretation of the scenario relies on extracting a frame based on some combination of information provided by the text of the joke and general world knowledge retrieved from long-term memory. Similarly, the connection between Mary having had her fourteenth child and running out of baby names is supported by knowledge that fourteen is an atypically large number of children.

However, the term "husband" is not consistent with the lexical-level expectations, or the message-level expectations that support them. This prompts the lexical reinterpretation of "names" as derogatory epithets directed toward the man Mary blames for her fourteen pregnancies, and triggers pragmatic reanalysis we have called frame-shifting. The initial interpretation related Mary's having fourteen children to naming fourteen children, and hence running out of baby names. In contrast, the shifted interpretation relates Mary's having fourteen children to being pregnant and giving birth fourteen times.

Although (11) functions as an example of how people utilize their knowledge of typical situations in their interpretation of utterances, in the final analysis, the scenario described is not at all typical. For one thing, it is (at least in the United States) quite atypical to have fourteen children. Second, swearing at the husband is presumably not part of

the scripts for pregnancy and/or childbirth. Furthermore, Mary's action suggests a negative attitude toward childbearing that (however warranted it may be) is not part of the scripted response to these events.

What (11) in fact suggests is that in the absence of information to the contrary, readers assume that typical assumptions obtain. However, *given* information to the contrary, readers display a remarkable ability to adapt the particular models that they have constructed to incorporate the new information. In (11), this involves understanding the implications of having fourteen children as opposed to the more typical 1.4. Exploration of the implications of an atypical situation is not confined to jokes: In (11), both the invited and the joke interpretations concern the implication of having an atypical number of children. But while the invited reading concerns its implication for naming the children, the joke reading concerns its implication for the mother's quality of life.

Interestingly, the shifted interpretation also relies on knowledge of the events surrounding childbirth – just slightly different knowledge than that required to construct the first interpretation. Besides the knowledge that fourteen is an atypically large number of children, the frame-shifted interpretation requires knowledge that giving birth is painful and that husbands play a causal role in their wives' pregnancies. The reanalysis of (11) thus demonstrates the way that context can affect the interpretation of a word by calling to mind particular aspects of its conceptual structure.

Similarly, Johnson-Laird (1993) contrasts the meaning of "tomato" in the following, where (12) highlights a tomato's characteristic shape, (13) its characteristic color, and (14) its propensity to be squashed.

(12) The tomato rolled across the floor.
(13) The sun was a ripe tomato.
(14) He accidentally sat on a tomato.

Barclay, Bransford, Franks, McCarrell, & Nitsch (1974) show that interpretation of unambiguous words varies with sentential contexts by rendering some aspects of its meaning more salient than others. Recall cues that mentioned relevant properties of events described in the study sentences were more effective than those that mentioned irrelevant properties. For example, "heavy" was a more effective cue for recalling a story about moving a piano than it was for recalling one about tuning a piano.

Further, Tabossi & Johnson-Laird (1980) have shown that questions such as (15) are answered fastest when they are preceded by a sentence

such as (16), which highlights the relevant aspect of a diamond. Thus subjects who read (16) took less time to answer the question than subjects who read (17). Moreover, subjects who read (17), a context that is relatively neutral with respect to a diamond's brilliance, took less time to answer the question than subjects who read sentences like (18), which highlight irrelevant aspects such as the diamond's hardness.

(15) Is a diamond brilliant?
(16) The mirror dispersed the light from the diamond.
(17) The film showed the person with the diamond.
(18) The goldsmith cut the glass with the diamond.

Further, Tabossi (1983) has shown that this phenomenon is not strictly lexical, by employing different words in the questions than in the sentences that precede them. For example, Tabossi reports that sentences like (19), which highlight a particular aspect of an object, promote faster responses to questions about the *relevant* dimension than to questions about an *irrelevant* dimension.

(19) The fountain pen left a spot in the desk drawer.
(20) Does ink leave a stain?

As in example (11), the phenomenon is not tied to activation caused by lexical items but relates to the construction of the overall message-level representation.

The idea that we create new frames in the course of meaning construction explains the partial nature of the models we construct. Because we're always creating frames for particular specific contexts, we don't *need* all the information that is potentially available in long-term memory. Rather, we construct models that contain the specific information relevant to the current context. Moreover, given all that we know about the many objects we encounter in the world, it makes sense that we do not retrieve the totality of this information.

Glenberg, Kruley, & Langston (1994), following a line of reasoning first articulated by Sanford & Garrod (1981), have proposed that limitations on working memory necessitate the partial nature of cognitive models. Although elements set up in models include only a few salient features, they include pointers to information in long-term memory that constrain reasoning about the entity in the mental model (Glenberg et al., 1994). For example, a cognitive model that includes an airplane would not include the entire airplane frame, but might include the fuselage

and the wings. However, information in long-term memory constrains elaboration of the model, including things such as how the plane moves.

Nonetheless, the issue of how particular partial structures are recruited to construct meaning and, further, how the reanalysis in frame-shifting serves to recruit different structures remains open. Along these lines, Langacker (1996) has made some suggestions about how abstract frames evoked by grammatical information can help the language user assemble the message-level representation. More specifically, Langacker shows how a number of different grammatical phenomena can be described as directing meaning construction by shifting the hearer's focus of attention within some context. Each time the hearer shifts focus, it creates a new context and makes new entities accessible as potential foci. The concept of attentional focus is known as *profiling* and plays a fundamental role in Langacker's (1987) theory of cognitive grammar. Langacker defines grammatical class in terms of the types of entities that an expression can profile. For example, nouns profile *things*, while verbs profile *processes* (see Langacker, 1987, for technical definitions of *thing* and *process*).

Langacker (1996) has subsumed a number of different linguistic phenomena under the rubric of reference-point relationships. Reference-point relationships invoke a cognitive ability to access a concept of one entity (the target entity) from another (the reference point). In possessives, for example, the possessor is the reference point and the possessed is the target. One feature of this account is that it explains the asymmetry of these constructions as a function of the asymmetry in the cognitive salience of the reference point and the target. Langacker contends human entities are more salient than nonhuman, wholes are more salient than parts, concrete entities are more salient than abstract ones, and visible entities are more salient than nonvisible ones. Because salient entities function as better reference points, human possessors, as in (21), are more acceptable than nonhuman possessors, as in (22); wholes, as in (23), are more acceptable than parts, as in (24).

(21) the man's wallet
(22) * the wallet's man
(23) the dog's tail
(24) * the tail's dog

Langacker (1987) suggests that the possessive construction prompts *constructional assembly*, a procedure for accessing information needed to comprehend phrases like (21) and (23). Constructional assembly begins

by profiling the reference point, thus making concepts in its dominion potentially available. Next the target is profiled. When the target is in the dominion of the reference point, its activation is facilitated by the prior activation of the reference point, and the utterance is grammatical. Likewise, when the reference point is less salient than the target, the utterance is ungrammatical, as in (22) and (24). In sum, the possessive construction functions as an abstract frame for assembling an unspecified relationship between two entities that differ in cognitive salience.

Information from the possessive construction helps constrain the meaning that the speaker will ultimately construct, but does not do so deterministically. Although the default frame recruited by the possessive construction is one of ownership, specific information about the reference-point and target entities can suggest alternative relational frames. Langacker (1996) points to examples such as "our bus," "your candidate," and "Lincoln's assassination" as demonstrating the flexibility of this construction. Similarly, Clark (1983) has noted the contextuality of seemingly straightforward instances of the possessive construction. For example, "John's dog" might be used to describe the dog that John owns, the dog that John is currently standing by, the dog that John fed once, the dog that chases John home from school each day, and so on.

Abstract frames suggested by morphemes such as *'s* manipulate the listener's attention in a way that helps constrain the construction of a meaningful interpretation. Moreover, as noted previously, the idea of a reference-point construction can be applied to a number of other linguistic phenomena to explain the way in which information such as word order can be exploited to direct the listener's attention from entity to entity in the discourse in a way that constrains the construction of the overall meaning representation. Besides the possessive construction, Langacker (1996) offers an account of nested locative constructions such as (25).

(25) Your camera is in the study, in the closet, on the top shelf, beside the pillows.

In this construction, the speaker traces a mental path in which each prepositional phrase serves as the context for the interpretation of the next. The entity being located is profiled with respect to a series of spatial contexts of increasing specificity. These spatial contexts serve as reference points from which the speaker can access more specific information about the target (the camera) and its relationship to other entities in the discourse.

While these suggestions are a far cry from the solution to the issue raised, Langacker's observations suggest how different patterns of profiling enable the language user to coordinate abstract frames with more specific frames constructed in response to particular words, perceptual information, and existing representations at the message level. In interpreting (11), for example, the construction, "By the time X, Y," invites the listener to create a frame that imputes shared causality in the occurrence of state X and state Y.* For example, in (26), eating fourteen donuts is related to feeling full afterward. Moreover, (27) is anomalous because it is hard to construe eating fourteen donuts as being related to feeling hungry.

(26) By the time he ate fourteen donuts, he was pretty full.
(27) # By the time he ate fourteen donuts, he was hungry.

In (11), the relationship between X (Mary having her fourteenth child) and Y (running out of names) initially involves the Naming slot in the childbirth script. However, "husband" suggests that the instantiation of the Naming slot in the childbirth script is erroneous. Presumably because "names" has been profiled as the head of the noun phrase modified by "to call her husband," the incongruous item first prompts lexical reanalysis of "names," which in turn leads the speaker to a new frame to relate X and Y. Information at various levels of abstraction helps to constrain the initial analysis and subsequent reanalysis.

In particular, interpretation of (11) involves recruiting and coordinating frames at various levels of abstraction, including the schematic relationship set up in the "By the time X, Y" construction, knowledge of the events surrounding childbirth, and novel frames for the relationship between Mary, her husband, and her children. Analysis of this example demonstrates how context can affect lexical interpretation by calling up particular facets of background knowledge about the domain in question. Moreover, while meaning construction clearly involves the recruitment of knowledge about typical situations, the message-level representation is constructed dynamically in response to contextual demands.

* I believe this account can be adapted to encompass the more complicated case in sentences such as, "By the time the United States intervened in Zaire, the damage was done," where it's not the intervention per se that causes the difficulty but the fact that X didn't occur sooner. The adaptation involves the construction of a counterfactual scenario in which some aspect of X does not occur and invites the inference that Y does not occur. So in the Zaire example, we set up a counterfactual scenario in which X occurs earlier and invite the inference that Y would have occurred.

2.2.2 Default Values

On many models, lexical and conceptual ambiguity prompts high-level semantic reinterpretation. However, frame-shifting in (28) is prompted neither by lexical nor by high-level conceptual ambiguity.

(28) Everyone had so much fun diving from the tree into the swimming pool we decided to put in a little water.

That is, we would not want to say that a scenario in which people have fun diving into a swimming pool is ambiguous between cases in which the pool *does* have water and cases in which it does *not*. Rather, the possibility that the pool is empty only arises because the speaker explicitly evokes it. An intuitive characterization of the aspects of the message-level representation that change as a result of reading the word "water" highlights the initial features of that representation, and suggests the need for processes that could both create and revise expectations appropriately.

In this case, the meaning of "water" in (28) implies that the Swimming Pool frame here has a slot *Contains*(x) that has been filled by a default value, Water. Moreover, the reanalysis prompted by "water" is not merely the reassignment of a filler to a slot. Information that the reader has already established about the scenario is also reanalyzed. The reader constructing a model of people jumping into an empty swimming pool will presumably fill slots for Cause and Result differently from one who is constructing a model of people jumping into a swimming pool full of water. The frame-shifted interpretation of (28) thus suggests an interdependence between the way slots are filled in a newly constructed frame.

Interpretation of (28) relies on knowledge of the typical backyard swimming pool. The first clause evokes a model of people having fun diving from a tree into a backyard swimming pool. Moreover, the second clause is initially interpreted as installing a piece of equipment commonly found near backyard swimming pools that might function analogously to the tree. Note how this interpretation relies on the reader having drawn a covert analogy between a nonstandard addition to the pool (the tree) and a possible manufactured addition that would serve the same function. However, the disjunctor "water" prompts the reader to revise the default assumption of the Backyard Swimming frame, that there was water in the pool. Revising this simple assumption has substantial implications for the consequences of diving from the tree into the pool, and for the mindset of those who enjoy such activities. The

last word of (28) ("water") is surprising, both because it violates the constraints on an unfilled slot in a frame that the listener has set up to represent the scenario, and because it contradicts the default assignment of *Water* to the x in the slot for *Contains*(*x*).

Of course, this latter contention is hardly controversial. All models of language comprehension acknowledge the importance of background knowledge at some stage in the process. However, in compositional models, the recruitment and coordination of background knowledge is a qualitatively different process from the mechanical assembly of the semantic representation. Moreover, in such models, discourse-level pragmatic processing relies to a certain extent on the output of the semantic processor. So, where compositional models appeal to background knowledge to complete the underspecified output of a compositional parser, constructivist models view the manipulation of background knowledge as an essential component of processing.

In the constructivist approach sketched here, the primary function of linguistic information is to facilitate the construction of idiosyncratic utterance meanings. The language of an utterance provides speakers with cues to help them recruit and coordinate background knowledge to construct idiosyncratic meanings that suit the context. Thus the need for frame-shifting in (28) was not prompted because the utterance describes an ambiguous scenario, or even (primarily) because it violates the default assumption that the pool contains water. The word "water" in (28) is inconsistent with the specific model the speaker is likely to construct, one in which the home owner decides to install some pool equipment, and, as noted previously, a number of models that the speaker has instantiated to get to this point in the story.

This is quite a leap, but competent speakers are able to adapt their existing representations via the process of frame-shifting. Semantic reanalysis in this example highlights the interdependence between the assignment of fillers to slots consistent with the predictions of models that construe comprehension as a process of constraint satisfaction. This capacity for rapid reanalysis is also consistent with the suggestion that language provides speakers with clues that help them to construct frames that can be revised in light of incoming information.

2.2.3 Reinterpretation

But what triggers the frame-shifting process? Presumably upon encountering the disjunctor, the reader or listener will search working memory

for something that can be reinterpreted. For example, in joke (29), the reader begins by evoking a barroom frame to structure the initial space.

(29) When I asked the bartender for something cold and full of rum, he recommended his wife.

Because customers often ask bartenders to recommend drinks (among other things), and because being cold and rum-filled are characteristics of drinks found in bars (though not of drinks in general), the object of the predicate Recommended is a slot for a drink. However, when the space is expanded in accordance with the second clause of the joke, the object of Recommended is a slot for the Bartender's Wife. Because this entity has been profiled, the structure that surrounds its counterparts in other spaces provides crucial input to the frame-shifting process.

The information provided in the phrase "something cold and full of rum" is reinterpreted so as to reconstrue the bartender's wife as a frigid alcoholic. This reconstrual is a nontrivial endeavor involving the creation of a scenario in which properties of being cold and full of rum might plausibly apply to the bartender's wife. The need for scenario creation is evidenced by the correspondence between "cold" and the resultant interpretation of frigid. There are a number of ways in which "cold" might be applied to a person: She might be experiencing the subjective sensation of cold, her body might have been frozen, she might be unfeeling or, vicious. Interestingly, the latter two possibilities are suggested by the following permutation of the joke in which the bartender is a woman and her spouse is a man.

(30) When I asked the bartender for something cold and full of rum, she recommended her husband.

This difference between the implications of being a cold husband versus being a cold wife reflects the way in which background knowledge is constantly recruited and exploited to construe particular scenarios.

In (29), the mapping suggested by coreference of "he" with "the bartender" requires little if any revision of established structure. However, in (30), the mapping suggested by coreference of "she" with "the bartender" may prompt revision of the default assumption about the bartender's gender. Further, as noted previously, the coreference in (29) of "something cold and full of rum" with "his wife" requires revision of the assumption that the object slot in the Recommendation frame will be filled by an alcoholic beverage. While the "size" of the leap differs

between these three cases, the creation of structure to plausibly connect anaphoric referents is qualitatively similar. Moreover, because the properties "cold" and "rum-filled" are more favorably applied to an alcoholic beverage than a woman, the bartender's speech act can be reconstrued as an insult. This helps to restore global coherence because the bartender's reply to the customer's question is no longer a drink recommendation.

Frame-shifting seems to be prompted by a violation of slot-filling constraints. However, rather than failing to interpret these sorts of examples, speakers respond by creating a new frame in which slot-filling can proceed. Further, the observations in this section suggest that the dramatic frame-shifting prompted by the punchline of the joke in (29) is a pronounced instance of a more general semantic reanalysis process necessitated by the fact that the meaning of utterances is not computed from linguistic representations, but motivated by them. Because language comprehension is a constructive process in which speakers have to activate and coordinate background knowledge selectively, meaning construction processes responsible for slot-filling can trigger the construction of a new frame.

2.2.4 Goal Presupposition

Discussion thus far has suggested that people construct cognitive models of the relevant scenarios that are at least loosely constrained by information in long-term memory. However, this is at best an impoverished version of what people do in the course of natural language processing. Language use is an integral part of social interaction and plays an organizing role in social relationships. Besides drawing on information about what sorts of entities are found in typical scenarios, people seem to use information about the goal of particular actions in their comprehension of sentences.

For example, joke (31) turns on the sort of assumptions people are likely to make about the point of the utterance.

(31) Arguments between couples are healthy; sometimes they even prevent marriage.

In this case, the invited interpretation is that couples want to stay together and that the point of the utterance is to provide guidance as to *how* couples can achieve this goal. Interestingly, the first clause in (31)

actually violates the default assumption that arguments are bad. To accommodate this information, the reader is invited to construct a frame where arguments fulfill a useful role with respect to the assumed goal of the couple. One possibility is that arguments defuse aggression that left unchecked might lead to a breakup or some other nondesirable effect with respect to the couple's goal.

However, "marriage" is completely inconsistent with the Defused Aggression frame, and so prompts frame-shifting. In fact, by setting up avoidance of marriage as a goal for couples, the shift involves a move back to the default interpretation of arguments, but away from a default interpretation of marriage. This example shows how default expectations are easily modulated and remodulated, and how knowledge about goals affects the meaning construction process.

Joke (32) turns on the sorts of assumptions people are likely to make about the topic of conversation.

(32) Now that the Hollywood couple has ironed out the divorce settlement, they can finally go ahead with their wedding.

Such examples point to a gap in current models of language comprehension in the need to understand goals at both topic and speech act levels. In fact, understanding the way speakers infer each other's goals may provide crucial information for models of meaning construction that can handle the creative contextuality of sentences that require frame-shifting. This suggestion is consistent with Barsalou's (1987) argument that concepts are temporary constructs in working memory that vary widely as a function of context, recent experience, and goals. Moreover, Boden (1988) notes the positive impact that computational notions of planning (such as goals, subgoals, and preconditions) could have for models that can understand utterances in context.

Clark (1983) suggests a relationship between the way we understand contextual expressions and the way indirect speech acts are understood. For example, (33) can serve as a reminder to bring an umbrella or wear a coat, or as a request to close the window.

(33) It's raining out.

Clark's goal is to develop a way of parsing indirect speech acts that will also work on other contextual expressions.

Clark (1979) argues that the interpretation of utterances involves the use of a hierarchy of goals the speaker is trying to attain. For example, if a hostess at a restaurant is asked, "Do you take American Express?" she

will infer the speaker wants to use his AmEx card and will answer with a simple "Yes" or "No." In contrast, if the hostess is asked "Do you take credit cards?" she will infer the speaker wants to use a credit card, and wants to know which ones are acceptable. Clark (1979) confirmed these intuitions empirically with a systematic telephone survey of restaurateurs. Hostesses were more likely to reply to the more general question with a list of acceptable credit cards, and the nature of the hostesses' reply often contained other clues that they had inferred the questioner was interested in dining that evening and paying with a credit card. For example, one reply to, "Do you take credit cards?" was, "Yes, we accept credit cards, but tonight we're closed."

Goal hierarchies are inferred by appeal to linguistic and nonlinguistic information alike. Clark (1983) notes a many-to-one mapping between the speaker's goal hierarchy and the constituents of a sentence. His notion of intentional parsing involves the idea that the point of parsing is to identify the goal hierarchy that motivates the speaker's utterances. Because of the indeterminate relationship between the sentence meaning and the speaker intentions, the traditional notion of parsing is necessarily inadequate for the task of inferring goal hierarchies. Clark (1983) suggests that parsers incorporate common ground that exists between speakers, shared background and contextual information, as well as linguistic properties of the utterance.

The idea of an intentional parser thus goes beyond the idea of sense selection by having the parser create meanings using the common ground and the goal hierarchy. This involves considering the apprehension of speakers' intentions to be a fundamental component of parsing. These ideas are supported by Gibbs' (1994) finding that speakers arrive directly at the intended interpretation of indirect speech acts rather than computing them from the output of a compositional parser. That is, listeners' initial understanding of, "Can you pass the salt?" is as a request rather than as a question about manual dexterity (see Gibbs, 1994, for review). Just as utterances prompt the construction of cognitive models that represent indicative content, they simultaneously prompt the construction of speech act frames. Consequently, I will occasionally employ additional mental spaces structured by speech act and other interactional frames whose elements are linked to counterparts in spaces that represent the more "literal" structure.

The nature of the frame-shifting necessary to understand joke (34) reinforces the idea that utterances are made with interactional goals in mind. Initially, the reader evokes structure in the Told space, the space

that represents what the speaker told the woman. Telling her that he "wasn't himself," he suggests he is upset and seeking her sympathy. Moreover, initial cues in her reply to the man presumably prompt the reader to complete the sympathy-seeking scenario with an offer of sympathy from the woman. However, at the disjunctor "improvement," the reader must create a scenario in which the presuppositions of improvement can coexist with the utterance "wasn't himself."

(34) He told her he wasn't himself and she said she'd noticed the improvement.

Joke (34) highlights the default assumption underlying the conventional meaning of "I'm not myself today." Namely, it's better to be in your normal state than an abnormal one. Joke (35) suggests that while improvement can be compatible with not being oneself, it seems to involve a scenario in which the person in question has been unwell.

(35) He told her he still wasn't himself, but she said she'd noticed (some) improvement.

In (34), the disjunctor "improvement" prompts a reinterpretation of "he wasn't himself." Because his abnormal state has been construed by the woman as an improvement over his normal state, it suggests his normal state is quite offensive. Consequently, the woman's remarks are no longer congruent with the contextually evoked Sympathy frame, but are best construed as an insult. The reinterpretation concerns the default evaluation of normalcy, and it has implications for the significance of the woman's remark in the social interaction between the two individuals mentioned in the sentence.

In conversation, utterances are made with interactive goals in mind. So, in their most natural context, language utterances function in an interactive context where speakers have a level of awareness of each other's goals. In keeping with the earlier proposal that language motivates structure but does not determine it, I have suggested that the same sort of flexible constructive retrieval mechanism posited to explain the contextuality of utterances (different aspects of conceptual structure recruited on different occasions) might be useful for explaining how utterance meaning varies with interactional goals. While pragmatic theorists have historically considered locutionary and illocutionary force (that is, what is *meant* versus what is *meant to result* from an utterance) as separate aspects of meaning, the same sorts of cognitive processes may well subserve them both.

2.2.5 Nonconventional Meanings

Another example of linguistic phenomena often overlooked in compositional accounts of meaning construction is people's ability to evoke and exploit nonstandard meanings for words, phrases, and other linguistic units. For example, to understand (36), the reader needs to reinterpret the conventional meaning of "redecorated" so as to accept "drunks" as an argument.

(36) They redecorated the bar with new drunks.

The conventional meaning of "redecorate" involves changing the appearance of an establishment or a dwelling by replacing furnishings. The initial use of this verb evokes an Interior Decoration frame that serves to constrain the unfilled slot to the type of decorations typically found in a bar. However, "drunks" does not meet the constraints specified by the Interior Decoration frame and thus triggers frame-shifting.

The frame-shifting process involves mapping the existing message-level representation into a new frame. The shifted frame contains some of the structure of the initial frame: Old aspects of the bar have been replaced with new ones. Because decorations are almost ubiquitously inanimate, patrons of a bar are not customarily considered decorations. However, stereotypical knowledge of the behavior of drunks in bars affords a mapping between drunks and decorations. Drunks who frequent a particular bar tend to have a regular seat where they sit when they drink. As drinking is (by definition) a favorite activity of drunks, a drunk can spend a considerable amount of time sitting in the same seat, thereby becoming a familiar figure in the landscape, much like a painting or a jukebox.

Linguistic cues suggest the need to abstract some schematic similarity between drunks and bar decorations. The listener responds by creating a common schema with which to license the link between the unfilled decoration slot and the drunks. The ability to *create* a model that allows the listener to construe a drunk as a barroom decoration absolves the speaker from the need to store such a mapping. Presumably, the partial nature of analogical matches occurs for similar reasons as the activation of partial (contextually appropriate) structure in other cases of meaning construction. We construct what's relevant to our goals and couldn't formulate any interpretation at all if we activated everything we knew.

In the initial Interior Decoration frame, the redecorators intended to improve (at least) the appearance of the bar and took concrete actions

to do so. However, in the shifted frame, there is no counterpart for this goal-directed activity. The "redecoration" that occurs in the joke consists of a change in the appearance of the bar that occurs presumably in the absence of intentional action on the part of the owners. Moreover, changing the clientele in the bar from one set of undesirable customers to another does not constitute the improvement suggested by the conventional meaning of "redecorated." The nonconventional use of "redecorated" thus results in substantial revision of assumptions about the quantity of and quality of the change to the bar.

Nonstandard Meanings of Nominal Compounds

Yet another problem jokes pose for compositional parsers is the seeming endless number of interpretation schemes for nominal compounds. Providing a compositional way of specifying the meaning of noun-noun compounds has long confounded linguists, although many have specified some semantic regularities in compounding (see, e.g., Downing, 1977; Levi, 1978; or Ryder, 1994). For example, jokes such as (37) proceed by evoking a meaning conventionally associated with a nominal compound such as "army surplus store" and then shifting to a nonstandard meaning assembled according to constructivist schemas not unlike that associated with other compounds.

(37) I opened an army surplus store but nobody wanted to buy surplus armies.

For instance, "army surplus store" conventionally denotes a store that sells items originally intended for use by the army, but that were not in fact used by them. In fact, army surplus stores frequently sell items that were *never* intended for use by the army, but are similar to things that are. Moreover, army surplus stores occasionally carry fashionable clothing items with little or no resemblance to the fatigues actually worn by the military. In the joke, however, "army surplus store" comes to mean a store that sells surplus armies, much as a surplus record store sells surplus records.

Similarly, in (38), the noun-noun compound "balloon tires" is initially interpreted as tires with balloonlike attributes (as in "hermit crab" or "background music"), but is shifted to represent tires for a balloon that functions as a vehicle (as in "car window").

(38) He told me he just bought a set of balloon tires and I told him that I didn't even know he owned a balloon.

Although the change in the interpretation of the nominal compounds in (37) and (38) is quite trivial, the pragmatic implications of this shift are not: (37) requires construction of a scenario in which the narrator markets large, organized bodies of military personnel; and, (38) requires the imaginative construction of the tires on a balloon that serves as a vehicle (perhaps a hot air balloon).

These jokes cannot be comprehended by retrieving prestored frames. In fact, these analyses suggest a process in which abstract frames are integrated with more specific frames to yield novel concepts. The fact that we can arrive at nonstandard understandings of noun-noun compounds implies that their meanings are not specified deterministically by the grammar. The difference between standard and nonstandard meanings of these expressions argues against an identical lexical semantic representation being used in both instances. However, the novel interpretation of noun-noun compounds is not completely unconstrained. Rather, it seems to appeal to alternative strategies used to construct the standard meaning of other, different, compounds. Similarities between the nonconventional interpretations of some compounds and the conventional interpretation of others suggest that speakers use the same set of interpretive strategies in both sorts of cases. These strategies are discussed in more detail in Chapter 5 on conceptual blending in noun phrases.

Nonstandard Meanings for Idioms

Idioms are another problematic phenomenon for compositional accounts of meaning. The way these data are usually accommodated is simply to list the meaning of each idiom in the lexicon. However, consider (39), which capitalizes on a nonstandard interpretation of the idiom "looks like a million (bucks)."

(39) She told him he looked like a million and she meant every year.

The reader begins by assigning conventional meaning to the expression, and evokes a Compliment frame. The second clause is thus interpreted as an expression of the woman's sincerity. At the disjunctor "year," the reader reinterprets "looked like a million" as "looked like a million years." However, this reinterpretation does not involve switching from a noncompositional (idiomatic) meaning to a compositional one (as in the difference between the idiomatic sense of kicked-the-bucket and the literal sense of kicked the bucket): What does a million years look like, after all?

Rather, the reader switches from one noncompositional meaning to another. The reinterpretation of "looked like a million" is what the man would (counterfactually) look like if he were very old – not a million years. The Compliment frame has thus been shifted to an Insult frame, and the woman's sincerity is assigned not to the literal meaning of her hyperbolic remark, but to her intention to insult the man. Because "looked like a million" does not match the lexical entry for the idiom "looked like a million bucks," it is doubtful that a compositional parser could compute the initial interpretation, let alone the more complicated hyperbolic one. Indeed, even if the parser had a gapping heuristic to get "looked like a million years," that parser would mark this utterance as anomalous.

Speakers' interpretation of idioms is more flexible than that of a compositional parser. For a parser to access the idiomatic meaning, the idiom must occur in its canonical form. However, speakers seem to be willing to construct idiomatic interpretations from utterances that are merely similar to the canonical form (as in "looks like a million"). Moreover, in the event that an idiomatic reading is rejected, parsers *must* switch to a literal meaning, while speakers are able to construct noncompositional alternatives. The discrepancy in the performance of people and parsers reflects the fact that while parsers compose meanings from prespecified meanings of their constituents, people construct meanings by flexibly retrieving knowledge in response to the content of utterances, existing knowledge of the context, and their understanding of the goals of the speaker, hearer, and other audience.

Nonstandard Meanings of Phrases

As noted previously, many jokes highlight how normal interpretive processes rely on knowledge of typical situations in order to construct a plausible interpretation. For example, contrast the meaning of "join him in" in (40) and (41).

(40) He asked if I'd join him in the pool.
(41) He asked if I'd join him in a cigarette.

An affirmative answer to (40) will result in both parties being contained in the pool, presumably swimming. In contrast, an affirmative answer to (41) will not result in both parties being contained in a single cigarette. Rather, the reader can invoke her general knowledge of people's interactions with cigarettes to construct a frame in which each party smokes

his own cigarette in close proximity to the other. In fact, it was phenomena like these that led Schank & Abelson (1977) to propose that people store information in terms of how it can be used in specific situations. However, as Schank (1983) has subsequently argued, an alternative explanation is that this information can be computed by combining abstract information stored as memory organization packets (MOP), with the particulars of a given scenario, as well as information in thematically organized packets (TOP).

Joke (42) capitalizes on the contextual modulation of the meaning of "join him in" to force the reader to shift from a plausible interpretation of this phrase to an implausible one.

(42) He asked if I'd join him in a bowl of soup but I told him we didn't have enough room.

The conventional interpretation, of course, relies on the knowledge that people eat soup. Because the second clause of (42) explains why the speaker feels the need to decline the invitation, the disjunctor "room" is not congruent with the initial interpretation of that request. The frame-shifting process thus maps the initial interpretation into a new frame in which spatial constraints might preclude two individuals from sharing a bowl of soup.

A similar case can be found in (43).

(43) When I asked him how long I should cook the noodles, he said at least ten inches.

Interpretation of the first clause relies on very basic knowledge about cooking: that it takes time. Moreover, because cooking noodles does not change their length, it's not relevant to seek information about their linear extension at this point in the process. In contrast, if we were making the noodles from scratch (God forbid!), one might plausibly ask how long to make the noodles, as in (44).

(44) When I asked him how long I should make the noodles, he said at least ten inches.

The relationship between "how long" and the disjunctor "inches" is set up through the sequence of speech acts that occurs in the course of the sentence. The speaker makes a request for information that is followed by a reply that presumably addresses that request. Because "ten inches" cannot be understood as providing a cooking duration, the content of the initial request is reinterpreted. The linear extension of the noodles

is congruent with the reply "ten inches," although, as noted previously, not with knowledge about cooking noodles.

Our ability to interpret jokes such as these suggests that in processing language, the violation of plausibility constraints need not result in the cessation of processing. Rather, it may prompt the processor to search for more abstract mappings. Recall the expression "looks like a million bucks" from the discussion of idioms. While a literal interpretation of this expression is quite implausible, the idiomatic interpretation is presumably motivated by an abstract correspondence (a mapping) between the way an individual might feel if she were looking at a million bucks, and the way she feels looking at the man in question. Moreover, the reinterpretation of "looks like a million" to "looks like a million years" (as in (39)) is similarly implausible on a literal reading. However, the intended reading in this case also recruits abstract information from the literal reading, namely that the man is very old.

The nonstandard readings constructed for both the phrases in this section and the idioms in the last reveal people's propensity to generate utterances whose descriptive content is only obliquely related to their rhetorical motivation. For example, the motivation for (43) might be to demonstrate the stupidity or rudeness of the man who was asked for cooking information. The creative meaning construction processes proposed here seem quite well suited for conveying these sorts of evaluative judgments and prompting speakers to construct various kinds of interactive frames. This is evident in jokes like (45), where nonconventional meanings for phrases are evoked to make a rhetorical point.

(45) My wife did natural childbirth: no makeup.

The structure of (45) suggests that the phrase following the colon will clarify the meaning of "natural childbirth," perhaps for emphasis. However, "no makeup" is not congruent with the conventional meaning of "natural childbirth." Frame-shifting in (45) involves constructing a frame for "natural childbirth" that is congruent with "no makeup." Because the word "natural" is associated with "no makeup" in other contexts, we can construct a novel frame for this particular context in which "natural childbirth" implies that the woman gives birth while wearing no makeup.

Understanding (45) also involves the use of *pragmatic scales*, a set of objects or scenarios that have been ordered along some dimension. In this case, the scale involves expectations about appropriate and inappropriate settings for wearing makeup. While it is not at all irregular

for women to give birth while wearing no makeup, the speaker's statement suggests that for his wife it is, indeed, an accomplishment. Once again the violation of expectations reveals their tacit operation. Further, this sort of evaluative knowledge is used in making judgments about other aspects of a person. As we will see in Chapter 9, speakers often make statements about one object on a pragmatic scale in order to imply things about other objects on the scale. So, if I asked whether the woman often wears makeup, someone might use the pragmatic scale of likely settings for makeup by saying, "Are you kidding? She wore makeup in the delivery room." The reasoning is that if she wears makeup in an especially unlikely setting, she wears it everywhere.

Besides shifting the conventional frame for Natural Childbirth to the No Makeup frame, the occurrence of this frame within the context of a statement of accomplishment suggests that the wife often wears makeup, but has decided to abstain for the current occasion. Knowing that the wife is unlikely to win over admirers in the midst of labor, the speaker in (45) suggests the delivery room is a particularly inappropriate setting for makeup, and casts the wife's explicit decision not to wear it in a ridiculous light.

Similarly, (46) evokes a nonconventional meaning for "invest in the country," perhaps to adopt a more favorable interpretation for bribery.

(46) One good way to invest in the country is to buy a congressman.

In this example, "congressman" triggers lexical reanalysis of "buy," which in turn triggers pragmatic reanalysis of the Investment frame evoked by "invest in the country." The conventional Investment frame involves committing money to earn a financial return. Moreover, "invest in the country" evokes a frame where the investor commits money to an investment that both earns a financial return for the investor and strengthens the domestic economy. Interpreted in this context, the object of "buy" is constrained to the set of investments that might serve these ends. Because "a congressman" is quite a different thing from (for example) a treasury bond, the meaning of "buy" must be revised.

Another sense of "buy" is readily available and results in the activation of the Bribery frame that prompts reanalysis of "invest in the country." The Bribery frame shares important structural characteristics with the frame evoked by the previous sense of "buy" that aid the mapping between the initial representation and the frame-shifted one. The Agent in the Invest-Buy frame maps onto the Agent in the Bribery frame; the Instrument in both frames is money; and the unfilled Object slot in the

Invest-Buy frame maps onto the congressman (the bribee) in the Bribery frame. While a compositional parser could easily generate the bribery interpretation of "buy," it would not be able to establish its relationship to "invest in the country."

2.3 CONCLUSIONS

Many cognitive scientists have assumed that listeners (at some level) retrieve a lexical entry for each word they hear, that syntactic operations are used to help compose individual word meanings into a semantic (that is, logical) representation, and that the role of background knowledge is to resolve lexical and syntactic ambiguities and to fill in gaps in the discourse. However, as Clark (1983) argues, normal language use is not explained very well by the assumption of a fixed lexicon. Language use is characterized not only by ambiguity, but also by semantic indeterminacy. Rather than having a determinate set of established meanings, most words can be used to evoke an indefinite number of contextually motivated interpretations. Because of the potentially infinite number of these nonce senses, they can't be stored in a finite lexicon. Either the lexicon will be filled with nonce senses, or it will include entries so vague that they don't begin to explain how speakers construct the rich, specific interpretations they formulate in particular instances of language comprehension.

The examples examined in this chapter suggest a complex interrelationship between lexical- and message-level meaning representations. The contribution of an appropriate word meaning depends upon the context in which it appears. Moreover, the message-level meaning is influenced by the particular words that occur in the sentence. Lexical reanalysis can trigger pragmatic reanalysis that results in substantial alteration to the message-level representation. Similarly, the choice of a new frame may change how we interpret the meanings of previously encountered words.

Importantly, the change in meaning that results from frame-shifting in these examples cannot be attributed to compositional mechanisms of reanalysis. While structural regularities in language play an important role in enabling speakers to construct meanings from utterances, syntactic variables considerably underdetermine the possibilities for meaning construction. Few of the examples in this chapter required syntactic reanalysis (as traditionally construed) to arrive at the correct interpretation. Moreover, while many examples required the reader to instantiate

alternative lexical semantic representations for a particular word, the lexical information considerably underdetermines the content of the resultant message-level interpretation.

Compositional accounts of natural language processing do not provide a mechanism for explaining how people are able to create novel meanings for words, phrases, and idioms. Nor do they provide a natural way of accounting for how frame-shifting can change the evaluative and sociocultural significance of actions. The affective consequences of semantic reanalysis suggest that frames evoked by language users are employed for more than cool reason. Besides enabling inference, frames cue affective and behavioral responses. Similarly, utterances are used to set up expectations in the hearer's mind about future actions, and as a way of establishing social relationships.

In sum, the ability to interpret so-called nonce senses represents a clear case of the productive use of language. It is rather unfortunate that the significance of this ability has been largely overlooked by researchers in psycholinguistics. These researchers have tended to focus on aspects of language comprehension that can be explained by a compositional parser. But, while the notion of a compositional parser is useful, it is not tenable as a comprehensive account of natural language processing. The contextual variability of word meaning surveyed in this chapter, including subtle meaning shading (as in the various aspects of "tomato") as well as the more radical contextuality evidenced by nonce senses and nonconventional meanings produced by frame-shifting, undermines the utility of a lexicon composed of a finite set of discrete meanings. The novel meanings that arise in sentences are better attributed to the activation of information represented in frames than to the algorithmic composition of semantic representations. While everyone agrees that background knowledge plays an important role in language comprehension, many investigators have misassessed the relative contributions of compositionality and construction from background knowledge.

CHAPTER THREE

Models of Sentential Integration

France's Christophe Rinero led a group coming off the first two peaks. Then Pantani began his charge toward the front as the leaders headed up 8,650-foot Galibier mountain.

He eventually built up a lead of nearly three minutes at the top of the mountain and even had time to stop and put on a plastic jacket to protect him from the wind and cold on the descent.

Although Ullrich cut the gap slightly on the downhill, Pantani moved away on the final climb to the ski station of Les Deux Alpes, almost 5,400 feet, with the rain worsening.

– Associated Press, *Pantani takes lead as Ullrich fades*

As you read this excerpt from a story about the 1998 Tour de France, perhaps you imagine cyclist Pantani coming from behind, pedaling alone on the ascent of Mount Galibier, and being chased more closely as he climbed to the end of the stage at Les Deux Alpes. Or, perhaps you did not realize that the passage was about a bicycle race, and wondered what the men in the story were doing on a rainy day in the French Alps without jackets. In either case, your comprehension of the passage involves the construction of a model of the events described therein. Moreover, the details of your model will depend both on what you know about the Tour de France bicycle race and what aspects of your knowledge were brought to bear on interpretation of the passage.

Given the obsession cognitive scientists have for combinatorial systems, perhaps it is not surprising that early attempts to account for comprehension of texts like the preceding one appealed to a building block type model – that is, a model that builds meanings of larger structures by combining the meanings of smaller ones. In such a model, sentence meaning is built via the composition of word meanings with syntactic

structure, and text meaning is built by concatenating representations of propositions expressed in the sentences and supplementing them with inferences derived from background knowledge. While the appeal of this sort of a model is obvious, empirical research on the way people understand texts suggests that the main goal of text processing is not to construct a representation of the text per se, but rather to model the scenarios the texts were designed to express (see Zwaan & Radvansky, 1998, for review).

Appreciation of the true goals of comprehension has led many in the text processing field to construe language as a set of processing instructions for setting up the discourse model (e.g., Gernsbacher, 1990). Similarly, the work of cognitive linguists suggests language provides speakers with a different sort of information than we might have assumed. Rather than deterministically providing an interpretation, language provides scattered clues that the speaker is invited to follow. For example, seemingly arbitrary grammatical information turns out to mark such things as perspective and epistemic certainty. Linguistic usages reflect entrenched metaphoric mappings (Lakoff & Johnson, 1980). Further, while the construction of the message-level representation is not determined by grammatical information, that information often provides important cues for how to structure mental spaces (Sweetser & Fauconnier, 1996).

In sum, speakers use linguistic information to help them construct simple cognitive models of the discourse situation by integrating linguistic information with background knowledge and local contextual information. Integral to the constructivist account pursued here is the idea that language cues the retrieval of frames from long-term memory for the construction of cognitive models of the message-level representation. However, while frames might seem like a good idea in the abstract, it is entirely possible that they are empirically inadequate. For example, they might prove to be computationally intractable. Alternatively, it might turn out that people's behavior in on-line language comprehension simply does not support the hypothesis that frames are deployed in meaning construction.

This chapter considers the empirical adequacy of a frame-based approach to sentential integration. I begin by examining evidence for frames and frame-shifting in text processing, where the need for "extra" information has long been appreciated. Section 2 considers evidence for frames and frame-shifting in sentential integration, the process by which individual words contribute to the overall meaning constructed

for a sentence. The final section briefly reviews some attempts to build frame-based computational models, and presents some ideas about how best to proceed. Striving for a comprehensive model of meaning construction, this chapter explores whether conceptual operations used to construct the message-level representation of texts might also be central factors in sentential integration.

3.1 FRAME-SHIFTING IN TEXT PROCESSING

Researchers have appreciated the importance of frame-based knowledge in text comprehension for many years (e.g., Bower, Black, & Turner, 1979; Graesser, Woll, Kowalski, & Smith, 1980). Rumelhart (1980), for example, tested the role of frames in story comprehension by asking participants to answer a series of questions after each sentence in a short vignette. Their progressive responses to these questions – for example, what is going on, who are the characters, and why are they behaving as they are – suggested a process of frame activation and refinement. Further, to allay skeptics' concerns that protocols do not reflect on-line comprehension processes, Rumelhart also collected reading time data for one particular vignette in which protocol analyses suggested that readers' interpretations of the story shifted during the final sentence. As the following example illustrates, two versions of the story were constructed: one consistent with most people's initial interpretation of the paragraph as being about a woman's pet, and another in which the reader learns the story is actually about a fur.

> *Dear little thing. It was nice to feel it again. She had taken it out of its box that afternoon, given it a good brush and rubbed life back into its dim little eyes. Little rogue! Yes, she really felt that way about it. She put it down/on. Little rogue, biting its tail just by her left ankle/ear. When she breathed something gentle seemed to move on her bosom. The day was cool and she was glad she had decided on her little pet/fur.*

Rumelhart gave one version of the story to each of two different groups and compared how long it took each group to read the first half of the story (which was identical in the Pet and the Fur versions), and the second half, which differed at three key points. He found that even when results were corrected for faster average reading rates in the Pet group, participants in the Fur group spent reliably longer on the latter part of the story, presumably because comprehension required reanalysis of their initial Pet interpretation. By violating our expectations, Rumelhart's Fur

passage underscores the role of frames in structuring experience and making inferences that go beyond what's immediately present.

Similarly, Garrod & Sanford (1980) performed a series of reading time studies that contrasted the predictions of a frame-based account with older propositional accounts of text processing (e.g., Kintsch, 1974). These investigators contrasted reading time for a target sentence ("The lawyer was trying to prove his innocence") that contained an element that either (i) had an explicit antecedent (e.g., "by a lawyer") and was suggested by the context (In Court); (ii) had *no* explicit antecedent but was suggested by the context (In Court); (iii) had an explicit antecedent ("by a lawyer") but was not suggested by the context (Telling a Lie); or (iv) neither had an explicit antecedent nor was suggested by the context.

Appropriate Scenario
Title: In Court
Fred was being questioned (by a lawyer).
He had been accused of murder.
Target Sentence: The lawyer was trying to prove his innocence.

Inappropriate Scenario
Title: Telling a Lie
Fred was being questioned (by a lawyer).
He couldn't tell the truth.
Target Sentence: The lawyer was trying to prove his innocence.

Garrod & Sanford (1980) found that when the title allowed the reader to set up the appropriate scenario, the presence or absence of an explicit antecedent made little difference in reading times for the target sentences. However, when the title suggested an inappropriate scenario, the absence of an explicit antecedent led to significantly slower reading times for the target sentence. These results suggest that readers use frames evoked by the appropriate titles to aid their on-line comprehension of texts. In the absence of frame-based cues, readers can utilize repetition of noun phrases to maintain a coherent representation of the text.

Further, Sanford & Garrod (1980) tested predictions of a frame-based account of text processing by using short texts that suggested particular roles for the characters. For example, when reading (1), many people assume John is a student.

(1) John was on his way to school.

A frame-based account thus predicts that any indication from the text that disconfirms this role assignment will prompt frame-shifting and

cause processing difficulty. Consequently, Sanford & Garrod compared reading times for target sentences consistent with the scenario-induced role assignment with those for sentences that required a revised role assignment.

No Role Change
John was not looking forward to teaching maths. [British English]
The bus trundled slowly along the road.
He hoped he could control the class today.

Role Change
John was on his way to school.
The bus trundled slowly along the road.
He hoped he could control the class today.

Sanford & Garrod (1980) found that people spent more time reading the target sentence when it suggested a role change than when it was consistent with the role suggested by the contextually evoked scenario. These data indicate that the message-level representation that speakers construct involves more than a listing of elements likely to occur in that setting. The processing cost exerted by role changing suggests that people entertain fairly specific expectations about the relationships among elements in a scenario.

Thus linguistic cues can be used to evoke frames that can be bound to contextually available elements. The computational utility of frames is to extend the domain of reference from explicitly mentioned entities to implied entities in the scene. Slots in a frame can be thought of as implicit hypotheses about the sorts of entities to which the discourse might refer. Thus the lawyer in Garrod & Sanford's examples need not be mentioned if the title sets up a Court frame. Since frames represent the relationships among their different slots, filling slots in a frame provides a way of understanding specific relationships among new and established elements. Moreover, when incoming information deviates from that predicted by the contextually evoked frame, it triggers frame-shifting to reestablish coherence.

3.2 FRAME-SHIFTING IN SENTENTIAL INTEGRATION

The import of frame-based knowledge is now fairly well established for text comprehension. However, its import for understanding how individual words function in overall sentence meaning is less appreciated. While researchers in a number of areas of psycholinguistics have been

interested in the role of context in on-line language processing, they have tended to assume one of two distinct perspectives. Researchers in text processing have been interested in conceptual operations that enable readers to assemble information across sentences, but have taken it for granted that readers can construct representations of individual sentences. Alternatively, researchers in sentence processing have focused intently on the details of local processing, treating context only insofar as it speeds these operations up or slows them down. Neither group has seriously addressed the issue of sentential integration (but see Foss & Speer, 1991; Hess, Foss, & Carroll, 1995; Sharkey & Mitchell, 1988, for exceptions to the rule).

Presumably the reason that sentential integration has received so little attention is that most psycholinguists think there is nothing to explain. Once the syntactic structure has been unambiguously identified and combined with contextually appropriate word meanings, sentential integration proceeds trivially. As an autonomous level of linguistic representation, syntax provides the crucial information for assembling an interpretation. For example, syntax provides the information that differentiates the meaning of "Jane loves John" from that of "John loves Jane." As Forster (1979) writes, "the whole point of a language having a syntax is to provide a clear and unmistakable indication of the correct interpretation of the sentence."*

However, as we saw in Chapter 2, sentential integration is a nontrivial process of adapting information associated with a word to form a coherent interpretation of the whole. While most researchers assume that ambiguity resolution is the chief problem in language processing, a full account of meaning construction will require solution of the more profound issue of semantic indeterminacy. The interpretation of an utterance differs depending on the speakers, the intended audience, the surroundings of its utterance, and world knowledge about the scenario in question (e.g., Clark, 1983; Sanford & Garrod, 1980; Travis, 1981). Contextual modulation of meaning can range from subtle shading to the construction of novel meanings. As we saw in the last chapter, the comprehension of even unambiguous sentences requires considerable creativity on the part of the language user to construct a specific representation of its meaning.

The creative nature of meaning construction results from the way that words both benefit from and contribute to the construction of the

* Forster, 1979: 55.

message-level representation. Words benefit from context because structure in the message-level representation facilitates the integration of elements and relations consistent with the contextually evoked frame (Hess, Foss, & Carroll, 1995). Because lexical-level expectancies are based, among other sources, on the frames active in working memory, the most easily processed lexical items will be those whose semantic contribution can easily be accommodated by the activated frame (Sharkey & Mitchell, 1988). Moreover, words also contribute to the message-level representation – by providing cues for addressing knowledge in long-term memory, for properly construing current conceptual content, and by triggering the frame-shifting process.

The relationship between a word and its context involves how it adds to the cognitive models active in working memory, as well as how it prompts the construction of new models. As noted previously, while frame-shifting cues clearly impact the processing of texts, its impact on the processing of words in isolated sentences is less appreciated. However, if this view is correct, it predicts that even in isolated sentences, words that trigger frame-shifting should present a challenge to the processor that differs from that presented by lexical violations consistent with the currently active frame. To test this hypothesis, Coulson & Kutas (1998) manipulated the relationship between sentence final words and their preceding contexts by comparing reading times for words that triggered frame-shifting, to the reading times for equally unexpected words consistent with a contextually evoked frame. The primary question was whether the time needed for sentential integration of unexpected words would vary as a function of the *type* of unexpected ending.

Because one-line jokes provide such an excellent source of identifiable frame shifts, we began by contrasting participants' self-paced reading times for jokes and nonjoke controls. Further, unlike many cases that present the need to shift frames in light of incoming information, in jokes the speaker clearly cues the need to shift. Joke theorists such as Attardo, Attardo, Baltes, & Petray (1994) call this cue the *disjunctor*. In (2), for example, the reader begins by evoking a frame where a busy professional pays an accountant to do his taxes.

(2) I let my accountant do my taxes because it saves time; last spring it saved me ten years.

The disjunctor "years" prompts the reader to reinterpret "time" so as to evoke a frame where a crooked businessman pays an accountant to conceal his illegal business dealings. The word "time" is called a

connector because it serves as a bridge between the two frames. Understanding that "time" refers to time in prison does not in and of itself explain why the accountant is doing the man's taxes, or how doing so will prevent a prison sentence. This requires recruitment of background knowledge about particular sorts of relationships that can obtain between business people and their accountants. The initial busy-professional interpretation is thus mapped onto a new space structured by the crooked businessman frame.

Stimuli like (2), in which understanding the joke requires semantic reinterpretation of meaning established earlier in the sentence, were assembled from joke books. In all cases, the disjunctor (the point at which the reader could, in principle, realize the joke) was a sentence-final noun. To establish the nonjoke interpretation of these materials, Coulson & Kutas stripped each joke of its disjunctor and normed them on the *cloze task.* In the cloze task, a large number of people are given sentence fragments and asked to complete the sentence with the first word or phrase that comes to mind. A word's cloze probability in a given sentence is the percentage of people who generated that particular word on the cloze task. By looking at the most popular response on the cloze task, Coulson & Kutas were able to establish the default nonjoke interpretations of the joke fragments.

Interestingly, stimulus pretesting revealed considerable variability in the cloze probabilities of the most popular response (that is, the default nonjoke interpretation). For instance, in sentence fragment (3), the most popular response was produced by 81% of the participants.

(3) I asked the woman at the party if she remembered me from last year, and she said she never forgets a (face 81%).

However, for sentence fragment (4), people produced a wide array of responses so that the cloze probability of even the most popular response for this sentence fragment was only 18%.

(4) My husband took the money we were saving to buy a new car and blew it all at the (casino 18%).

Consequently, sentence fragments were divided into two groups: high-constraint sentence contexts – that is, fragments in which the cloze probability of the most popular response was greater than 40% – and low-constraint contexts (less than 40%).

To test the hypothesis that frame-shifting affects the processing of individual words, Coulson & Kutas (1998) compared word-by-word

reading times for joke and nonjoke endings of low- and high-constraint sentences such as the following:

Low-Constraint Sentence
My husband took the money we were saving to buy a new car and blew it all at the
Nonjoke Ending: tables. (Cloze probability: 3.6%)
Joke Ending: movies. (Cloze probability: 3.6%)

High-Constraint Sentence
I asked the woman at the party if she remembered me from last year, and she said she never forgets a
Nonjoke Ending: name. (Cloze probability: 3.6%)
Joke Ending: dress. (Cloze probability: 3.6%)

Although the sentence-final words in this study were equally unpredictable, people spent appreciably longer on sentences that ended as jokes than those that had the nonjoke endings. Both joke and nonjoke endings had a low cloze probability. But while nonjoke endings were constructed so that they would be congruent with a contextually evoked frame, joke endings required the initiation of frame-shifting. The researchers' observation of longer reading times for joke endings is consistent with the suggestion that semantic reanalysis prompted by the integration of the joke ending exerts a greater processing cost than that associated with integrating an unexpected word congruent with the contextually evoked frame (Coulson & Kutas, 1998).

However, the frame-shifting effect was confined to the high-constraint sentences (such as (3)) that suggested a particular lexical item as the best completion for the sentence. In spite of the fact that participants in this experiment never encountered the *expected* completion, results suggest that high-constraint sentence contexts aided sentential integration of unexpected items consistent with the message-level frame. Because high-constraint sentences were those that strongly suggested a particular lexical item on the production task, it may be that this task is an indirect index of the likelihood a given reader will commit to a particular frame. Perhaps high-constraint sentences allow the reader to commit to a frame with which to structure language input, whereas low-constraint sentences engender less frame commitment.

Researchers interested in lower-level aspects of language processing – such as word recognition or parsing – would do well to consider how those processes are driven by higher-level communicative goals

(Sanford & Garrod, 1980). Processing costs that psycholinguists have attributed to low-level processes such as word recognition may instead reflect differences in the complexity of the high-level processes of meaning construction. The joke endings were harder to process not because (or, at least, not *simply* because) context served to "inhibit" them (although this may be true), but because understanding the last word of a joke involves recruiting a new frame to structure the message-level representation and mapping existing information into the frame. Similarly, phenomena investigated under the rubric of sentence processing (that is, parsing) are undoubtedly influenced by the demands of meaning construction supported by contextually activated frames.

Pickering & Traxler (1995) present eye tracking data that suggest people are more hesitant to abandon a semantically plausible misanalysis than an implausible one. Comparison of people's gaze durations for sentences such as (12) and (13), with sentences that included a clarifying comma, suggested participants initially misanalyzed "the magazine" as being the object of the verb "edited" in (12) and of the verb "sailed" in (13).

(12) As the woman edited the magazine about fishing amused all the reporters.

(13) As the woman sailed the magazine about fishing amused all the reporters.

In spite of the implausibility of sailing a magazine, Pickering & Traxler's data suggested that "sailed" in (13) was initially analyzed (incorrectly) as a transitive verb. However, their results also suggested that people were more willing to abandon their initial misanalysis in (13) than in (12). Comparison of gaze duration to individual words in (12) and its unambiguous control suggested readers began to experience processing difficulty at the verb "amused." However, a similar comparison in (13) suggested readers began to experience processing difficulty at "magazine." Although both sentences present the same parsing ambiguity, readers initiated reanalysis in (13) sooner than in (12). Evidently, readers were more committed to the misanalysis in (12) because it was more plausible.

Further, Pickering & Traxler report a similar pattern for sentences in which the plausibility manipulation involved the use of contextually motivated nonstandard meanings for words. For example, they contrasted gaze durations in sentences such as (14), preceded by a context that mentioned bronze statues *of* particular professors, with those preceded by a context in which a janitor was polishing statues *for* particular

professors.

(14) While the janitor was polishing the professor of physical chemistry prepared a lecture in his office.

As in the previous experiment, the pattern of gaze durations suggested people were less willing to abandon the polish-the-professor garden path when context suggested a plausible reading for this misanalysis. Given that, prima facie, it is quite implausible to polish a professor, this finding is all the more remarkable. Apparently commitment to particular interpretation depends less on general background knowledge than on context-sensitive plausibility. As such, these data underscore the need for a creative model of meaning construction.

Moreover, Pickering & Traxler's findings point to a possible explanation for Coulson & Kutas' report that the frame-shifting effect – the difference between reading times for jokes and nonjokes – was more robust in the high-constraint sentences. Because jokes are merely syntactically unambiguous sentences misanalyzed on semantico-pragmatic grounds, the greater effect of frame-shifting in high-constraint stimuli could have been due to a commitment effect. Perhaps readers were more hesitant to abandon their initial frames in high- than in low-constraint sentences. Indeed, on Langacker's (1987) theory of cognitive grammar, grammatical phenomena reflect the operation of schematic frames, and thus traditional garden-path phenomena represent frame-shifting at a very abstract level. Likewise, Jurafsky (1994) proposes a parsing algorithm based on construction grammar that predicts similar reanalysis effects for all sorts of constructions, be they lexical or grammatical.

Although most psycholinguists have approached the issue of meaning construction from the perspective of how low-level processing of words is influenced by the developing message-level representation, it is also important to consider how the processing of individual words influences the *creation* of the message-level representation. Priming effects usually attributed to mechanisms that facilitate or inhibit low-level processes such as word recognition and lexical access may instead (or also) reflect differences in the complexity of high-level processes of meaning construction. Indeed, a number of results argue against the traditional view that semantic priming reflects spreading activation in the lexicon (c.f. Meyer & Schvaneveldt, 1971). For example, McKoon & Ratcliff (1986) have shown that priming occurs for newly learned associates, and need not reflect semantic memory. Moreover, Hess, Foss, & Carroll (1995) demonstrate effects of global context on naming latencies that strongly suggest an extralexical locus of priming effects.

In the framework proposed here, word processing involves much more than activating static representations in the lexicon. Words serve as triggers for the creative construction of cognitive models in working memory. Because words trigger context-sensitive retrieval of information from long-term memory, the conceptual structure recruited in response to lexical input will not always be the same. Consequently, contextual manipulations serve to modulate the complexity of the processes that a given word initiates, as well as the complexity of obligatory processes active in all word processing. This suggests that frame-shifting effects such as those identified by Coulson & Kutas ought not to be confined to jokes, but should be evident in nonjoke materials that require frame-shifting.

To test this hypothesis, Coulson & Kutas (1998) compared participants' self-paced reading times for humorous and nonhumorous examples that required frame-shifting. Examples of the four different kinds of sentences follow:

Humorous Materials (Joke Versus Nonjoke)
I asked the woman at the party if she remembered me from last year, and she said she never forgets a
Nonjoke Ending: name.
Joke Ending: dress.

Nonhumorous Materials (Shift Versus Nonshift)
She couldn't get herself in because she had lost her
Nonshift Ending: pass.
Shift Ending: memory.

Collapsed across the frame-shifting manipulation, people spent longer reading the last word of the joke/nonjoke than the shift/nonshift materials. However, in both the humorous and the nonhumorous stimuli, people spent about 100 milliseconds longer reading the word that triggered frame-shifting. These data suggest that frame-shifting is not confined to one-line jokes, but is a more general phenomenon. Nonhumorous materials that required frame-shifting also elicited increased reading times relative to equally unexpected controls designed to be consistent with contextual expectations at the message level.

The results of Coulson & Kutas (1998) support the psychological reality of frame-shifting, even in single sentences. The results are thus broadly consistent with the frame-based approach discussed previously. In view of the computational complexity of this operation, perhaps the

most remarkable finding was not the existence of an effect of frame-shifting but that it was on the order of hundreds of milliseconds. Given the dramatic reorganization of the conceptual representation implied by frame-shifting, this relatively small difference in reading times is quite amazing. The next section briefly considers the issue of whether a frame-based system can truly accommodate the computational demands of frame-shifting.

3.3 FRAME-BASED MODELS

The constructivist approach to comprehension sketched in Chapter 2 relies extensively on the idea that the meaning construction process uses frames. The appeal of frames lies chiefly in their ability to account for all the "extra" information readers infer in the course of meaning construction. Because slots in a frame can be thought of as hypotheses about the sorts of entities that might be encountered in a given scenario, a frame-based approach can deal nicely with the fact that people develop expectations in the course of language comprehension. More importantly, a frame-based approach can capture the extent to which readers conceptualize the relationships that obtain among various objects, actions, and events described in natural language utterances.

Further, the semantic reorganization involved in frame-shifting would seem to require that the message-level representation have at least *some* properties of frames. In particular, the computational challenge of connecting an initial interpretation to the reinterpretation seems at least to require the representation of causal and relational information, slot/filler organization, and the existence of default values. However, given various critiques of frames (e.g., Allen, 1987; Wilensky, 1986), it seems relevant to question whether frame-based models can accommodate the flexibility needed to implement frame-shifting. This section considers several frame-based models with respect to how well they can handle the computational requirements of frame-shifting.

The primary problem with frame-based systems is that they are overly brittle. Wilensky (1986), for example, complains that scripts are rigid data structures and cannot accommodate events that are out of the ordinary. He argues convincingly that while the restaurant script works well when the events in the restaurant coincide with those in the script, the system cannot adequately represent unexpected events. Of course, a script-based system is not completely unable to deal with unexpected events. It can, for example, create a new slot in response

to the unexpected event. However, because its inferencing capacity is based in knowledge represented in the script itself, it cannot compute the relationship between the unexpected event and its more normal counterparts. Moreover, since frame-shifting often requires adapting models to accommodate unexpected or atypical events, traditional implementations of frames are clearly inadequate.

3.3.1 Frames and Constraint Satisfaction

Rumelhart, Smolensky, McClelland, & Hinton (1986) have also argued that there is a gap between the explanatory goals that motivate cognitive scientists' conception of frames and schemas as knowledge structures, and the way those ideas have been implemented in computational models of cognitive processing. The need for a sufficiently flexible implementation of frames has driven these investigators to explore the adequacy of subsymbolic processing in neural networks (see McClelland & Rumelhart, 1986, for review). Rumelhart et al. argue that frames are not static knowledge structures but stable states in the energy landscape of constraint satisfaction networks. Units in the network represent semantic microfeatures, and the weights between units encode correlations between those microfeatures.

To illustrate, Rumelhart et al. describe a network that represents features of a number of household rooms such as the kitchen or the living room. Units in the network represent the presence of semantic microfeatures, including things like has-a-ceiling, has-walls, has-a-toaster, and so forth. The weight between two units is a function of the conditional probability of their cooccurrence. The network is set up to promote excitatory weights between microfeatures that cooccur and inhibitory weights between microfeatures that do not. Consequently, when one of the units is clamped on by the programmers, a gradient descent algorithm is used to fill in the frame via the activation of correlated microfeatures and the inhibition of others until the network settles into a stable state.

Stable states in the energy landscape of the network correspond to instantiated frames. In the network described by Rumelhart et al., the energy landscape allowed for the possibility of 2^{40} stable states. However, because of the configuration of the weights, their network always settles into one of five states corresponding to five different room frames. Although frames in the network are not stored in the traditional sense, they are implicit in the way weights between units reflect the cooccurrence of the various microfeatures.

Filling slots in these networks involves filling in subpatterns that vary from one instantiation of the schema to another. Default values are those variable subpatterns that are automatically activated by the settling process. Moreover, because filling in subpatterns is a function of the weights in the network, the value of one variable can constrain the values of the other variables. While this sort of nonmonotonic slot-filling can be hard to achieve in symbolic implementations, it arises naturally out of the algorithm used to train the network. In fact, this sort of context sensitivity – for example, knowing that different motives underlie diving into a pool of water rather than into an empty pool – accords well with the observations discussed in Chapter 2.

The rigidity of a frame will depend on the correlations among the activations of its constituent microfeatures. If the network has experienced a high correlation among the mutual activation of stove, refrigerator, and kitchen, it will result in a narrow energy landscape that the network is unlikely to settle into unless the input matches the specifications of the frame very closely. On the other hand, if the concomitant activation of these features displays more of a family resemblance pattern (e.g., kitchen *with* stove but *without* refrigerator; kitchen *with* refrigerator but *without* stove, etc.), the energy landscape will be broader, reflecting the increased set of inputs that can satisfy the frame.

The main advantage that PDP (parallel distributed processing) implementations of frames have over traditional symbolic ones is flexibility in accommodating inputs that do not conform to previously instantiated frames. In the model described by Rumelhart et al., even central elements of a frame can be missing and the network will nonetheless settle into a stable state. But while the metaphor of interpretation as a constraint satisfaction process seems a profitable way to approach sentential integration, the particular network described by Rumelhart et al. would fail miserably in the representation of information in the jokes discussed in Chapter 2. Although one might design a network of this sort that could fill in various aspects of the fishing script, for example, it would not contain any mechanisms for high-level inferences that relate frames to one another.

3.3.2 The Story Gestalt Model

An example of a PDP model that can draw high-level inferences is St. John's (1992) Story Gestalt model, which combines information from multiple sources to fill in aspects of the message-level "gestalt"

representation. For example, given a story about dinner in a restaurant, the model can combine information about who paid the bill, whether that character is cheap or extravagant, and the quality of the restaurant, in order to predict the size of the tip. Rather than the perceptron employed by Rumelhart et al., St. John uses an architecture more suited to time series predictions used in script-based reasoning (the single recurrent network). But, like Rumelhart et al.'s (1986) model, the Story Gestalt model can infer default information based on cooccurrence frequencies in its input. Coherence inferences (e.g., inferring that a customer ordered from the information that he paid) and predictive inferences (e.g., inferring that the customer will tip from the information that he paid) both occur because information associated with characters and actions in the text are automatically activated.

The Story Gestalt model employs graded constraint satisfaction to represent text, resolve pronouns, draw inferences, and revise assumptions. St. John (1992) views the text as imposing a set of constraints on the formulation of the message-level "gestalt" interpretation. The virtue of this approach is that it provides a mechanism for combining multiple sources of information. The graded nature of constraints affords a means of representing the uncertainty that marks much of our interaction with the world. Moreover, revision and updating can be accommodated by modification of the weights, as opposed to computationally costly backtracking and path evaluation procedures.

For instance, take the Bar script:

[Person1] decided to go to the bar.
[Person1] made a polite/obnoxious pass at [Person2].
[Person2] gave a kiss/slap to [Person1].
[Person1] rubbed lipstick/cheek.

Given the information in the second event in this script, the model should be able to predict the probability of the last action as 70:30 chance of rubbed cheek (obnoxious pass) and 30:70 chance of rubbed lipstick (polite pass). However, with the knowledge of the third event in the script, the model should be able to predict the last event deterministically with perfect validity. Moreover, when the model encounters an unexpected event, such as a kiss after an obnoxious pass, it automatically modifies its prediction from "Person1 rubs cheek" to "Person1 rubs lipstick."

St. John (1992) thus provides a model that illustrates how reanalysis of local predictive inferences can be constrained by higher-level

information in a script. However, this reanalysis is not an example of true frame-shifting, as only one higher-level frame is being recruited. Moreover, in spite of the flexibility afforded by connectionist-style modeling, the frame-shifting ability of the Story Gestalt model is limited in much the same way as symbolic implementations. Because it is unable to compute the relationships between different higher-level representations, St. John's (1992) model is incapable of combining information from different scripts in any sensible way. While networks with distributed representations are well suited for learning categories or other functions that can be induced from cooccurrence statistics, they often lack the representation of structure needed to understand conceptual relationships between elements in a discourse (Feldman, 1989).

3.3.3 Structured Connectionism

In an attempt to combine the representational capacities of symbolic natural language processing systems with the processing advantages of parallelism, Lange & Dyer (1989) argue for an alternative framework known as *structured connectionism*. In a structured network, each node stands for a distinct concept, and connections between nodes represent structural relationships between concepts. Lange & Dyer's (1989) model ROBIN (ROle BInding Network) uses connections between nodes to encode semantic knowledge represented in a frame-type data structure. Each frame has one or more roles or slots, and slots have constraints on the type of fillers to which they can be bound. The relationships between frames are represented by excitatory and inhibitory connections between nodes, and pathways between corresponding slots. Once initial slot assignments have been made, ROBIN propagates evidential activation values to compute inferences from the information the programmers have provided.

Inference occurs as spreading activation propagates evidential activation across the excitatory and inhibitory connections that exist between related frames. For example, connections between Transfer-Inside and Inside-of allow the system to infer that Inside-of(Bread, Oven) results from Transfer-Inside(Seana, Bread, Oven). In ROBIN, frame selection is entirely a matter of spreading activation. Because each slot has a number of binding nodes, all the meanings of an ambiguous word can serve as candidate bindings. Candidate bindings can be propagated simultaneously, and the binding node with the greatest evidential activation eventually wins out. Because multiple frames are activated in parallel,

contextual information can further activate an already highly activated node (or set of nodes), thus confirming an initial interpretation. Alternatively, contextual information can activate a previously less active interpretation, thus implementing frame-shifting.

While the use of spreading activation to implement frame-shifting looks promising (as does the related model presented by Shastri & Ajjangadde, 1993), Lange & Dyer's model has some serious deficiencies. First, ROBIN can't represent scenarios that require recursion. Further, it has no mechanism for adding new concepts, other than having a programmer add new frames and their connections to all the other frames in the model. Moreover, while ROBIN can cope with lexical ambiguity, it is completely unable to deal with what Clark (1983) has called semantic indeterminacy. ROBIN is unable to handle metonymyic or metaphoric uses of language. Nor is it able to adapt frames to accept new fillers (as in the concept "same-sex marriage," cf. Turner & Fauconnier, 1995). Finally, the representation of each concept with a single node limits the model's ability to represent the meaning shading found in natural language.

However, Lange & Dyer suggest one improvement to the system that is well worth pursuing: the use of distributed rather than localist representations for signature activations. This would give some semantic content to the fillers that would allow for meaning shading and would be amenable to relaxation of binding constraints. Distributed representations might help the system accommodate contextual expressions rather than merely dealing with computational problems caused by ambiguity.

3.3.4 Space Structuring

We have considered several computational models that implement frames and found all of them wanting in one respect or another. As noted previously, the main problem with frames is that they are too rigid to represent the relationship between expected and unexpected events. However, Schank (1982) introduced the idea of MOPs, memory organization packets, in order to deal with the "brittleness" issue. MOPs are generalized clusters of *scenes*, or the high-level components of scripts (e.g., a set of Entering scenes for entering a grocery store, a dentist's office, a restaurant, and so on). Rather than employing encapsulated scripts such as the restaurant script, Schank (1982) uses a combination of general MOPs that deal with exchanges and services, and more specific MOPs that have knowledge about particular restaurants.

Schank (1982) found that encoding information at various levels of specificity maintained the representational advantages of scripts while allowing for more combinations and recombinations of scenes at various levels. Further, Kellermann, Broatzmann, Lim, & Kitao (1989) have suggested that MOPs are more psychologically plausible memory representations. Citing data on memory confusions between stories that recruit distinct (though similar) scripts, Kellermann et al. contend these data argue for the use of abstract MOPs rather than the highly encapsulated packets of information originally proposed by Schank & Abelson (1977).

Consequently, I suggest a generalization of Schank's (1982) proposal in which the representations underlying language – even at the most rudimentary levels – have important properties of frames such as slot/filler organization, default values, and hierarchical structuring. As stated elsewhere, grammar does not algorithmically specify context invariant meanings, but provides clues to help the language user construct a message-level representation (Fauconnier, 1994; Gernsbacher, 1990; Turner, 1991). Grammatical phenomena reflect the operation of very abstract frames for construing the conceptual content evoked by linguistic utterances (Langacker, 1987; Kay, 1997). The role of grammar in this approach is to provide cues for the coordination of abstract grammatical frames evoked by particular words, perceptual information, and the existing message-level representation.

Moreover, the account of meaning construction sketched here goes somewhat beyond traditional constructivist accounts (Rumelhart, 1980; Sanford & Garrod, 1981) in its acknowledgment of the creative nature of language comprehension. Speakers do not just recall frames from long-term memory, but construct partial frame-based representations in mental spaces. As discourse unfolds, speakers add elements and relations to existing spaces in working memory. Moreover, while the structures in spaces are partial, their operation is constrained by information in long-term memory (as in Glenberg, Kruley, & Langston, 1994). Speakers set up new spaces in response to space builders (such as "I think") as well as to less explicit cues. For example, if extending a space would result in the representation of contradictory or irrelevant information, a new space is set up. In cases of frame-shifting, information in the original space is mapped into the new one.

While these suggestions lack the explicitness of a computationally implemented model, they do suggest three features that characterize meaning construction: partial representational structure, the importance of mapping, and partitioning into spaces. Data reviewed in this chapter

support the suggestions of partial structure and the importance of mapping. Moreover, data from text processing literature are broadly consistent with the suggestion that language comprehension involves setting up and structuring spaces. For example, we might expect that delimiting one mental space and setting up a new one would utilize processing resources. Accordingly, people slow down when they encounter sentences in a text that change the topic (Haberlandt, Berian, & Sandson, 1980), point of view (Bower, Black, & Turner, 1979), location, and temporal setting (Anderson, Garrod, & Sanford, 1983).

Further, we might expect structure in the delimited spaces to be less accessible than in the focus space where meaning is currently being built. Indeed, previously read material becomes less accessible after grammatical events that signal an end to space construction, including clause boundaries, phrase boundaries, and sentence boundaries (Caplan, 1972; Jarvella, 1973; von Eckhardt & Potter, 1985). Moreover, Haenggi, Gernsbacher, & Kintsch (1995) have demonstrated that the accessibility of information is affected by the location of a fictional character. For example, participants took longer to verify whether or not various objects were "in the ballroom" after they read, "The king left the ballroom," than when they read, "The king was in the ballroom."

Mental space theory, so construed, is largely compatible with Gernsbacher's Structure Building framework (reviewed in Gernsbacher, 1998). On Gernsbacher's framework, information in the developing substructure is most accessible to speakers. Likewise, on the space structuring model, the focus space is most accessible. Further, Traxler, Sanford, Aked, & Moxey (1997) show that the processing difficulty of "because" constructions such as those in (4) and (3) depends on whether explicit space builders cue the correct configuration of mental spaces.

(3) The streets are wet because it's raining.
(4) It's raining because the streets are wet.

Without any prior context, causal statements such as (3) are read faster than diagnostic statements such as (4), suggesting that the default interpretation of "because" is as a causal connective. However, by employing a space builder such as "John thinks," it is possible to bias a diagnostic reading of "because," and consequently reverse the pattern of reading times. That is, (4') is read faster than (3').

(4') John thinks it's raining because the streets are wet.
(3') John thinks the streets are wet because it's raining.

Thus when properly motivated, the diagnostic statement in (4) is processed faster than the causal statement (3) in the same context.

In sum, the suggestion is that linguistic utterances cue the retrieval of abstract grammatical frames that speakers integrate with frame-based structures evoked by lexical and contextual information in mental spaces. Meaning construction thus consists of mapping cognitive models from space to space while keeping track of the links between spaces and between elements and their counterparts. Data reviewed in this chapter are broadly consistent with these suggestions. Moreover, the role of background information in such a system goes far beyond that of filling in gaps: The creative integration of current experience with background knowledge is the very essence of comprehension.

CHAPTER FOUR

Frame-Shifting and the Brain

Scales fall from the eyes, the light dawns, the structure is suddenly apparent, and so forth, sometimes on a timescale of seconds.*

– Paul Churchland, *A Neurocomputational Perspective*

We have seen that language comprehension involves a complex set of processes in which listeners actively construct meanings by integrating linguistic input with background knowledge and local contextual information. Pursuing a constructivist approach, I have suggested that language cues the retrieval of frames from long-term memory for the construction of cognitive models in working memory. Language input, perceptual input, and current conceptual content all influence the recruitment of information as well as the construction of models. Comprehension proceeds by assembling a series of simple cognitive models and establishing the mappings, or correspondences, that exist between the elements and the relations in different spaces, or partitions of working memory.

When a word cannot be integrated into the model currently being built in the focus space, it may trigger *frame-shifting*, or reanalysis of information that has come before. Frame-shifting occurs when the currently activated frame does not adequately represent the relationship among two or more objects, actions, or events. While jokes provide a particularly salient example of the frame-shifting phenomenon, the need to construct novel frames to accommodate new information is quite common. Because sentential integration involves the coordination and accommodation of different frames, people presumably appeal to

* Churchland, 1989: 236.

mechanisms of frame-shifting in the course of everyday language understanding.

However, as Churchland (1989) observes, frame-shifting can be accompanied by an experience of sudden insight. Can this remarkable event really be related to the more common processes of on-line meaning construction? Traditional models suggest the answer is no. In such models, the semantic leaps in jokes represent a departure from normal, literal language use. Literal language use relies on meanings specified in the lexicon, and recruits grammatical information to combine word meanings into sentence meaning. Nonliteral language use, on the other hand, relies on literal meanings and appeals to encyclopedic world knowledge, recruiting general reasoning processes to derive what is meant from what is said.

However, psycholinguistic research in the past two decades has consistently undermined traditional distinctions between literal and nonliteral language processing (see Gibbs, 1995, for an extensive review). Similarly, previous chapters have shown how the interdependence of meaning and context is an inherent feature of meaning construction. We have seen how the productive character of language use results from constructive processes that speakers use to assemble, link, and adapt simple cognitive models. The extent to which language use involves appeals to background and contextual knowledge suggests that the difference between so-called literal and nonliteral language use may be quantitative rather than qualitative.

In contrast, neuropsychologists have suggested that the old distinctions between literal and nonliteral meaning, semantics and pragmatics, are supported by differences in the character of language breakdown caused by damage to the left versus right hemisphere of the brain. Damage to the left hemisphere can result in severe disruption of language competence affecting capacity for phonological, syntactic, and semantic aspects of language. In contrast, damage to the right hemisphere is more likely to cause subtle communicative deficits involving the relationship between an expression and its context of utterance. Given these differences in the stereotypical profile of communicative disorders, neuropsychologists have assumed that normal left hemisphere function corresponds to phonology, syntax, and semantics, while normal right hemisphere function corresponds to pragmatics; left hemisphere to literal, context-invariant aspects of meaning, and right hemisphere to nonliteral, derived meanings.

This difference in the treatment of the reputed distinction between semantic and pragmatic aspects of meaning presents a puzzle that is explored (though not solved) in this chapter. The first section considers data from neuropsychology that suggest a critical role for the right hemisphere in frame-shifting. Neuropsychologists' attempts to map traditional linguistic distinctions onto the brain are certainly motivated by real distinctions in the character of language deficits resulting from different sorts of brain damage. However, the interpretation of language breakdown along traditional lines might also reflect a lack of alternative frameworks.

Because it is always difficult to infer the nature of normal brain function from the behavior of patients suffering the effects of brain damage, section 2 turns to data from event-related brain potentials (ERPs) recorded from neurologically intact individuals. A continuous on-line measure of language processing, ERPs can address the issue of precisely when the brain begins to differentiate different categories of stimulus events. Further, ERPs can potentially discriminate qualitatively different processing events from those that differ only quantitatively. Data suggest that while the marked frame shifts in jokes are more complicated than for other cases of sentential integration, the processes recruited for meaning construction may be essentially similar.

4.1 NEUROPSYCHOLOGY

The standard division of language into literal and nonliteral, semantics and pragmatics, has been supported by observations of functional dissociations in patients with unilateral lesions in either the left or the right cerebral hemisphere. While patients with left hemisphere damage (LHD) can experience difficulty with even the most basic aspects of language production and comprehension, patients with right hemisphere damage (RHD) exhibit communicative deficits that are far more subtle. Right hemisphere patients often make socially inappropriate remarks, and their speech is frequently tangential, and marked by digressions of topic (Joanette, Goulet, & Hannequin, 1990). Moreover, they have a tendency to underutilize, or to utilize inappropriately, nonlinguistic communicative clues such as facial expressions and prosodic intonation (Ross, 1981). Overall, RHD patients' spontaneous speech suggests an impaired ability to modulate linguistic output to meet the cognitive and social demands of context.

Interestingly, difficulty in understanding jokes has been specifically associated with brain damage in the right hemisphere. Discrepancies between the ability of patients with RHD and LHD to understand jokes were first noted by Gardner, Ling, Flamm, & Silverman (1975). Moreover, subsequent research has suggested that joke comprehension deficits associated with RHD involve the conceptual demands of frame-shifting. For example, Brownell, Michel, Powelson, & Gardner (1983) gave RHD patients the first few sentences of a joke and asked them to pick the punchline from an array of three choices: a straightforward ending, a non sequitur ending, and the correct punchline. While age-matched controls had no trouble choosing the punchlines, RHD patients tended to choose the non sequitur endings.

Drawing on Suls' (1983) two-stage model of humor, Brownell et al. argue that RHD patients retained the *surprise* component of joke processing, in which a listener recognizes the incongruency between the punchline and the expectations engendered by hearing the first part of the joke. However, the RHD patients were impaired on the *coherence* stage in which coherence is reestablished between the narrative and the punchline. Suls' coherence stage thus coincides with the need to map information in the initial message-level representation into a new frame evoked by the disjunctor.

The pattern of deficits in RHD patients differs dramatically from those evidenced by LHD patients whose communicative difficulties are seemingly more severe. To compare the performance of LHD and RHD patients on the comprehension of jokes, Bihrle, Brownell, & Gardner (1986) used both verbal and nonverbal materials. In addition to using jokes of the sort employed by Brownell et al. (1983), these investigators also used four-frame cartoons with the narrative structure of a joke. Whether patients received verbal or nonverbal materials, they were asked to pick the punchline (or punch-frame for the cartoons) from an array of four choices: a straightforward ending, a neutral non sequitur, a humorous non sequitur, or the correct punchline.

While both sorts of patients were impaired on this task, their errors were qualitatively different. RHD patients, tested on both the verbal and nonverbal materials, showed a consistent preference for non sequitur endings over straightforward endings and correct punchlines. In contrast, LHD patients (who participated only in the nonverbal task) chose the straightforward endings more often than either of the non sequitur endings. These data indicate that the deficits RHD patients experience in the comprehension and production of jokes are not attributable to

the emotional problems that have been associated with some kinds of RHD (e.g., Gainotti, 1972; Goldstein, 1948). Although the RHD patients in Brownell et al.'s study were unable to choose the correct punchline, they did appreciate the slapstick humor depicted in the non sequitur endings.

Moreover, Brownell, Potter, Bihrle, & Gardner (1986) have also demonstrated that RHD patients have difficulty interpreting *nonjoke* materials that require what we have been calling frame-shifting. Brownell and colleagues gave patients vignettes such as the following:

> *Sally brought a pen and paper with her to meet the famous movie star. The article would include comments on nuclear power by well-known people.*

When reading this passage, RHD patients persisted in the original Autograph interpretation, and seemed unable to integrate this information with the Interview frame suggested by the second sentence. However, when the misleading information was presented in the second sentence rather than the first, the RHD patients' performance improved considerably. These observations suggest that the RHD patients' difficulty with jokes is cognitive rather than emotional, and specifically involves the inferential reanalysis of frame-shifting.

While some researchers have suggested RHD patients have difficulty drawing *coherence inferences* that connect sentences in discourse (Beeman, 1993; Brownell et al., 1986), most data suggest these patients have the opposite problem: They draw too many inferences, and are unable to reject the irrelevant ones. Using a recognition task, MacDonald & Wales (1986) found RHD patients were able to make simple inferences to establish coherence between sentences, as in:

> *The woman held the little girl's hand. Her daughter was only three years old. Inference: The woman held her daughter's hand.*

Because RHD patients were less able than controls to reject false inferences, MacDonald & Wales argued that the patients' deficit was in retrieving information rather than in making inferences.

Confabulations in the spontaneous speech of RHD patients also suggest a preserved ability to draw coherence inferences, albeit not the contextually relevant ones. In fact, data indicate patients with RHD draw coherence inferences when normal individuals would not. For example, given a story about a boss who catches his employee sleeping on the job, and an incongruous sentence, "He offered him a salary increase," an RHD patient might justify the incongruous sentence by saying that the

raise was related to an increase in the cost of living (Wapner, Hamby, & Gardner, 1981). Wapner et al. stress the qualitative difference in the behavior of RHD patients from the behavior of the normal individuals and LHD aphasics who tended either to ignore such incongruous statements or normalize them in their retellings of the story. These findings indicate RHD patients can identify incongruous elements but cannot reject them.

Accordingly, Joanette & Goulet (1987) read short narratives to RHD, LHD, and normal participants, and compared their performance on a verification task for information that was either presented explicitly, presented implicitly, or presented without any plausible information at all. The three groups performed similarly on the first two tasks, but RHD patients had an increased tendency to accept wrongly the plausible statements. Consequently, Joanette & Goulet argued that RHD patients have a preserved ability to draw inferences, but experience difficulty when contextual factors suggest competing interpretations.

The literature from neuropsychology, then, seems to suggest that frame-shifting necessary to understand jokes is a specialized process that recruits the right hemisphere. But while commonalities in the literature do suggest some sort of a role for the right hemisphere in frame-shifting, they do not point to a special frame-shifting region in this portion of the cortex. First, there was no evidence of double dissociation between frame-shifting and more typical meaning construction involved in sentential integration. While the behavior of LHD patients differed from that of RHD patients, neither showed a preserved ability to perform the semantic reanalysis necessary for frame-shifting. Further, neither of the groups of patients studied – those with left hemisphere and those with right hemisphere damage – are homogenous populations. Among both groups of patients studied, the locus and extent of damage vary greatly, as does the precise characterization of their deficits.

The matter is further complicated by the brain's propensity to reorganize to compensate for deficits caused by brain damage. Consequently, it is always speculative to employ the character of language deficits in brain-damaged individuals to infer the role damaged regions play in normally functioning individuals. A full understanding of this latter issue requires that we supplement data from neuropsychology with information about the role of different brain regions in healthy adults with normal language capacity. Moreover, because traditional distinctions suggest pragmatic processing occurs after the computation of literal meaning, addressing the relationship between frame-shifting and sentential integration requires a measure of brain function that can

address the relative timing of these processes. Accordingly, the next section considers evidence from one technique that is particularly sensitive to the timing of neural events – event-related brain potentials.

4.2 EVENT-RELATED BRAIN POTENTIALS

Event-related brain potentials (ERPs) are one measure of ongoing brain activity currently used by psycholinguists. The physical basis of the ERP signal is the fact that when large groups of neurons (on the order of tens of thousands) fire synchronously, they create an electrical field in the brain that can be detected with electrodes at the scalp. ERPs are patterned voltage changes in the ongoing electroencephalogram (EEG) that are time-locked to particular sorts of processing events. The ERP is obtained by recording participants' EEG and averaging across events within experimental categories. Because the averaging process presumably cancels out the EEG that is not related to the experimenter's categories, the remaining signal represents the brain activity caused by cognitive events associated with the processing of the various categories of experimental stimuli. By comparing the ERPs to different sorts of stimuli, the researcher can assess how changing the nature of the cognitive task modulates the brain response.

For example, in a classic ERP language experiment, Kutas & Hillyard (1980) contrasted ERPs elicited by visually presented sentences that ended congruously with ERPs elicited by sentences that ended incongruously.

Best Completion
(1) I take my coffee with cream and sugar.
Semantic Anomaly
(2) I take my coffee with cream and socks.

Because eye movements necessary for normal reading produce artifacts in the EEG, ERP reading experiments typically involve presenting sentences one word at a time in the center of a computer monitor. EEG can thus be time-locked to the onset of each word on the monitor, and the resultant ERP represents brain activity associated with reading. In their experiment, Kutas & Hillyard averaged EEG elicited by the last word of all of the sentences like (1) to yield the ERPs elicited by Best Completions; similarly, EEG elicited by the last word of all of the sentences like (2) was averaged to yield ERPs elicited by Semantic Anomaly.

The ERP is a waveform with a series of positive and negative peaks (often called *components*) that can be correlated with various types of motor, sensory, and cognitive activity. Components are generally labeled by reference to their polarity (*P* for positive-going and *N* for negative-going) and when they occur relative to the onset of the stimulus event or relative to other ERP components. However, an important part of the definition of an ERP component involves its *scalp distribution*, or relative amplitude over different regions of the scalp. Differences in the waveforms elicited by the two sorts of events suggest differences in the underlying brain processes. Further, ERPs can differentiate between processing differences that are quantitative from those that are qualitative. *Quantitative* differences are indexed by ERPs that have the same polarity, morphology (or wave shape), and scalp distribution, but differ in amplitude. *Qualitative* differences are indexed by ERPs that differ in polarity, morphology, and scalp distribution.

Offhand, the finding that electrical signals recorded from the scalp are a manifestation of ongoing cognitive events strikes many people as counterintuitive. However, when we consider that the brain is an electrochemical organ whose main means of information processing involves neural firing, the relationship between mental activity and the electrical activity in the EEG signal appears more straightforward. Moreover, over thirty years of research have revealed reliable relationships between the nature of various stimulus and task manipulations designed to alter participants' cognitive state, and the corresponding modulations of ERP components. For example, the P1 component is the first positive deflection in the ERP elicited by visually presented words. This component, evident 70–100 milliseconds after the word is shown (or post-word onset), reflects early sensory and vision-related attentional processing (Hillyard & Picton, 1987, for review). Mangun, Hillyard, & Luck (1993) have proposed that the P1 component reflects a gating mechanism responsible for modulating the width of the attentional spotlight. P1, N1, and P2 components elicited during reading probably reflect the visual feature extraction necessary to relate the visual stimulus to information in memory (Kutas & King, 1996).

However, most ERP studies of language processing have focused on longer latency components. In particular, the N400 component, a negative-going component that peaks at approximately 400 milliseconds after stimulus onset, has proven particularly useful as an on-line index of language processing. The N400 was first detected in an experiment that compared the ERPs elicited by the last word of

sentences like examples (1) and (2) (Kutas & Hillyard, 1980). While ERPs elicited by Semantic Anomalies like "socks" contained a negative-going wave that peaked about 400 milliseconds post-onset, Best Completions like "sugar" elicited a positive-going wave in the same time window. The N400, the monophasic negativity elicited by "socks," begins about 200 milliseconds after the onset of a word or pseudoword, peaks approximately 400 milliseconds post-onset, and continues for another 100 milliseconds. Subsequent research indicates that N400 is elicited by all words, spoken, read, or signed. Its amplitude varies chiefly as a function of semantic factors, though it is modulated by nonsemantic factors as well.

In short, ERPs provide a continuous on-line index of the processing that occurs upon presentation of a linguistic event. Because they are multidimensional, ERPs are well suited for addressing questions regarding what sorts of information people are sensitive to and when.

4.2.1 Lexical Versus "Post"-Lexical Processing

Traditionally, psycholinguists have divided processing into so-called lexical and postlexical stages of processing. For a word to be processed, first it must be recognized – that is, matched to a template stored in memory – and then accessed – that is, information about the word must be made available for later, postlexical stages of processing. One relic of this explanatory framework is the distinction between lexical priming thought to result from activation within a fixed lexicon, and sentential priming that arises from integrative processes that take place only after the word has been accessed. In Chapter 3, we saw that much of the behavioral evidence for lexical priming has been undermined. Findings in the ERP language literature also point to the inadequacy of traditional distinctions between lexical and postlexical stages of processing.

Because the N400 was discovered in an experiment comparing congruous sentences with incongruous sentences, many people think it is a marker of semantic anomaly. However, subsequent research indicates N400 is elicited by all words, congruous and incongruous, regardless of whether they appear at the beginning, in the middle, or at the end of a sentence (see Kutas & Van Petten, 1994). Generally, the best predictor of N400 amplitude for a word in a given sentence is its cloze probability. N400 amplitude is largest for items (such as semantic anomaly) with low cloze probability, and smallest for items (such as best completions)

with high cloze probability. Semantic anomaly thus represents the end point on a continuum of expectedness in a particular context.

The issue of whether the N400 indexes the violation of expectancy or the congruence of the target with the preceding context was addressed by Kutas & Hillyard (1984). To tease apart the possibilities, these investigators crossed several levels of *contextual constraint* – the degree to which the preceding context constrains the possible completions of a sentence – and *cloze probability* – the probability that a given word will actually be generated in that context. Kutas & Hillyard found that N400 amplitude was the same for low cloze items in high-constraint contexts as it was for low cloze items in low-constraint contexts. So, when cloze probability is matched, N400 amplitude is unaffected by contextual constraint. Thus N400 amplitude is an inverse indicator of the congruence of the actual target with the preceding context rather than a direct indicator of disparity between actual and predicted targets. In general, the amplitude of the N400 component is proportional to the difficulty of integrating a given word into established context and can be used as a reliable indicator of semantic priming (see Van Petten & Kutas, 1991, for review).

Besides its sensitivity to factors associated with sentential integration, N400 is also sensitive to factors traditionally associated with putatively lexical aspects of processing. For example, while pseudowords such as "blick" elicit very large N400s, nonwords such as "blckhg" do not elicit N400 at all. This suggests that the process or processes that generate N400 are sensitive to regularities of orthography and phonology, yet unable to filter out pseudowords. Moreover, a word's overall frequency in the language has been seen to modulate N400 amplitude in both word lists and in sentences (Smith & Halgren, 1987). Just as people respond faster to high- than low-frequency words in naming and lexical decision paradigms, N400 amplitude is smaller for high- than for low-frequency words.

Using the known correlation between the amplitude of the N400 component and semantic expectancy, Van Petten & Kutas (1990) compared the effects of word frequency and ordinal position in sentences to see whether the two factors were additive or interactive. The assumptions of the additive model are that if two factors (that is, processes) are independent of one another, they will each influence the response measure autonomously. While the applicability of this model to behavioral data can be questionable, physical characteristics of the ERP signal (Helmholz's law of superposition) ensure that electrical fields generated in different brain regions will summate where they intersect. Consequently, if

discrete lexical and postlexical processes underlie the respective frequency and position effects on N400, their influence on N400 amplitude should be additive.

In fact, Van Petten & Kutas (1990) found main effects of frequency and position, and an interaction between the two. Low-frequency words elicited larger N400 components than did high-frequency words. Moreover, when Van Petten & Kutas examined data from the sentence intermediate words more closely, they found a downward linear trend in N400 amplitude. However, the two effects were nonadditive, as the frequency effect was most prominent in words in the beginning of sentences, somewhat attenuated in sentence-intermediate words, and entirely absent from sentence-final words. Van Petten & Kutas argued that this interaction, especially the fact that sentence-final words exhibit no N400 frequency effect whatsoever, suggests the buildup of contextual knowledge influences the initial activation and retrieval of lexical items, as well as integrative processes typically referred to as postlexical.

Given N400 sensitivity to variables attributed to both lexical and postlexical processes, one might question whether this distinction is a valid one to make. Besides frequency effects, lexical processing is thought to be characterized by priming effects resulting from spreading activation between semantically related items in the lexicon. Analogously, postlexical processing is thought to underlie sentential priming effects in which contextually congruent words are integrated more easily into the developing representation.* As noted previously, because they are multidimensional signals that index brain activity, ERPs can potentially distinguish between the operation of qualitatively different processes and the differential operation of a single underlying process.

Accordingly, to determine whether lexical and postlexical priming effects were the manifestation of qualitatively distinct processes, Van Petten (1993) compared lexical and sentential context effects by presenting the same associated pairs (for instance "moon" and "stars") in both a congruent (3) and an anomalous (4) sentence. Besides varying the type of sentential context, Van Petten also varied the lexical context by including *unassociated* word pairs (for instance, "insurance" and "refused") in congruent and anomalous sentential contexts, as in (5) and (6).

* As noted in Chapter 3, this category tends to be a dumping ground for effects that don't fit under the rubric of lexical priming. However, few models have been proposed to explain how these postlexical priming effects might be related to comprehension processes besides the (relatively impoverished) demands of the lexical decision task.

Congruous Sentence/Associated Words
(3) When the (moon) is full it is hard to see many *stars* of the Milky Way.
Anomalous Sentence/Associated Words
(4) When the (moon) is rusted it is available to buy many *stars* of the Santa Ana.
Congruous Sentence/Unassociated Words
(5) When the (insurance) investigators found out that he'd been drinking they *refused* to pay the claim.
Anomalous Sentence/Unassociated Words
(6) When the (insurance) supplies explained that he'd been complaining they *refused* to speak to the keys.

Participants' ERPs were recorded as they read sentences like (3)–(6) to compare the effects of lexical and sentential context. In the Congruous Associated condition, the target word could benefit from both lexical and sentential context alone; in the Anomalous Associated context, the target could benefit only from lexical context; in the Congruous Unassociated context, the target could benefit only from sentential context; and, in the Anomalous Unassociated condition, the target could benefit from neither.

Van Petten (1993) found that ERP effects elicited by manipulation of lexical and sentential context were strikingly similar. Observed similarities in polarity, wave shape, and distribution of ERPs elicited by target words in varying sentence contexts indicate the two sorts of context effects are produced by a similar underlying mechanism. Although the effect of sentential context lasted longer than the effect of lexical context, the two effects *began* at the same point in time. Van Petten's findings clearly argue against a lexical stage of processing that is distinct from processes that underlie integrative effects.

In a similarly motivated experiment, Kutas (1993) compared lexical and sentential priming by using sentences and word pairs. During the first half of the experiment, ERPs were recorded as the participants read sentences for comprehension. Half of the sentences were like (7) and ended with the most predictable response. The other half, like (8), ended with a word of low cloze probability.

Best Completion Sentences
(7) Before exercising Jake always stretches his muscles.
Low Cloze Probability Sentences
(8) Fred put the worm on the table.

In the second half of the experiment, ERPs were recorded as participants read word pairs derived from the sentences. The second (target) word in every word pair was a sentence-final word from the sentential stimuli. Moreover, the first word (the prime) in each word pair was designed to be similar in semantic content to the sentence fragment that preceded each target. So, the prime derived from the best completion sentences was a semantic associate of the sentence-final word. The word pair derived from (7), for example, was BICEPS–MUSCLES. The prime derived from the low cloze probability sentences was the best completion for that sentence. Thus, the word pair derived from sentence (8) was HOOK–TABLE.

This design allowed Kutas to compare ERPs to the same word in roughly similar lexical and sentential contexts. Assuming that qualitatively different mechanisms are indexed by qualitative differences in the ERP, we can use the ERP measure to reveal the operation of different mechanisms in contexts that require higher-level processing and those in which only lexical processing is needed. If lexical and sentential context effects are due to discrete processes, there should be a qualitative difference in the ERPs elicited by stimuli in the different context conditions. However, if lexical and sentential context effects are due to the same underlying process, one would expect only quantitative differences in the amplitude and timing of the two effects.

In fact, Kutas (1993) found that ERPs elicited by sentence-final words were qualitatively similar to those elicited by the second word in each word pair. Further, the priming effect was larger and began earlier in the *sentence* condition than in the word pairs. Nonetheless, the peak latencies of the two N400s were nearly identical, as were their scalp distributions. Because these differences were quantitative, they undermine a qualitative distinction between lexical and sentential processes. Citing similarities in the wave shape, scalp distribution, and experimental modulation of the N400 priming effect in lexical contexts as in sentential contexts, Kutas argues for the existence of a single process underlying priming effects in both.

4.2.2 What's in a Word?

Curiously, while 90% of the variance in N400 amplitude can be explained by cloze probability, the two factors dissociate when it comes to items that are completely unexpected. Kutas & Hillyard (1984) compared ERPs elicited by high cloze endings for sentences such as (9) with ERPs elicited

by two sorts of anomalous endings: *related anomalies*, such as "drink," which were semantically related to the best completion, and *unrelated anomalies*, incongruous endings with no semantic or associative relationship to the best completion.

(9) The pizza was too hot to . . .
Best Completion: eat.
Related Anomaly: drink.
Unrelated Anomaly: cry.

Related anomalies elicited N400 components whose amplitude was *greater* than that elicited by the best completions, but *less* than their cloze-matched controls. Kutas & Hillyard (1984) interpreted the difference in the amplitude of the N400 component elicited by related and unrelated anomalies as reflecting the spread of activation between related word nodes in a semantic network. They argued that the related anomaly effect suggests the context-sensitive process or processes that generate N400 are also sensitive to relationships that obtain between words.

Of course, a mechanism sensitive to the buildup of semantic context *and* to contextually irrelevant interword relationships does not fit well with existing models of language processing. Until relatively recently, it had been assumed that semantic priming reflects the organization of word representations in the lexicon. Assuming that related words are often encountered in the same sentences, if those words were stored "close" to one another in the lexicon, spreading activation might serve to prime the entries for words the speaker was likely to encounter. The functional utility of semantic priming, then, was to bias the lexical access of related terms and thus facilitate processing in sentential contexts.

While it is not necessary for every psychological mechanism to have an obvious utility value, sometimes it pays to reevaluate the logic behind our assumptions. In the case of semantic priming, this includes two assumptions. First is the idea that priming reflects the organization of lexical items in the lexicon. Second is the assumption that preactivating a word's representation in the lexicon facilitates subsequent processing. However, when we consider that the demands of meaning construction go far beyond the activation of an entry in a mental dictionary, the utility of preactivating such a lexical representation is highly questionable.

Moreover, in spite of its sensitivity to some lexical variables (such as word frequency and lexical associations), it would seem that the process underlying N400 priming effects is primarily aimed at sentential integration. The ERP literature reviewed in this section suggests that

sentential priming effects are more robust than those of lexical priming. Further, while ERP effects that index sentential priming usually begin at about the same time as those that index lexical priming, Kutas (1993) found that sentential priming effects actually began earlier than the N400 priming effect for word pairs.

Perhaps semantic priming effects are not to be interpreted as a function of lexical access, but rather as a fast-acting integrative process. While the priming effect indexed by the N400 is not automatic, it does appear to act obligatorily: N400 is elicited by all words, irrespective of whether they appear in sentences or in word lists. The N400 component might index a process that takes word forms as one input and working memory contents as another, and initiates a search for relevant semantic information for inclusion in the current message-level representation. Rather than reflecting semantic relationships between word nodes in a semantic network, the related anomaly effect might reflect constraints on slots that help the processor flesh out the message-level representation.

Moreover, the frame-based approach to meaning construction developed in this section suggests an alternative interpretation of the related anomaly effect described earlier in this section. Once a frame has been activated, words associated with that frame are easy to integrate into the developing representation of meaning. Because "eat" and "drink" are associated with the same frame, it's easier to integrate information associated with "drink" than with "cry." The related anomaly effect, then, might reflect dynamic properties of meaning construction rather than spreading activation in a fixed network.

Further, these ideas are supported by the report that N400 is also sensitive to global contextual factors. St. George, Mannes, & Hoffman (1994) report N400 amplitude is sensitive to global context provided by a title. St. George and colleagues recorded ERPs as participants read ambiguous paragraphs that either were, or were not, preceded by a disambiguating title. Although the local contextual clues provided by the paragraphs were identical in the titled and untitled conditions, words in untitled paragraphs elicited greater N400 amplitude. These data indicate that frames evoked by the title facilitated sentential integration on-line as indexed by N400 amplitude.

4.2.3 ERP Indices of Frame-Shifting

The ERP literature reviewed thus far suggests that even the initial stages of semantic processing are integrative in nature. However, what of the

relationship between frame-shifting in jokes and the more general processes of sentential integration?

In Chapter 3 we saw that readers spend more time on passages, sentences, and even individual words that prompt frame-shifting (Coulson & Kutas, 1998; Rumelhart, 1980; Sanford & Garrod, 1981). These findings indicated that words that trigger frame-shifting present the processor with a challenge that differs from that presented by sentences that don't require revision of high-level expectations. But because reading time indexes the end stage of processing, some important questions remain unanswered. For instance, at what stage of processing is frame-shifting initiated? On traditional accounts, frame-shifting should occur as part of a pragmatic processing stage initiated after the computation of literal meaning. However, the approach sketched here argues against a late-acting stage of pragmatic processing that differs qualitatively from basic processes of sentential integration.

To test the relationship of the processes underlying frame-shifting and the integration of other sorts of unexpected words, Coulson (1997) recorded ERPs to sentences that ended either as jokes or with equally unexpected nonjoke endings designed to be congruent with the contextually activated frame. Stimuli were taken from those reported by Coulson & Kutas (1998) in the reading-time study of one-line jokes discussed in Chapter 3, and included sentences such as the following:

Low-Constraint Stimuli
The city planners built the train station on the edge of town so that it would be near the
Nonjoke Ending: planes.
Joke Ending: tracks.
High-Constraint Stimuli
She read so much about the bad effects of smoking she decided she'd have to give up the
Nonjoke Ending: habit.
Joke Ending: reading.

If frame-based activation of background knowledge is a driving factor in meaning construction, ERPs elicited by jokes should differ from nonjokes that do not trigger frame-shifting. Moreover, the interesting question concerns the relationship between ERPs elicited by jokes and their nonjoke controls. A qualitative difference would suggest that fundamentally different processes underlie the processing of words that trigger frame-shifting. In contrast, a quantitative difference would

argue for similar mechanisms underlying the integration of both sorts of endings.

Coulson (1997) found that joke endings elicited greater amplitude N400 than did nonjoke endings. Although the preceding sentence fragments were identical, and both sorts of ending types consisted of lexical items that were unexpected in those sentence contexts, the joke endings elicited greater amplitude N400 (see figure 4.1). This finding indicates that joke endings were more difficult to integrate than nonjoke endings. Moreover, as in the reading-time study described in Chapter 3, the joke effect was far more pronounced in high-constraint sentences that suggest a particular sentence-final word than in the more open-ended contexts provided by the low-constraint stimuli (see figure 4.2).

However, one difference between the ERP study and the earlier reading-time study involved performance on the comprehension probes. While all participants in the self-paced reading study scored well on comprehension probes, there was much more variability in the comprehension scores of participants in the ERP experiment. Although reading comprehension was quite good for filler materials and nonjoke stimuli, not all of the participants "got" the jokes. Consequently, participants were categorized as good and poor joke comprehenders on the basis of their performance. As can be seen in figure 4.3, good joke comprehenders generated joke effects for both high- and low-constraint sentences, while poor joke comprehenders generated smaller joke effects that were largely confined to high-constraint sentences.

These data suggest good comprehenders constructed message-level representations that could benefit from the fact that nonjoke endings, while unexpected words, were congruent with the contextually evoked frame. In contrast, poor comprehenders seemed less able to capitalize on contextual information in order to commit to a frame at the appropriate level of specificity. Perhaps the poor comprehenders were unable (or unwilling) to capitalize on clues in the sentences to construct a rich message-level representation. Overall, however, differences in the two groups' ERPs support a functional connection between N400 modulation and successful frame-shifting: The presence of joke effects was dependent on participants' successful comprehension of the jokes.

While N400 amplitude indexes the difficulty of integration, it does not differentiate between words that prompt the reader to construct a new frame, and words that prompt the refinement of frames already evoked. Although the former tend to elicit greater amplitude N400, Coulson's (1997) results suggest that the underlying neural processes operate

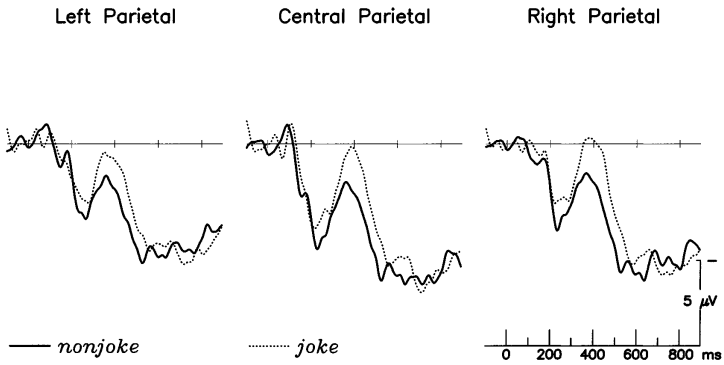

Figure 4.1 Joke Effect. Grand average of all participants' ERPs elicited by jokes (dotted) that required frame-shifting versus nonjokes (solid) that were unexpected but congruent with the contextually evoked frame.

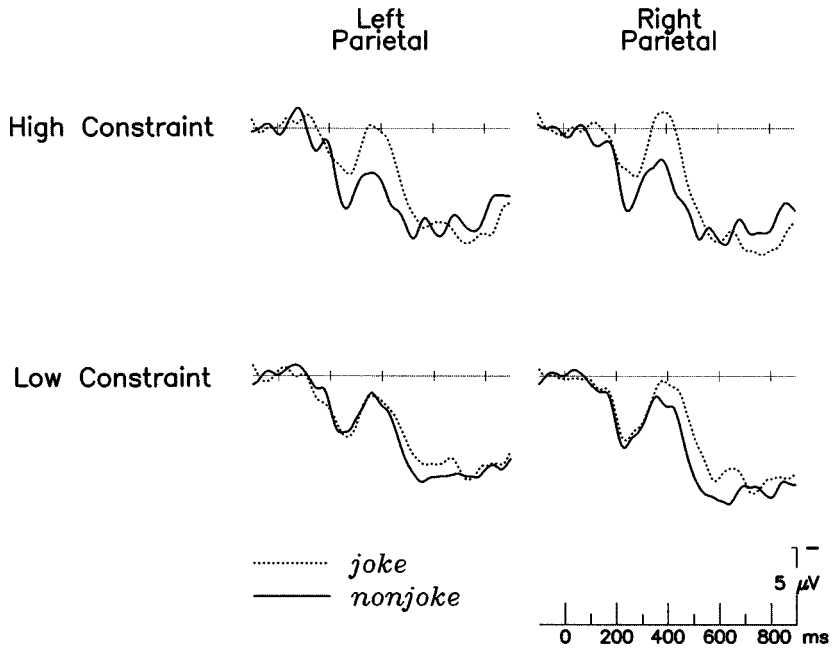

Figure 4.2 Ending Type x Contextual Constraint Interaction. The grand average of all participants of ERPs elicited by the last word of Low versus High Constraint fragments at one left and one right parietal scalp site. At the right parietal site, jokes were more negative than nonjokes in both High and Low Constraint sentences. However, the effect is much larger in High Constraint sentences.

obligatorily in both kinds of meaning construction. In fact, N400 amplitude modulation is probably best attributed to a reduction in N400 elicited by nonjoke stimuli rather than increased N400 for jokes. Thus construed, these data seem to argue against a *cost* associated with frame-shifting, pointing more toward a facilitative effect for words consistent with the contextually evoked frame. Rather than revealing frame-shifting as a special process, these data suggest that something like frame-shifting happens more generally.

Sustained Positivity

Besides detecting the N400, Coulson (1997) detected another ERP effect that was closely associated with the hypothesized frame-shifting process. Compared to joke endings that required frame-shifting, nonjoke endings elicited a sustained positivity of 500–900 milliseconds post-onset focused in left lateral anterior sites. Moreover, this effect was observed only in the ERPs generated by people who consistently got the jokes (see figure 4.4). Further, among good joke comprehenders, the

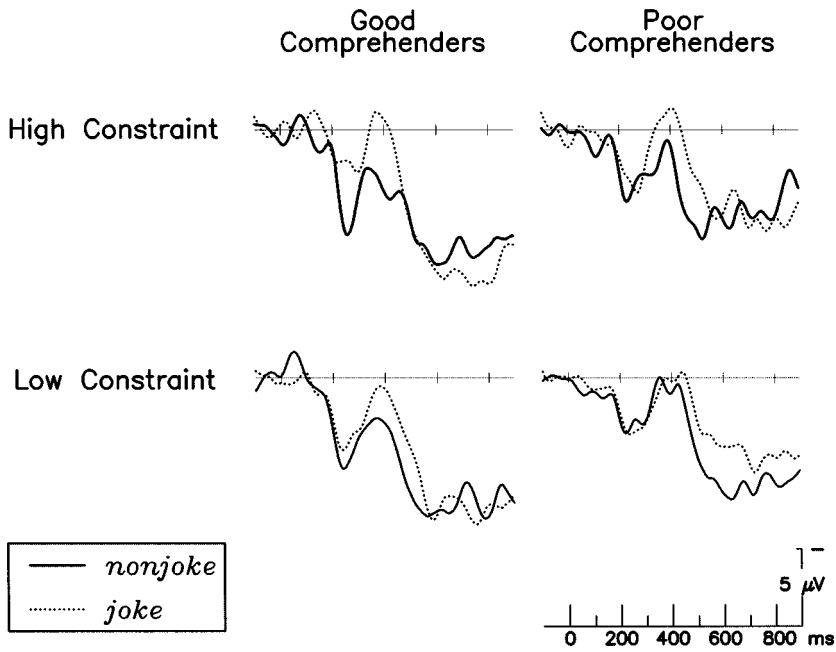

Figure 4.3 Joke Effects in Good and Poor Comprehenders. ERPs from the vertex electrode elicited by Joke versus Nonjoke Endings in High and Low Constraint contexts.

sustained positivity was evident for both high- and low-constraint stimuli. Although localization of ERP effects is inherently difficult, the extremely focal nature of this effect is at least consistent with a neural generator in the dorsolateral prefrontal cortex.

Robin & Holyoak (1995) have proposed that the prefrontal cortex is important for the creation and maintenance of explicit relational representations of the sort required for the space structuring model sketched in Chapter 3. Robin & Holyoak argue that the prefrontal cortex subserves three functions invoked in processing relational representations. The first, maintaining representations of elements in working memory, is thought to be subserved by the dorsolateral prefrontal cortex. A second function is learning conditional contingencies. This operation is thought to involve the periarcuate prefrontal cortex. Finally, attentional control is necessary to prevent cross-talk between each element and its role assignments and for the suppression of task-irrelevant aspects of the situation. The orbitofrontal cortex is thought to be important for inhibiting responses that are irrelevant to the task. One possibility is that the sustained positive shift indexes prefrontal activation implicated in the operation of verbal working memory.

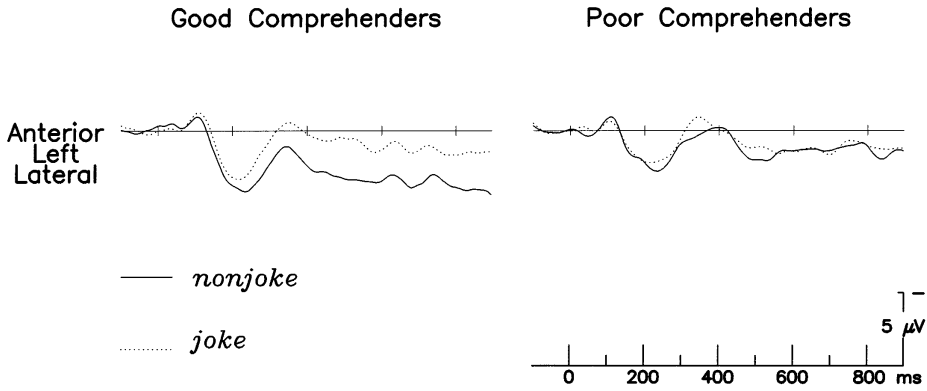

Figure 4.4 Sustained Positive Shift in Good and Poor Comprehenders. This shift is evident from the grand average ERP elicited by Joke and Nonjoke Endings at one anterior lateral left hemisphere scalp site.

In an experiment in which participants read simple transitive clauses in a paradigm similar to that employed by Coulson (1997), Kutas & King (1996) report a slow-rising cumulative positive drift at anterior recording sites, particularly over the left side of the scalp, that was more noticeable in ERPs generated by good than poor comprehenders. Kutas & King suggest the time course of the slow positive drift may reflect dopaminergic neuromodulatory activity associated with working memory processes in monkeys (Sawaguchi & Goldman-Rakic, 1991). They interpreted this positivity as reflecting integration between items in working memory with information in long-term memory to form a mental model as in the structuring of a single mental space.

Similarly, King & Kutas (1995) found differences in the amplitude of the slow positive shift in a comparison of ERPs elicited by subject- and object-relative clauses. The anterior shift grew most quickly for the simplest sentences (simple transitive clauses), least quickly for the most difficult object-relative sentences, with subject-relative sentences of intermediate difficulty falling in between these two extremes. Further, King & Kutas did not observe the positive drift in poor comprehenders' ERPs, supporting the interpretation of this positivity as an index of *successful* message-level processing. Data are thus consistent with King & Kutas' hypothesis that the positive drift is associated with the buildup of a coherent representation at the message level.

4.2.4 Summary

It seems that successful frame-shifting is more difficult than the mere integration of an unexpected word congruent with the contextually evoked frame. Results suggest that frame-based retrieval of background knowledge plays an important role in general processes of sentential integration. However, results argue against the idea that there is a special role for the right hemisphere in frame-shifting. If anything, the sustained positivity over the left anterior scalp points to an important role for brain areas in left prefrontal cortex in frame-shifting. However, as the strength of the ERP technique is temporal rather than spatial resolution, localization of function is currently unknown. Nonetheless, with the increasing use of neuroimaging techniques such as positron emission tomography (PET) and functional magnetic resonance imaging (FMRI), this issue could be resolved within the next decade. But based purely on the ERP literature to date, frame-shifting in jokes appears to differ only *quantitatively* from other sorts of sentential integration.

PART TWO

CONCEPTUAL BLENDING

In Part II, we will see the importance of *conceptual blending* for the integration of knowledge structures and the development of new concepts. Conceptual blending is a set of noncompositional processes in which the imaginative capacities of meaning construction are invoked to produce emergent structure (Fauconnier & Turner, 1998). Although blending is frequently employed for sophisticated feats of reasoning, its intermediate products are cognitive models whose plausibility spans the gamut from chimerical, to merely bizarre, to downright trite. Analyses in Chapters 5 through 7 show how cognitive models built in blended spaces can yield productive inferences in spite of, and, sometimes even *because of*, their strange properties.

II.1 TRASHCAN BASKETBALL

Imagine a scenario in which two college students are up late studying for an exam. Suddenly one crumples up a piece of paper and heaves it at the wastepaper basket. As the two begin to shoot the "ball" at the "basket," the game of trashcan basketball is born. Because it involves the integration of knowledge structures from different domains, trashcan basketball can be seen as the product of conceptual blending. In conceptual blending, frames from established domains (known as *inputs*) are combined to yield a hybrid frame (a *blend* or *blended model*) comprised of structure from each of the inputs, as well as unique structure of its own. For example, in trashcan basketball, the input domains are trash disposal and (conventional) basketball, and the resultant blend incorporates a bit of both domains. Moreover, emergent structure – that is, properties of trashcan basketball that differ from properties of the input

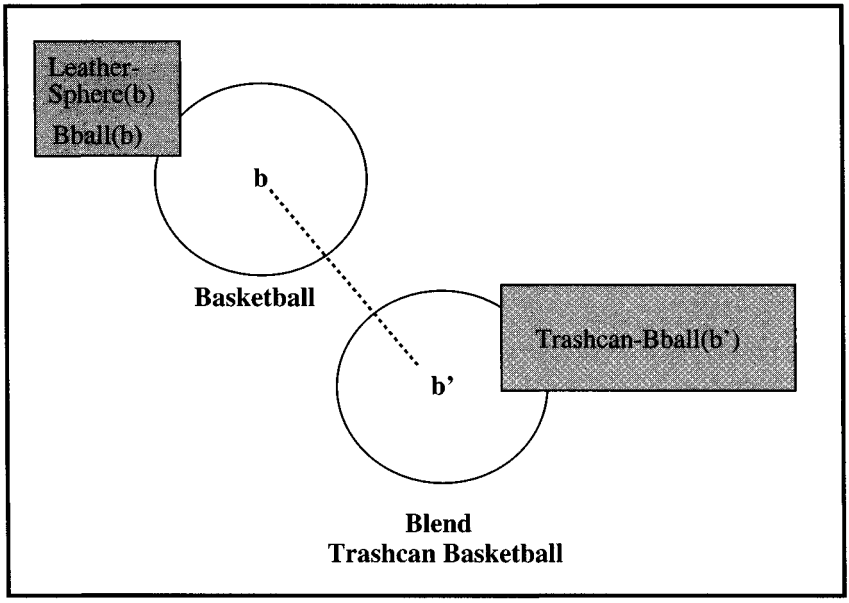

Figure II.1 Analogical Link Between Conventional and Trashcan Basketballs.

domains – need not be explicitly computed, but arise from affordances in the environment.

In trashcan basketball, the elements are inherited from the Disposal domain (that is, you play trashcan basketball with paper and a trashcan); however, they are *understood* as elements in the Basketball domain. One way of representing the relationship between the balls used in the games of conventional and trashcan basketball is to employ the framework of mental spaces. This involves using a different mental space to represent each type of basketball game, and linked elements to represent each of the balls. This sort of a representation is meant to capture the sense in which our two college students understand the analogical relationship between the balls in conventional and trashcan basketball (see figure II.1). Moreover, mental space theory can represent the fact that our trashcan basketball can also be understood to be a crumpled-up piece of paper. Again, mental spaces are used to represent each of the domains, while linked elements represent the identity of the ball in trashcan basketball with the crumpled paper in the disposal domain (see figure II.2).

In fact, mental spaces prove to be quite useful in representing a number of aspects of conceptual blending. For example, one feature is the

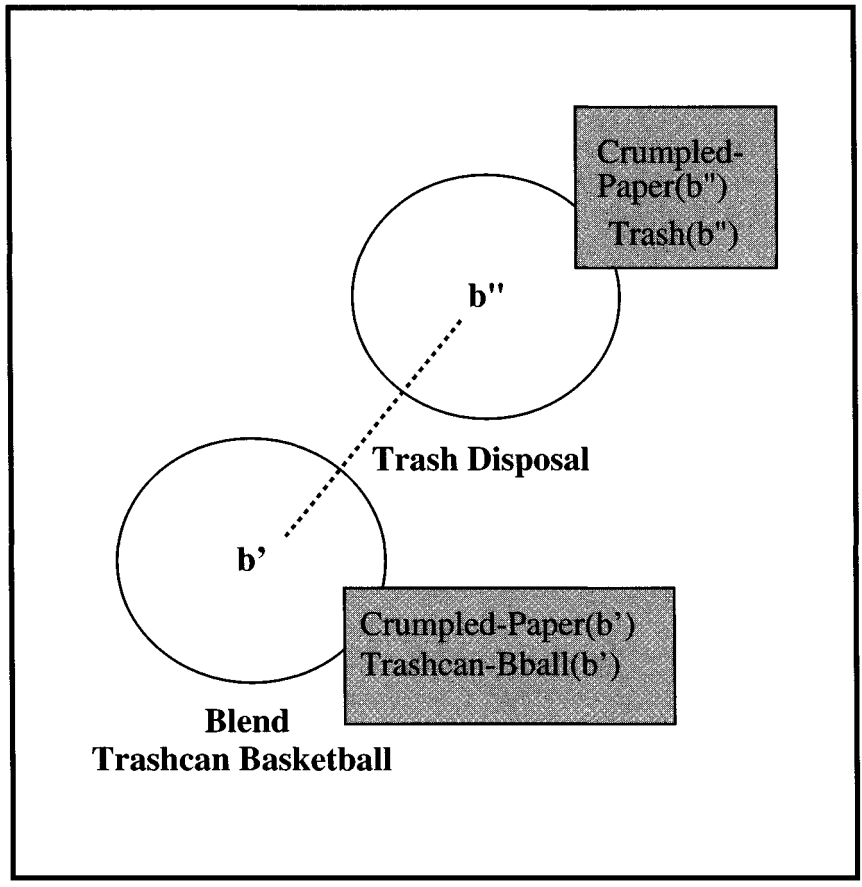

Figure II.2 Identity Link Between Trashcan Basketball and Crumpled Paper.

partial nature of their representations, which consist of elements, their attributes, and higher-order relations between the elements in a space. Moreover, while these representations are quite minimal, their elaboration is constrained by a substantial body of background knowledge. This, coupled with the way in which information evoked in one mental space can spread to others, lends these representations a flexibility that exceeds that of traditional frame-based representations. Moreover, the Access Principle, which licenses the use of a term from one space to refer to its counterparts in linked spaces, makes mental spaces a good framework for discussing metaphoric language, as we will see in Chapter 6. Indeed, mental space theory is suitable for discussing cross-space mappings based on all sorts of pragmatic functions, and suggests

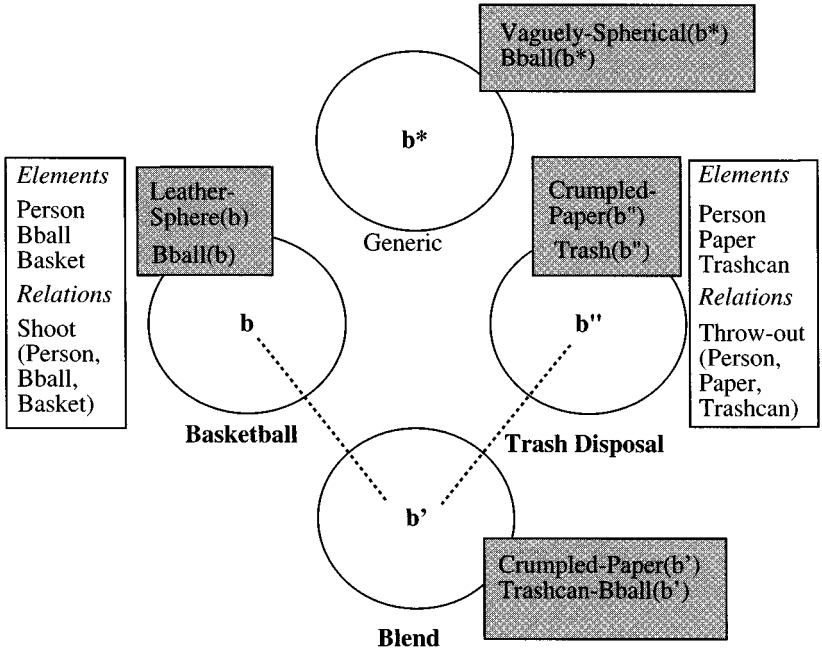

Figure II.3 Conceptual Integration Network for Trashcan Basketball.

commonalities in the comprehension of metaphors and counterfactual constructions, as discussed in Chapter 7.

II.2 BLENDING AND INTEGRATION NETWORKS

Because a proper understanding of the game of trashcan basketball involves the simultaneous apprehension of the relationship between knowledge recruited from established domains as well as aspects of the novel concept, Fauconnier & Turner (1994;1998) have argued that conceptual combination occurs in *conceptual integration networks*. A conceptual integration network is a network of mental spaces structured with frames that the speaker constructs from contextual information and background knowledge. The prototypical integration network is comprised of four mental spaces: one for each of the input domains, one for the blended domain, and a *generic* space that represents abstract properties that apply to structure in all the spaces. For example, figure II.3 represents the conceptual integration network for the concept of trashcan basketball. The generic space helps establish mappings between the inputs that enable integration to proceed. Further, people

often appeal to this sort of abstract, schematic information when extending their concepts.

Because simple conceptual blends often rely on a more general system of mappings between the input domains (either because they are conventionally established or because speakers in particular contexts are able to generate a novel set of mappings), analysis of a single blended model frequently suggests a systematic set of cross-domain mappings within which a particular blended model functions. For example, in trashcan basketball we can speak of conceptual blending that contributes to the concept of the ball in trashcan basketball. Moreover, we can speak of a more general blend between the domains of trash disposal and the game of basketball – a blend that alludes to the whole family of concepts that can be constructed by processes of conceptual integration.

To capture the existence of extensive counterpart structure in the conceptual blends discussed in the following chapters, in Part II I employ a slightly different notation for mental spaces and their representational structure. Rather than using circles with attached frames, I have employed tables that represent distinct cognitive domains in either a single column or (occasionally) a single row. Table II.1 shows an example. The content of these tables is intended to be illustrative of structure in each domain, alluded to by a list of contextually relevant elements and relations. Of course, neither the tables nor the circles are appropriate for a real-time model of on-line meaning construction processes. Because the theories of mental spaces and conceptual blending are intended to capture informational aspects of these conceptual combination processes, they do not generate fine-grained predictions about their neural implementation. However, see Shastri & Ajjanagadde (1993) for a modeling framework potentially compatible with this approach.

Table II.1: Mabel is the Daughter of Paul

Input1	**Input2**	**Blend**
Elements	**Elements**	**Elements**
Mabel	[Daughter]	Mabel
Paul	[Father]	Paul
	Relations	**Relations**
	Daughter-of (Daughter, Father)	Daughter-of (Mabel, Paul)

II.2.1 Single Framing Networks

The simplest kind of integration network is the single framing network, a network comprised of two input spaces that contain the information to be combined, and a blended space that incorporates information from both of the inputs. In the single framing network, one input has an abstract frame with no fillers, while the other has elements with no organizing frame. For example, the network for representing (1) would involve one input space with an abstract frame with no fillers, while the other input contained elements for Mabel and Paul. Information from each of the input spaces is combined in the blend with no accommodation at the frame level.

(1) Mabel is the daughter of Paul.

Emergent properties of blending are represented in the network and involve frames at various levels of specificity. The term *frame-level* structure is used to refer to attributes and relations that have been attentionally focused, while *specific-level* structure refers to the values of those attributes. Moreover, because of the hierarchical organization of frames, the expansion of activated frames can result in the activation of more detailed *incidental-level* structure.

Integration that occurs in a single framing network is a particularly tame sort of a blend in which the only accommodation is at the specific level in which, for example, Mabel and daughter are assumed to be female and animate. These sorts of networks can also represent cases in which one input is a frame constructed from background knowledge, and the other consists of elements constructed from the perception of a specific contextual situation. As in (1), integration consists of filling the slots in the frame with the specific attributes of the situation and accommodating the attributes accordingly.

II.2.2 Frame Networks

More complex integration occurs in a *frame network*, a network of spaces that all share an organizing (frame-level) frame. As noted previously, a frame network typically consists of four mental spaces: two input spaces, a blended space, and a generic space that contains abstract information shared by the inputs. Because the generic space shares structure with both of the inputs, it defines a partial mapping between representations

in the input spaces. As in the single framing network, the blended space contains a representation constructed from structure in both of the inputs.

Processes of conceptual blending project elements and relations from each input so as to form an integrated representational structure that may differ substantially from that in the input spaces. Together, the four spaces form a conceptual integration network that can lead the language user to novel conceptualizations of actions and events. In a *one-sided* network, the blend inherits frame-level structure from one of the input spaces and specific-level structure from the other input space. Accordingly, the blend will share much of the logic of the input space that projects frame-level structure. By contrast, in a two-sided blend, both input spaces contribute frame-level structure to the blend. Consequently, the resultant logic of the two-sided blend can be different from either of the input frames.

For example, a one-sided network would be employed to represent the concept of a hook shot in trashcan basketball. The Hook Shot frame from conventional basketball is integrated in the blend with elements from the disposal domain on the basis of a systematic cross-domain mapping. As in the examples previously discussed, blending here relies on the establishment of a cross-domain mapping between the wad of paper and the basketball, and between the trashcan and the basket. In a one-sided network, one of the domains provides the structure (in this case, conventional basketball) and the other provides the elements. Although the input domains share topology and abstract structure (of putting an object into a container), the organizing frame in the blended hook shot concept comes from the conventional Hook Shot frame (the outstretched arm, the upward propulsion of the ball, and the arclike trajectory of the ball into the basket), except that instead of applying to a basketball and a basket, the frame applies to the wad of paper and the trashcan. In this case, the emergent structure of the blended hook shot (e.g., the slightly different muscles recruited, the different rate at which the trashcan basketball travels, etc.) comes from affordances of the environment.

A two-sided frame network might be employed to understand a "dunk" in a variant of trashcan basketball played by a sketch artist and his son. Imagine that the artist makes a mistake in the course of sketching and wads up the paper into a ball. Seeing his son, he asks the boy to "dunk" the ball for him. Taking the paper, the boy runs across the room

to the trashcan and forcefully throws away the crumpled-up sketch. In this case, both elements and relations in the input domains have been put into correspondence with one another. As before, the ball maps onto the paper, and the trashcan maps onto the basket. Moreover, dunking the basketball has been put into correspondence with disposing of the paper. Note that specific-level aspects of the frame for dunking have been accommodated to suit the trash disposal domain: Rather than jumping *up* to dunk (as in conventional basketball), the boy looks *down* into the wastepaper basket. Moreover, unlike the previous example, certain frame-level aspects of the blended dunk come from the trash disposal input – namely the artist's goal to dispose of the bungled sketch.

II.2.3 Blending Processes

Blending involves three processes – composition, completion, and elaboration – each of which provides for the possibility of emergent structure. *Composition* involves attributing a relation from one space to an element or elements from the other input spaces. A good example of composition can be found in the single framing network, described previously, where the elements (Mabel and Paul) are integrated with the frame for Daughter-of. Another example comes from the domain of trashcan basketball, in which the frame for Ball is applied to a crumpled-up piece of paper (originally) from the disposal domain. Although emergent structure in composition appears to be due to the mere juxtaposition of predicates and arguments from different input domains, it actually arises from contextual accommodation of a predicate from one domain to apply to elements from a different domain.

Completion is pattern completion that occurs when structure projected from the inputs matches information in long-term memory. For example, in trashcan basketball, if one student shoots and the other attempts to defend the goal, pattern completion could result in evoking the frame for a one-on-one basketball game. Completion is closely related to *elaboration*, a process that involves performance and/or mental simulation of the event in the blend and is constrained by the logic (or illogic) of the blended domain itself. For example, one might employ elaboration to understand the concept of moon rock basketball – basketball played on the moon with moon rocks.

Formally, elaboration involves the same sort of conceptual integration as completion, evoking novel structure in the blend in response to the coalescence of information. The activation of novel structure

either can be done by mental simulation, or, as in the trashcan basketball examples, may rely extensively on interaction with the environment as construed with existent blended models. For example, astronauts might elaborate moon rock basketball through action in the world. In this case, elaboration might proceed as a series of completions that result from consistently integrating concepts from the domain of basketball with their current activity. Alternatively, earth-bound students might elaborate the same concept through creative mental simulation. Elaboration, then, is more of a creative process than completion, and is potentially more cognitively taxing due to the demands of mental simulation.

II.2.4 Outline

Conceptual blending is a very general set of processes and is evident at many levels of analysis, including the morphological structure in single words (Lehrer, 1996), conceptual structure of single concepts (Turner & Fauconnier, 1995), and grammatical structure in sentences (Mandelblit, 1997), and can even be said to operate at an overarching level across conceptual domains (Fauconnier & Turner, 1994). To demonstrate the general utility of conceptual integration networks, Chapter 5 addresses the application of conceptual blending theory to some problems of noun modification. Chapter 6 outlines the blending in various metaphorical utterances, and demonstrates how both entrenched metaphor and novel analogical mapping involve conceptual integration. Chapter 7 considers blending in conditionals and counterfactuals, and shows how and why these constructions are often used rhetorically.

CHAPTER FIVE

Conceptual Blending in Modified Noun Phrases

In this chapter, we will tackle one of the central problems of cognitive science: how people combine concepts in order to yield new ones. One way that psychologists and linguists have addressed this issue in the past has been to look at the interpretation of simple noun phrases like "square peg" or "trashcan basketball." But while the appeal of the noun phrase is presumably its simplicity, the meaning construction that underlies these phrases turns out to be very complex indeed. Ironically, while it appears to be a construction – a predicate and an argument – that would present the simplest case for compositionality, the meaning of noun phrases is rarely compositional. Much like inserting a square peg into a round hole, some previous attempts to account for conceptual combination have taken compositionality as a given, and formulated mechanisms for accommodating noncompositional phenomena. In contrast, in conceptual blending theory, the goal is to formulate an account of conceptual combination that is general enough to encompass both compositional and noncompositional phenomena.

This chapter addresses the application of blending theory to concept combination coded by modified noun phrases. In section 1, traditional assumptions about predication in nominal compounds are contrasted with those offered in conceptual blending theory. In section 2, I point to similarities between the difficulty associated with accounting for people's understanding of *predicating* adjectives and that of accounting for *nonpredicating* adjectives and modified noun phrases. Section 3 discusses the treatment of privative adjectives such as "fake" or "toy," which appear to predicate the absence of certain essential properties. We shall see that to understand even simple cases of noun modification, we must appeal to processes operative in apparently more

complicated phenomena of indirect reference, metaphor, and counterfactuals.

5.1 NOMINAL COMPOUNDS

Historically, nominal compounds have comprised an unruly set of phenomena for linguists to explain (see, e.g., Jesperson, 1942). The problem is that it is difficult to provide an abstract characterization of the conceptual relationships that exist between nouns in a noun-noun compound. For example, Hampton (1991) notes the contrast in the contribution of the modifier noun in "hand wound" and "gun wound." A hand wound is a wound *to* the hand, while a gun wound is inflicted *by* the gun. One might speculate that as a modifier "hand" serves as a patient, while "gun" serves as an instrument. However, to counter this suggestion, Hampton offers "hand repair," in which the *hand* serves as the instrument, and "gun repair," in which the *gun* serves as the patient. Similarly, Shoben (1991) notes the difference in the contribution of "rain" to "rain forest" and "rain cloud." Moreover, while the contribution of "mountain" is locative in both "mountain town" and "mountain river," the town is likely to be small, while the river is likely to be large.

One of the most ambitious attempts to deal with nominal compounds in generative linguistics is by Levi (1978). Levi proposes that nominal compounds result from a syntactic transformation in which relative clauses such as "a wound caused by a gun" are converted into nominal compounds such as "gun wound." Consequently, the semantic relationship between the nouns in a nominal compound can be characterized as encoding one of nine recoverably deletable predicates such as "cause," "have," "make," "use," "be," "in," "for," "from," and "about." Although these predicates have been deleted at the surface structure, we can interpret these compounds semantically by using the reconstituted relative clause. Hence the name *recoverably deletable predicates.*

However, many examples of nominal compounds are hard to classify on Levi's taxonomy, either because they fall into more than one category or because they don't fit under any of them. For example, "dog collar" might be paraphrased either as a "collar *used* by a dog" or a "collar that a dog *has*" (Lehnert, 1988). Moreover, while a "rain forest" is a forest that has rain, this is not what makes it a rain forest or what differentiates it from other forests. While Levi's taxonomy might look explanatory, the predicates she picks are so general, they don't necessarily map on to unitary concepts. For example, while both "dog collar" and "college

student" arguably involve the predicate "in," the two kinds of containment are different.

Levi's goal is to account for the productive semantic content of nominal compounds, and, for Levi, this means the content in their semantically equivalent relative clauses. While Levi's account satisfies her *own* goals, it leaves important aspects of meaning construction unexplained. Levi basically succeeds in reducing the meaning of compounds to relative clauses. However, the productive character of Levi's vague relative clauses itself depends on speakers' ability to coordinate and integrate abstract notions such as containment with more specific information about the container and the thing contained. In fact, because recovering the predicate doesn't begin to explain the specific interpretations that speakers give for nominal compounds, from the cognitive science perspective the explanatory content of Levi's scheme is actually quite limited.

In contrast, Downing (1977) suggests the interpretation of nominal compounds depends on the retrieval of substantive information about component nouns. She gave subjects novel (though attested) noun-noun compounds without contextual support and asked her subjects to interpret them. The definitions people generated often included domain-specific causal and relational information. For example, for "frog-slime," one person wrote, "the slime that frogs exude to keep from dehydrating" (Downing, 1977: 826). For "pea princess," seven out of ten people referred to the fairy tale "The Princess and the Pea," as in "a genuine princess, who passes the test of a pea under twenty mattresses." Responses such as these suggest people utilize abstract features of related knowledge structures in order to integrate novel compounds.

Moreover, the incorporation of world knowledge into representations that enter into the integrative process has led cognitive modelers to increased success in building systems that can "understand" novel compounds. For example, Leonard (1984) based a cognitive model of noun-noun compounding on a corpus search of two thousand compounds from literary works that have appeared in the past two to three centuries. Leonard attributed the relatively high success rate of her model (which generated a correct interpretation approximately 75% of the time) to the inclusion of various types of specific information not commonly listed in the lexicon. Lebowitz (1983) describes a system that can incorporate novel compounds (in a restricted technical field). This system begins by searching for previously encountered relationships between the two component nouns. If this fails, the model accesses

knowledge structures about related words and can generalize from previous examples.

The difficulty in finding an algorithmic procedure for interpreting nominal compounds is thus similar to the problems associated with generating an account of sentential integration (discussed in Chapters 2–4) and the nontrivial nature of understanding connections that obtain between sentences in connected text and discourse (as we will see in Chapter 8). Speakers do not (and cannot) explicitly encode everything that they know, or even everything that they wish to convey, in the strings that comprise their utterances. Consequently, compound interpretation is highly context-dependent, and greatly influenced by pragmatic and discourse factors. Downing (1977: 839) writes:

> these kinds of facts . . . derive in large measure from cultural values associated with the various entities to be named, and from the primary linguistic fact that speakers code what is salient to them within a given context.

5.1.1 Conceptual Blending and Nominal Compounds

The semantics of nominal compounds thus present a clear case for the application of conceptual blending (previously argued by Oakley, 1995; and Turner & Fauconnier, 1995). As discussed previously, the resultant meaning of nominal compounds often requires an appeal to causal and/or relational information contained in frames from each of the component nouns' domains. Moreover, the main difficulty that nominal compounds pose for generative theorists is the impossibility of deterministically predicting emergent properties that are either absent from or underspecified in the input domains.

One problem with formulating rules for the semantic interpretation of nominal compounds is that while the component nouns clearly bear on the interpretation of the resultant compound, the flexibility that speakers evidence in the production and comprehension of these compounds defies any parsimonious classification scheme. Although there are clearly regularities to be found in the data, speakers have an uncanny ability to generate productive new compounds that fall outside of existing classification schemes. Similarly, Ryder (1994) argues against accounts of compounding that provide a list of rules and include a category such as "other" or "miscellaneous" to cover examples that cannot be accounted for (e.g., Hatcher, 1960). Working from the theoretical framework of cognitive grammar (Langacker, 1987), Ryder argues instead that nominal

compounds form a continuum from less to more productive patterns that speakers exploit in both the production and comprehension of these constructions.

In fact, much of the groundwork for Fauconnier & Turner's (1994; 1998) conceptual blending theory is laid in Langacker's (1987) theory of cognitive grammar. Rather than positing a static lexicon, Langacker argues for the existence of extensive *semantic networks* associated with individual words. Besides providing access to the usual sorts of lexical semantic information, Langacker's semantic networks include encyclopedic and even episodic information associated with a given word. Further, Fauconnier & Turner's (1994;1998) processes of conceptual blending are presaged by Langacker's *integration* of *component structures* into a *composite structure*. Langacker (1987: 76) describes the production of emergent structure in these processes as *accommodation:* "one component may need to be adjusted in certain details when integrated with another to form a composite structure." For example, the "house" in "doghouse" is slightly different from the default model of "house."

The component nouns in a noun-noun compound, and indeed the abstract frame for the noun-noun construction, serve as prompts for speakers to construct a composite structure. As a general account of conceptual integration, conceptual blending analysis can be applied to frames at various levels of abstraction, including the integration between lexical semantic frames and the more abstract frames of construction grammar. For example, Fauconnier & Turner (1995) discuss how mechanisms of conceptual blending apply to the caused motion construction (Goldberg, 1995), as in (1):

(1) He sneezed the napkin off the table.

Work on blending in grammar by Mandelblit (1997) and Fauconnier & Turner (1996) involves assuming that constructions represent abstract scenarios that are blended with lexical information to evoke a resultant construal both more specific than the abstract input and more complete than the construal a speaker might construct given lexical input alone. For example, in (1), a speaker might appeal to elaboration to comprehend how sneezing could cause the napkin to leave the table. Emergent properties of on-line meaning thus result directly from the creative processes of blending.

These processes of meaning construction in nominal compounds are well captured by the conceptual integration network. Frames associated with each of the component nouns are evoked in the input spaces of the

network. The generic space contains abstract commonalities, and the blended space contains elements and relations selectively projected from both of the input domains. Further, Leonard's (1984) modeling work predicts specific information about component nouns to be important for determining what information is projected from the inputs into the blend. For example, one of the most useful types of representations in Leonard's system was *secondary features* derived from verbs related to the nouns in the compounds. This sort of information included whether the verb was transitive or causative, and information about the thematic role of the nominal with respect to the related verb.

In "gun wound," for example, the first input might contain a frame in which "gun" participates, and the second input might contain a frame for "wound." Leonard's work suggests one strategy for determining which of the many possible frames to evoke for "gun" and "wound" are the default action frames associated with each domain. Action frames can be used to set up more comprehensive mappings between the domains that are being integrated. Moreover, information such as transitivity and thematic roles can help to constrain the mapping process. The generic space in this network contains a generic Violent Act frame with a cause and a result, while the blend contains a more specific Shoot frame structured with information from the input spaces. An illustrative set of cross-domain mappings is outlined in table 5.1.

Table 5.1: Cross-Space Mappings in the Gun Wound Blend

Input1	**Input2**	**Blend**	**Generic**
Gun	Wound	GunWound	ViolentAct
Elements	**Elements**	**Elements**	**Elements**
Agent	DangerousAct	Agent	ViolentAct
Target	Human	Human	Patient
Gun	Means	Gun	Means
Result	Wound	GunWound	Damage
Relations	**Relations**	**Relations**	**Relations**
Shoot (Agent, Gun, Target)	Cause (DangerousAct, Means, Human)	Shoot (Agent, Gun, Human)	Cause (ViolentAct, Means, Patient)
Result (Result)	Result (Wound)	Result (GunWound)	Result (Damage)

From the "gun" frame, the blend inherits information about the cause. From the "wound" frame, the blend inherits information about the result. The two inputs share (among other things) a common associated event sequence, although they profile different aspects of it. In the resultant blend, the wound has been profiled just as it was in the input space for the head noun, "wound." Moreover, the composition of the two frames results in the activation of more specific information about the nature of wounds caused by guns. For example, one might be inclined to infer that the wound was serious, that there was internal tissue damage, or that a bullet is lodged in the patient's body.

One might argue that the compound "gun wound" fits nicely into Levi's (1978) taxonomy of recoverably deletable predicates as a "cause" relationship. However, the predicate "cause" works only because elements likely to be named as causes correspond to prominent roles in frames. Given the philosophical quandary surrounding the proper definition of a "cause," it's unlikely that this concept adds anything to a cognitive account. For example, in "gun wound," is it really the gun that causes the wound? One might argue that the bullet *causes* the wound. Alternatively, one might argue that it's the shooter who *causes* the wound. On the cognitive account, one need not appeal to an objective definition of causality, but instead appeal to an idealized model of the event sequence surrounding getting shot. Regardless of the actual cause, we understand the compound "gun wound" because the cognitive model makes "gun" cognitively accessible. Similarly, frames for tennis enable us to understand "tennis injury" as an injury to a tennis player caused by *playing* tennis (rather than, for example, an injury caused by the abstract existence of tennis), because playing tennis is a salient aspect of the Tennis frame.

Further, this suggests that we might expect to find relationships between input frames in two-sided frame networks that, although not *causal* in the usual sense, suggest the operation of cognitive models. In this light, it is interesting to contrast the causal connection represented in the frames evoked by "gun wound" with the more elaborate causality invoked in the nominal compound "caffeine headache." A "caffeine headache" is a headache caused not by caffeine, but by its absence. Understanding this compound thus requires knowledge of caffeine as an addictive substance, including the symptoms of caffeine withdrawal. Meaning construction in this example recruits structure from an implicit counterfactual scenario in which the experiencer ingests caffeine and does not experience the headache.

Table 5.2: Inputs to the Caffeine Headache Blend

Actual Present	**Actual Past**
Elements	**Elements**
Agent	Agent
Headache	
	Caffeine
Relations	**Relations**
Experience (Agent, Headache)	Not-Ingest (Agent, Caffeine)

The inputs to this blend involve the actual scenario in which the agent experiences a headache, and a past space in which the agent does not ingest caffeine (see table 5.2). To understand the relationship between the noningestion and the headache as an integrated event frame that might be classified as causal, it's necessary to evoke a salient counterfactual scenario in which the agent engages in typical ingestion patterns. Consequently, understanding the compound "caffeine headache" involves the construction of not one blended space, but two. One blend integrates the actual past (in which no caffeine was ingested) with the actual present (headache), while the other integrates the counterfactual (typical) past in which the agent does ingest the caffeine, with a counterfactual present in which the agent has no headache (the Input1 space in table 5.3). Some

Table 5.3: Cross-Space Mappings in the Caffeine Headache Blend

Input1	**Input2**	**Input3**	**Blend**
Headache	Counterfactual	Caffeine	CaffHeadache
Elements	**Elements**	**Elements**	**Elements**
Behavior	Caffeine	Caffeine	No-Caffeine
Headache	No-Headache	Awareness	Headache
Relations	**Relations**	**Relations**	**Relations**
Causes (Behavior, Headache)	Causes (Caffeine, No-Headache)	Causes (Caffeine, Awareness)	Causes (No-Caffeine, Headache)

of the cross-space mappings in the caffeine headache blend are outlined in table 5.3. Note that the typical goal of caffeine ingestion is *not* the prevention of headaches, but its effect as a stimulant. This information, as well as the existence of a generic schema (not listed in the table) that relates food or drug ingestion to the experiencer's phenomenological state, can be viewed as background knowledge that helps structure the integration network.

Understanding what a caffeine headache is involves the idea that the actual headache is caused by the failure to ingest caffeine, and the counterfactual reasoning schema that if the experiencer had ingested caffeine there would be no headache. Similarly, other addictive substances can be named in compounds that describe withdrawal symptoms. Take, for example, "nicotine fit." Compounds such as "caffeine headache" and "nicotine fit" can be understood because cognitive models of addiction and withdrawal symptoms license links to a counterfactual space in which the agent has engaged in typical ingestion patterns. Moreover, mappings between actual and counterfactual spaces are by no means a drug-induced phenomenon. Turner & Fauconnier (1995) point to modifiers such as "child-safe" and "sugar-free" that also invoke a disanalogy between actual conditions and a salient counterfactual alternative.

5.2 PREDICATING AND NONPREDICATING ADJECTIVES

Given that modified noun phrases involve a predicate and an argument, we might expect a simple compositional approach to account for their combined meaning. For example, one might approach concept combination in these cases as a matter of set conjunction. On such an account, the combination of concepts C and D occurs via the conjunction of the set of c's with the set of d's. This can be done either *extensionally*, by picking out the entities in the world that occur both in the set of c's and the set of d's, or *intensionally*, by appealing to the union of the set of defining attributes of C and those of D. For example, a "brown cow" could be defined extensionally by finding the intersection of brown things and cows, or intensionally by the union of the set of features that define "brown" with the features that define "cow." Because the extensional approach lacks psychological plausibility, most psychologists interested in conceptual combination have looked to intensional conjunction.

Besides its formal appeal, viewing conceptual combination as the conjunction of features confirms our intuition that adjectives predicate the

nouns they modify. Unfortunately, actual language use abounds with nonpredicating adjectives. Lakoff (1987) points to a number of examples in which the meaning of a noun phrase is not the set intersection of the component parts. These include "electrical engineer," "topless bar," and "heavy price," among others. However, the meaning of nonpredicating adjectives is usually so straightforward that their presence goes unremarked by all but linguists with the preconceived idea that adjectives predicate the nouns they modify.

Take, for example, the term "hot lids," which is often seen on signs in convenience stores that sell both hot and cold drinks. The lid for the paper cup that contains a cold drink is made of a thin plastic, and has a hole through which a straw can be poked. However, because hot drinks are more dangerous when spilled than cold drinks, and because a hot drink might melt the thin plastic lid, manufacturers make a more sturdy plastic lid for hot drinks. A "hot lid" is thus a lid for a cup that contains a hot drink. Because a hot lid is not a lid that is hot, the speaker is unlikely to construct the correct blend for "hot lid" based on default strategies.

While blending theory does not provide a detailed account of how abstract information represented in the inputs is accommodated in the blended space, it suggests that speakers utilize constructive facets of retrieval processes (see Fruzzetti, Toland, Teller, & Loftus, 1992, for review) to recruit specific information to structure the blended space. Moreover, blending theory offers hypotheses about what sorts of information are appealed to and what sorts of constraints are operating on the integrative mechanism. For example, conceptual blending theory predicts that information at various levels of abstraction is involved in meaning construction for the compound "hot lid," that the inputs contain relatively more abstract structure than the blended space, and that the mapping of elements from Input1 into the blended space will be constrained by the topology of the organizing Container frame.

Conceptual blending in "hot lid" involves a two-sided frame network in which Input1 contains a frame for Hot which occurs in the larger domain of Temperature (see table 5.4). Because a "lid" is a movable cover for the opening of a container, Input2 is the Container frame in which Lid has been attentionally focused, or profiled. The need for a frame for Hot can be seen in the extent to which the temperature denoted by Hot reflects knowledge about the substance whose temperature is in question. Hot tea is a different temperature from hot food. On the coast, it's a hot day when the temperature goes above 80 degrees Fahrenheit.

Table 5.4: Temperature Domain

Temperature
Hot (Substance)
Solid/Liquid/Gas
Warm (Substance)
Solid/Liquid
Cold (Substance)
Solid/Liquid

However, in the desert, it's not hot unless the temperature goes above 100 degrees Fahrenheit. In 1994, the context-sensitivity of "hot" was reinforced by a court ruling in which an Albuquerque, New Mexico, woman was awarded $2.9 million* in punitive damages after being scalded by coffee that was deemed to be too hot. Although the coffee container was marked with a warning that the contents were "hot," the count decided that coffee over 200 degrees Fahrenheit was hotter than a reasonable person should expect "hot coffee" to be.

Integrating the two frames to understand "hot lid" involves recruiting contextual knowledge about containers and lids as well as the implicitly available contrast between "hot" and other aspects of the temperature domain (see table 5.5). In this case, structuring the network involves information rooted in the activity in which the term "hot lids" is meaningful – buying coffee in a convenience store. This process involves getting a cup, filling it with coffee, and putting a lid on the cup. At the point in the process when the need for a lid becomes salient, the abstract Container frame will have been bound to the speaker's perception of the particular cup of coffee in hand. The metonymic relationship between the contents of a container and the lid affords the mappings that go into the Hot Lid blend (see table 5.5).

Faced with a choice of lids, where one pile is labeled "hot lids," contextual information in this case promotes ascription of the property Hot to the liquid already in the cup. Similarly, the cross-space mappings in this integration network license the blend for "cold lid," in which the cold end of the temperature frame is ascribed to the liquid in the cup,

* The award was later reduced to a mere $480,000 by Judge Robert Scott.

Table 5.5: Metonymic Relations in the Hot Lid Blend

Input1	**Input2**	**Blend**
Temperature	Container	HotLid
Elements	**Elements**	**Elements**
	Container	Cup
Substance	Substance	Liquid
	Lid	Lid
Relations	**Relations**	**Relations**
	Contains (Container, Substance)	Contains (Cup, Liquid)
	Covers(Lid, Container)	Covers(Lid, Cup)
Hot(Substance) Solid/Liquid/Gas		Hot(Liquid) Lid[Sturdy, Plastic]

and the associated lid is evoked by metonymy. In both of these cases, the activity of choosing which lid to put on the coffee is crucially involved in the construction of the blended model. Moreover, the Hot Lid blend is possible because of speakers' ability to set up the more extensive mappings between temperature and lids outlined in table 5.6.

In this respect, the blend for "hot lid" does not differ in kind from the more compositional-looking compounds "small lid," "medium lid," and "large lid." Just as the scalar value that can be assigned to the

Table 5.6: Cross-Domain Mappings in the Hot-Lid Blend

Input1	**Input2**	**Blend**
Hot(Substance)	Hot(Liquid)	Lid1
Solid/Liquid/Gas	Coffee/Tea	[Sturdy-Plastic]
Warm(Substance)	Warm(Substance)	
Solid/Liquid		
Cold(Substance)	Cold(Liquid)	Lid2
Solid/Liquid	Coke/7up/Lemonade	[Translucent]

Table 5.7: Small Lid

Input1	**Input2**	**Blend**
Size	Container	SmallLid
Elements	**Elements**	**Elements**
	Container	Cup
	Substance	Liquid
Object	Lid	Lid
Relations	**Relations**	**Relations**
Small(Object) Medium(Object) Large(Object)	Size(Lid) Size(Lid) Size(Lid)	Small(Lid) Medium(Lid) Large(Lid)
	Size(Cup)	Small(Cup)
	Contains (Container, Substance)	Contains (Cup, Liquid)
	Covers(Lid, Container)	Covers(Lid, Cup)

temperature of a hot substance is modulated by expectations about the normal temperature of the substance, the actual size of "small," "medium," and "large" is relative to the object of the frame for Small, Medium, and Large, and, of course, relative to one another (see table 5.7).* Moreover, in the context of drinks bought at the convenience store, there is a systematic correspondence between the relative size of the lids and the relative amount of liquid. Thus, "small lid" does not pick out lids of a particular size by virtue of some objective measure of smallness. Rather, the correct understanding of "small lid" in the convenience store context depends upon the speaker's ability to match lids to containers on the basis of their relative size. Similarly, "hot lid" picks out lids of a particular design in a way that relies upon the speakers' ability to match lids to containers on the basis of the relative temperature of their contents. While not compositional, the term "hot lids" is *analyzable* in virtue of the systematic way in which the Temperature domain can be mapped onto the container's contents, and how – especially in the

* Note how the bigger-is-better American cultural model has led some stores to offer Large, Extra-Large, and Super Size instead of Small, Medium, and Large.

situated context of use – the contents are a salient aspect of the Container frame metonymically evoked by "lid."

5.2.1 Predication and Noun Modification

Conceptual blending thus provides a framework for addressing meaning construction in cases where compositional approaches fail. However, adjectives can be divided into predicating and nonpredicating varieties. The advocate of compositional approaches to language comprehension might argue that while conceptual blending may account for the behavior of nonpredicating adjectives, compositional mechanisms are better suited to accounts of predicating adjectives. In contrast, I argue that predicating adjectives pose some problems of their own, and will show how conceptual blending can accommodate these problems nicely.

As noted previously, a compositional account of noun modification has a certain intuitive appeal. The most obvious explanation of predicating adjectives is that they specify one of the predicated object's attributes. However, understanding noun modification in this way entails a number of empirical predictions that are not borne out by existing data on conceptual combination. First, the compositional approach relies on outdated notions of category structure defined by necessary and sufficient conditions. Second, a strictly compositional account incorrectly predicts the effect of predication to be describable via monotonic logic. Finally, a compositional approach cannot account for the ways in which typicality effects of a combined concept are related to typicality effects accruing to its components.

A key assumption of compositional approaches is that of "classical" categories defined by necessary and sufficient conditions. However, discoveries in several subdisciplines of cognitive science have revealed prototype, or typicality, effects in people's categorization behavior (see Lakoff, 1987, for extensive review). While most psychologists acknowledge the reality of prototype effects, they have not interpreted these effects as invalidating objectivist paradigms. Rather than abandoning old assumptions about category structure, many cognitive psychologists have instead sought to incorporate prototypical representations into a compositional account. For example, an early attempt to deal with typicality effects involved an appeal to fuzzy logic (Zadeh, 1965). In fuzzy logic, truth is a continuous variable ranging from 0 (exclusive) to 1 (inclusive): 0 is absolutely false, and 1 is completely true. We can

apply fuzzy logic straightforwardly to typicality effects by representing the truth value of a category membership statement such as "a dog is a pet."

Fuzzy logic can also be used to predict typicality effects that result from conceptual combination. The truth value of a conjunction is given by the minimum rule: The truth value of (p&q) is equal to the smaller of the individual truth values for p and q. Consequently, Osherson & Smith (1981) argued that the typicality of an exemplar in a combined concept such as "pet fish" should be predictable from its typicality in the category "pet" and its typicality in the category "fish." However, they found that while people considered the guppy to be a good example of a "pet fish," it was a poor example of both a "pet" and a "fish." Osherson & Smith concluded that the guppy effect indicates people's typicality judgments violate the minimum rule in fuzzy logic, and consequently cast doubt on prototype theory.

While Osherson & Smith were willing to admit that *some* aspects of people's categorization behavior could not be captured by classical categories with necessary and sufficient conditions, they were not willing to drop the idea completely. Instead they proposed a distinction between a concept's *core* and its *identification procedures*. Osherson & Smith propose a dual theory for handling data on categorization and conceptual combination. All concepts have a *core* representation that is involved in meaning and conceptual combination, and *identification procedures* that are used in categorization and typicality judgments. This argument was also made by Armstrong, Gleitman, & Gleitman (1983) to explain the fact that people understand that all even numbers fall in the category *even numbers*, though numbers 2 and 4 are consistently rated as being more typical than other even numbers.

In another experiment, Osherson & Smith (1983) showed people pictures and asked them to judge how good each picture was as an example of an apple, a red apple, a brown apple, a red thing, and a brown thing. Just as the nominal compound "pet fish" yielded the guppy effect, this study found that people considered a picture of a brown apple to be a better example of a "brown apple" than a "brown thing" or an "apple." Data such as these suggest that noun modification via predicating adjectives can yield a concept whose attributes differ from those of its constituents. Moreover, sometimes a combined concept has unique features that can't be attributed to either of the constituents. For example, Murphy (1988) finds people rate "is pulled over your head" to be typical of "casual shirts" but not of "casual" things, nor of "shirts."

Another factor that contributes to these sorts of effects is the fact that set theoretical reasoning just doesn't seem to obtain. Hampton (1988) found that 94% of people queried agreed with statements of the form, "Garden furniture is a type of furniture." However, only 34% agreed with "All garden furniture is furniture," instead preferring "Some garden furniture is not furniture." This suggests that status as a subordinate member of a category (paradoxically) does not entail category membership.* Further, Hampton finds that subjects deny that a screwdriver is a weapon, but assert that it's a weapon that is a tool and that it is a tool that is a weapon. In a similar vein, Medin & Shoben (1988) found subjects rated "round" as a typical property of "watches," but "rectangular" as a typical property of "digital watches."

Medin & Shoben had people make typicality judgments for exemplars in simple categories, such as Games, and combined categories, such as Board Games. The study found that the correlation between the ratings for each simple category and its derived category depended on whether the derived category involved a typical or an atypical attribute of the category. For example, "board games" are typical exemplars of the category Games, and the typicality ratings for both categories were highly correlated. In contrast, "alcoholic beverages" are atypical examples of the category Beverages, and the typicality ratings for the two categories were negatively correlated. Finally, "raw vegetables" are not typical of Vegetables, but they are not atypical either (e.g., contrast them with tropical vegetables); there was no correlation between typicality ratings in "vegetables" and "raw vegetables."

The advocate of compositional approaches might object that these findings are irrelevant because the emergent properties discovered by Murphy (1988) and Medin & Shoben (1988) can be attributed to people's experience with the objects named by the noun phrase. Murphy thus argues for the use of background knowledge in conceptual combination. Rips (1995), however, sees this argument as illegitimate. For Rips, the case where the speaker knows something about the resultant output of a nominal compound is not relevant to whether the combination process is compositional. Because compositionality is invoked principally to

* However, proper interpretation of Hampton's (1988) results requires consideration of how people treat natural language quantifiers such as "all." Moxey & Sanford (1993) have argued quite convincingly that the use of natural language quantifiers is motivated by speakers' desire to convey various sorts of conversational implicatures, and does not correspond well to the quantifiers of first order logic.

explain speakers' ability to understand and produce novel expressions, the critical test would involve a case where the speaker knows the meanings of the component nouns, but not of the resultant compound.

In fact, Yeh & Rips (1992) conducted one such test in which they asked subjects to list properties for adjective-noun combinations such as "smoky chimney," and "smoky apple," as well as individual words "smoky," "apple," and "chimney." Yeh & Rips found people listed emergent properties for both the familiar and the novel noun phrases. For "smoky apple," for example, a number of subjects listed "dried" and "bad tasting," even though neither of these properties was listed as pertaining to the individual words "smoky" and "apple." These data suggest that the concept of "smoky apple" cannot arise from the monotonic composition of "smoky" and "apple." Rips (1995) suggests that the observed pattern of results might have been caused by invoking knowledge structures about apples and about smoky things and performing some sort of elaborate causal reasoning.

5.2.2 Conceptual Blending and Predicating Modifiers

The account of noun modification offered here grows out of the more general philosophy of cognitive linguistics that individual elements contribute to the formulation of the overall meaning of an utterance – although not always (or even often) deterministically. Although there is substantial regularity in the varieties of meaning construction afforded by language, the meaning of the whole is *not* computable from the meanings of the constituents. Because the function of language is to enable the listener to participate in the interactive frame set up in a shared context, a noun phrase need only provide the listener with enough information about the element in question to connect the phrase with contextual information and/or background knowledge.

According to blending theory, both those cases where people have foreknowledge of the resultant compound and those in which knowledge is lacking are relevant. The nature of the blend will depend on the amount and types of knowledge one has of the input domains. In the case in which the listener has little knowledge about the resultant compound, default models of the input domains will be integrated. However, in the case of contextually defined compounds, the listener must integrate abstract information from memory with perceptual and conceptual information of her current surroundings. Further, once a space has been set up and structured, it is quite simple to recruit all sorts of

information from long-term memory, and in varying degrees of detail. This information can include episodic memories regarding particular interactions with, for example, digital watches, or semantic memories about, for example, casual shirts.

Pet Fish

Conceptual integration of "pet" and "fish" is an example of a relatively tame blend. The first input space is structured with a frame in which "pet" participates, while the second is structured with a frame in which "fish" participates. Unlike some of our previous examples (such as "hot lid"), "pet fish" is a case in which the two named elements ("pet" and "fish") are conceptual counterparts and map onto the same element in the blended space. The network in this example is a *two-sided network* in which both of the inputs contribute frame-level structure to the blend. Frames having to do with pet ownership are projected from Input1, while frames that concern the attributes of being a fish are projected from Input2. The resultant blend contains elements and relations from both of the input domains. The differences at the specific level arise because of the processes of composition and completion that help structure the blended space.

Hampton (1991: 110) notes how "pet fish" inherits some attributes from the "pet" domain and others from the domain of "fish":

> Pet fish inherit habitat and not being eaten from the pet concept, because the attribute values for PET are more criterial for PETS than are attribute values of FISH criterial for being a FISH. On the other hand, the possession of gills and fins is more critical for fish-hood than is the possession of lungs and legs for pet-hood. Hence the composite prototype for PET FISH is a hybrid – having some attributes and lacking others of each parent.

While Hampton's example is good for explaining how these concepts are combined in a default context, it is not readily extendable to cover the many different ways that "pet fish" might be used in context. For example, an old man who regularly feeds the carp in a public pond might consider them to be his pet fish. The input spaces in this blend would resemble those sketched in table 5.8. However, its emergent structure would differ from the structure in the default blend since the Lives-in predicate would not need to be completed with Tank, but would instead be filled by the contextually available Pond.

Similarly, consider the case of a field biologist studying shark behavior. While she does not live with her shark subjects or feed them, she

Table 5.8: Pet Fish

Input1	**Input2**	**Blend**
Elements	**Elements**	**Elements**
Pet	Fish	PetFish
Owner	Water	Owner
House		House
		Tank
Relations	**Relations**	**Relations**
Feeds(Owner, Pet)	Lives-in (Fish, Water)	Feeds(Owner, Fish)
Loves(Owner, Pet)		
	Swims(Fish)	Swims(Fish)
		Lives-in(Fish, Tank)

might well consider them to be her pet fish. Unlike the case of the default context discussed by Hampton (1991), contextually available information here suggests projection of a different structure from the default "Pet" domain (e.g., Loves(Owner, Pet)), as well as completion with structure recruited from the subframe of Shark in the domain of Fish. Besides the animal that we designate with the term "pet," there are a whole set of associated activities that go along with having a pet: giving an animal a name, ascribing it with a personality, and doing other sorts of anthropomorphic projection. Although not central to the prototypical pet concept, it is these associated properties that license our biologist to categorize the sharks as her pet fish. Because these associated, peripheral properties can also license categorization, an account of the typical cases is not sufficient.

The preceding examples are intended to demonstrate the tremendous variability in the various instantiations of the same concept. However, besides making the negative point that intensions don't determine their extensions (contra Frege, 1970; 1892), I have endeavored to make the positive point that cross-space mappings play an important role in conceptual combination. In predicating adjectives, conceptual combination tends to rely on established mapping schemas, while in nonpredicating adjectives, conceptual combination depends on our ability to recruit relevant domain-specific knowledge to establish novel mappings.

These observations suggest that rather than viewing conceptual combination as operations on features and sets, it should be understood as conceptual integration of frames constructed from background knowledge and specific information available from context.

5.3 PRIVATIVE ADJECTIVES

Kamp (1975) defines a privative adjective as one for which, given privative adjective A and any noun N, the claim, "No AN is an N," is necessarily true. For example, "fake" is privative because the statement, "No fake gun is a real gun," is true by definition. To explain how "fake" acts as a modifier, the semanticist's job has been to provide a mechanism for specifying the similarities and differences between real and fake guns. Because, at least at first blush, privatives pose a greater problem than, say, "brown cow," Franks (1995) has argued that privative constructions should be employed as a test for models of conceptual combination. One such model is Franks' (1995) sense generation model. Another is offered by conceptual blending theory. While both models are motivated by the observation of substantial flexibility and context-sensitivity in actual instances of meaning construction, they each deal with these phenomena in fundamentally different ways. The following section compares these two models to reveal the inherent problems with compositional attempts to explain the concept combination in modified noun phrases.

5.3.1 The Sense Generation Model

Franks (1995) argues that the sort of conceptual combination that occurs in privatives is in fact no different from conceptual combination of nouns with other sorts of adjectives (cf. Kamp, 1975). Thus Franks proposes a model that accounts for functional and proper privatives, including privatives with a negating function, as well as the privatives whose function is equivocating. The model generates a sense for a particular combination of privative and head by ascribing features to a discourse referent set up to represent the combined concept.

Franks' account acknowledges the importance of perspectival relativity in classification judgments and how it serves to limit the generalizability of those statements in particular ways. As in mental space theory, perspectival relativity implies that we can adopt various perspectives

on the same discourse referent. This allows us to understand a particular apple as a dessert, a cricket ball, or a still life, depending on what perspective we have adopted. In the sense generation model, perspectival relativity is handled by *pegs*, or discourse referents, and *alecs*, perspective-relative instantiations of pegs (Landman, 1986).

Franks' treatment of pegs (discourse referents) and alecs (perspective relative instantiations of pegs) is generally compatible with the way in which mental space theory establishes elements so that those different spaces have different properties. The main differences are, first, the use of feature-based rather than frame-based representations and, second, that in Franks' system the relationships between the different alecs for a peg are derivable but not explicitly represented or involved in ongoing meaning construction – and thus are missing one of the central insights of mental space theory.

As in conceptual blending, meanings are built up in context, based on information available from long-term memory. In-context meanings, or *senses*, consist of attribute-value structures (AVS). Senses of privative constructions are built up similarly to senses for other sorts of compounds: by recruiting attribute-value structures for the modifier and the head, and performing operations on the representations in order to yield the resultant compound sense. The core of the sense generation model is its three processes for sense generation: unification, priority union, and metonymic type coercion.

Unification is a monotonic operation that simply summates the two sets of features. Consequently, unification occurs only when there is no conflict between the values of shared attributes of the component concepts. The second process, *priority union*, is more useful because of its flexibility. Simply put, priority union is unification with overrides. When combining concepts A and B to form AB, the modifier A has *priority*. Consequently, in cases where A and B have conflicting values for a given attribute (as in the Q attribute in table 5.9), the combined concept AB will inherit the value from A. Moreover, AB will inherit all nonshared attributes and nonconflicting attribute/value pairs directly (see the S, U, and P attributes in table 5.9).

More relevant to privative constructions, *metonymic type coercion* (MTC) is a process in which the modifier A coerces the attribute value structure of the head B to one of its diagnostic features while negating its central features. Franks proposes two forms of metonymic type coercion, *MTC with rebuttal*, to deal with negating privatives such as "false

Table 5.9: Priority Union: An Example (Adapted from Franks, 1995)

A	B	AB
Q:r	Q:m	Q:r
S:t	S:t	S:t
U:v		U:v
	P:l	P:l

eyelashes," and *MTC with undercutting*, to deal with undercutting privatives such as "alleged criminal." In both sorts of MTC, the attribute value structure is divided into *central* and *diagnostic* attributes. In MTC with rebuttal, the central attributes of the head are negated, but its diagnostic attributes are inherited. In MTC with undercutting, the central attributes are left unspecified, and the diagnostic attributes are inherited.

An additional feature of the model, known as *implicit attachment*, can modify explicitly evoked senses by unification with contextually evoked concepts. If there is an informational demand for further specification, the attribute-value structure of an implicitly attached noun can be unified with explicit senses in the compound. For example, if you are in the park and someone says, "stone lion," the AVS for "stone" and "lion" could be unified with the AVS for an implicitly evoked noun such as *statue*.

5.3.2 Fake Guns

Franks (1995) demonstrates his sense generation model with the celebrated example of a "fake gun." Because "fake" is a negating privative, it is handled by MTC with rebuttal. Consequently, the central attributes of "gun" are negated, while its diagnostic attributes are maintained. Franks proposes Fires(Bullets), MadeOf(Metal), and Function(Kill) as central features of the concept gun, and Trigger(+), Barrel(+), Handle(+), and Size(Portable) as diagnostic features. In this case, MTC with rebuttal results in a set of attribute value structures that characterize something that looks like a gun, but doesn't fire bullets.

However, the use of the null context here can be quite deceptive. Besides the object that we designate with the term "gun," there is a whole set of associated activities in which guns play a prominent role. While Franks' sense generation model generates a plausible sense for

"fake gun," there is more to being a fake gun than just not firing bullets. Consider the case where we're comparing two drawings to figure out which of the guns is fake. Neither can fire bullets, but we might well designate one as real and the other as fake. Examples such as this suggest that the diagnostic features we ascribe to guns depend on our cross-space mapping ability. In a default scenario, diagnostic features correspond well with objective features. However, it is quite simple to come up with scenarios in which this is not the case (see Coulson & Fauconnier, 1999, for examples).

The insight that is captured by Franks' metonymic type coercion algorithm is that a "fake gun" and a "real gun" share many of the same features. Although a "fake gun" is not a "real gun," the mappings between the integrated (fake gun) concept and its counterpart in the gun domain allow us to think and talk about fake guns in many of the same ways that we think and talk about real guns. Although Franks' system accounts well for some default scenarios associated with the Gun frame, the property of being a fake is not just a matter of the properties of the gun. It also has to do with the intent of the actor who uses a fake gun and the reaction of the actor's audience. Important aspects of the meaning of "fake" involve the actor's intent to deceive, and the audience reacting to a fake gun as it would have reacted to a real gun.

Speakers treat "fake" as a space builder that prompts a mapping between an actual scenario in which the actor employs the fake gun, and a counterfactual scenario in which his audience reacts as if the gun were *real.* For example, in a case where an actor uses a plastic gun in a robbery, the object in the blended space inherits the property of being *plastic* from the space that represents the actor's own knowledge, and the property of being a *gun* from the victim's belief space. This sort of a mapping goes beyond the presence or absence of central versus diagnostic features, and relies on the speaker's ability to coordinate frames in actual and counterfactual spaces via cross-space mappings.

In the default scenario that Franks uses to demonstrate the sense generation model, the mappings between the fake gun model and its counterpart in the gun domain are similarity mappings between features in object frames. However, being fake need not entail similarity between the fake object and its counterfactual counterpart. The important thing is that its properties could potentially induce the victims to believe that the counterfactual scenario obtains. In the right context, a single feature might produce the desired belief. For example, a pipe in the back of the

neck, a hairbrush in the raincoat, or a balloon popping in a room full of hostages with their heads to the floor might all serve as fake guns. In 1997, in a testament to the adage that truth is stranger than fiction, Carlos Diaz was sentenced to eighteen years in prison for stealing $20 and a watch while armed with a zucchini concealed in his jacket.

To determine what features are relevant in a given example, it is necessary to construct a model of the actor's actions and the victim's perception of those actions. Mappings are constrained by the understanding that in the faker's mind there is a causal connection between his own actions and the victim's beliefs. Of course, the victim need not *actually* be deceived for the fake object to be fake. However, the characteristics of fake objects arise because of the way in which the intent to deceive is central to the concept of "fake." Conceptual combination in each case is driven by the way in which theory of mind determines cross-space mappings between the actor's intentions and the victim's would-be belief. Diagnostic features of guns, both fake and real, result from apprehension of the more extensive set of mappings between the relevant spaces.

Consider the case in which James Bond's engineer Q builds a gun that looks like an umbrella. In this case, our fake umbrella would, indeed, lack the central features of an umbrella while maintaining its diagnostic features. However, what if Q were very clever about the detonation device and constructed the gun so that it could open like a real umbrella and protect you from the rain. We then have a fake umbrella that has central features of both a gun and an umbrella, as well as diagnostic features of the latter. Further, imagine that a movie producer decides to use an umbrella as a prop for the fake umbrella in her James Bond movie. In this case, we have a fake umbrella with both central and diagnostic umbrella attributes, and a gun that lacks the central attributes of a gun.

Goffman (1974) has pointed to the way in which we employ multiple levels of embedding, and, further, has emphasized the extent to which being real or fake is socially constructed. Sometimes the categorization of an object as a gun is directly related to the builder's intention to make a gun. Similarly, sometimes the categorization of an object as a gun is the result of the producer saying the gun is a gun. Because human activity involves the projection of partial structure from domain to domain, there is far more counterpart structure in the world than we typically realize. The object frame for a gun is only one of a potentially infinite set of frames in which a gun participates. Because classification of a

feature as central in a given context depends on the situation frame as well as the gun, there are an infinite number of potentially different configurations of central versus diagnostic features. In conceptual blending, the ascription of features is only a side effect of coordinating the representational structure in the various spaces. Consequently, projection, representation, and rerepresentation can all make objective features of the gun largely irrelevant.

Socially defined activities such as play present similar problems for a feature-based account. For example, if I'm playing cops and robbers, I might stick my thumb in the air and point my index finger at my playmate saying, "Give me the money or I'll shoot." After he complies, I might laugh, saying, "Hah, it's just a fake!" In this case, whether I had a "real" gun or not, bullets wouldn't come out of my finger. In such a scenario, the difference might even involve whether my playmate *pretended* to die after being shot by my gun. Because conceptual blending theory relies on the establishment of mappings based on pragmatic functions such as similarity, identity, and analogy, it predicts the use of these relational counterparts as well as the similarity-based ones.

Conventions of play are such that configuring the hand in a particular way licenses a blend of frames from actual reality (in which the object is a finger) and counterfactual reality (in which the object is a gun). In the blended play spaces, the finger assumes the features of a gun (such as having a barrel, being able to fire bullets, and being a weapon) via partial composition. Moreover, via pattern completion, the finger can be used to evoke other frames associated with guns, such as the robbery scenario discussed previously. Participating in the game means reacting to the finger gun as one might react in the counterfactual reality space. Nonetheless, participants in the game are not deluded – the frames in the reality space are always available.

The gun in this example is not fake because it is made from flesh and blood, or because it is otherwise lacking in central gun features. The term "fake" suggests a possible discrepancy between the actor's knowledge and the victim's beliefs. Within the context of the game, the gun counts as fake, first because one of the participants (the actor) has stipulated that it is so, and second because the victim, in surrendering the play money, has reacted to the gun as if it were a real (play) gun. Examples like this, in which the definition of the scenario depends on participants' construals, present a formidable challenge for models (e.g., Franks, 1995) that rely on context to specify the exact nature of a meaning built from the bottom up.

5.3.3 Stone Lions

Besides specifying "fake" and "alleged," Franks (1995) points to a third type of privative construction, the *functional privative*. Functional privatives are adjectives (or adjectival nouns) that behave as privatives in some constructions but not in others. For example, "stone" acts as a privative in "stone lion," but not in "stone bridge." To account for functional privatives, the sense generation model appeals to the mechanism of implicit attachment in which the attribute value structure (AVS) of a contextually evoked word can contribute to the meaning of the combined concept via priority union. Because implicit attachment is such a potentially powerful mechanism, it is subject to certain constraints. First, an implicit attachment (NI) is required to be situationally appropriate. Further an implicit attachment must be subsumed by the sense of the phrase it instantiates. In his discussion of how the implicit attachments (NIs) are derived, Franks (1995: 463) writes:

> The NI lexical concept is then the concept associated with the head noun of a phrase that individuates the referent. For example, if a discourse or text states that, *A girl was sketching a stone lion in the park*, the instantiation of *stone lion* would likely be as a statue of a lion, and not an ornament in the shape of a lion. The NI lexical concept selected would then be *statue*, and not *ornament*. This is one constraint on NI selection: In addition to being situationally appropriate, an AVS description of the instantiation must be subsumed by the sense AVS of the phrase that it instantiates; in this case, the AVS for *statue of a lion* would be subsumed by one for a *stone lion*.

Although Franks' formalism provides a very clever mechanism to provide the meanings for functional privatives in which the head is (physically) constructed from the modifier (as in "stone lion" or "chocolate rabbit"), the requirement that the implicitly attached lexical concept be subsumed by the phrase it instantiates limits the generality of the mechanism. Ironically, Franks himself provides an example with the sentence, "A girl was sketching a stone lion in the park." The reading discussed by Franks is the one where "stone lion" refers to the statue in the park. However, "stone lion" can also be interpreted as referring to the girl's drawing. A sketched stone lion is not a lion any more than a stone lion is a lion (from the reality perspective, something that Franks would dub "Type I Perspective"), but because the lexical concept for "sketch" isn't subsumed by that for "stone lion," it cannot serve as an implicitly attached concept.

The obvious solution to this problem is to incorporate recursion into the account of functional privatives so that instead of unifying "sketch" with "stone lion," we can combine the concepts using metonymic type coercion (with rebuttal) of "stone lion" and priority union of "sketch" and the coerced "stone lion." The possibility of applying the processes of conceptual combination to an implicitly evoked concept makes Franks' system immensely powerful. However, the main weakness of Franks' model is the requirement of a preexisting relationship between the content of the implicitly evoked concept and the existing information about the referent.

Franks (1995: 471) writes:

> To reiterate, for a categorisation of the appropriate type to be warranted, it must be an instance of the appropriate schema. Following communicative requirements, *c* and NI must be a part of the shared or available information for the agents, present in or derivable from a text or discourse context. Given this availability, perspectival content can, and usually does, remain implicit. Additionally, of course, there is the requirement that the information encoded in the perspectival description subsumes information available about the referent (i.e., that the criterion employed for classification *c* subsumes the information available about the discourse referent Y).

This requirement that the perspective-appropriate concept (i.e., the "sketch") subsume information about the referent has some importunate implications. First, it necessitates the inclusion of a huge (potentially infinite) number of features in the lexical concept for nouns such as "sketch" that might serve as an implicitly attached concept for all things sketchable. Since the set of implicitly attached concepts is presumably large, this requirement would severely compromise the elegance of the partial representations appealed to in Franks' current model. Further, it fails to capture the way that features of the discourse referent accommodate (in Langacker's 1987 sense) the implicitly attached concept. Moreover, because alternative categorizations always depend on the currently ascribed features of the discourse referent, Franks' model is unable to account for the emergent features that arise in cases of frame-shifting.

The crucial factor in meaning construction in functional privatives is not feature unification, but conventions for mapping. Just as Picasso and Michelangelo utilize different representational conventions, so too do different modes of representation, such as sketches, photographs, dramatizations, and the like. The deep reason for the existence of these

constructions is that we have a general notion of *representation* as well as experience mapping between different sorts of representations. Understanding functional privatives is not as much a problem of semantics as of understanding the representational conventions that allow us to map between spaces with different degrees of iconicity. Consequently, positing general blending mechanisms constrained by conventions of representation can make a special feature mechanism superfluous.

A general blending mechanism can also account for the way in which features accommodate (in Langacker's 1987 sense) to the different domains in which they occur. For example, the sense in which a trashcan basketball is "round" differs from the sense in which a leather basketball is "round." In conceptual blending, to say that both balls are "round" means that the element set up to represent each ball can be integrated with a frame for "round." The jagged character of the roundness of a trashcan basketball is an emergent property that can be perceived by participants and conceived by people who can imagine a crumpled-up piece of paper. This sort of composition need not entail the ascription of objective properties. Rather, it relies on the speaker's ability to integrate conceptual structure on the basis of both features and relational information.

5.3.4 Land Yachts and Luxury Cars

A further weakness of Frank's sense generation model is that it doesn't extend to cases that involve analogical components. For example, on Frank's (1995) scheme, "land yacht" would be classified as a *functional negating privative* because of conflicting central features for the lexical concept for "land" and "yacht." Just for the sake of argument, we'll say that "land" and "yacht" conflict on the value for Natural (see table 5.10). This triggers the rebutting MTC operator, which negates the central features of "yacht," yet maintains its diagnostic features (see table 5.11). The next step is to perform priority union on "land" and the type-coerced "yacht." After priority union, we get the representation in table 5.12.

In step two of Franks' procedure, the implicitly attached noun is invoked to create a more specific representation. To come up with the right set of central features, "car" apparently would have to be chosen to serve as the implicitly attached lexical concept. However, without a mechanism for analogical mapping, the sense generation model would either generate an entity that has both central features of a "car" and

Table 5.10: Franks' Hypothetical Representation of Land

Land
Central
Solid: +
Natural: +
Diagnostic
Surface: +
Dry: +

of a "yacht" (see table 5.13 for "land yacht car") or be hampered by the requirement that implicit attachments be subsumed by the sense of the phrase they instantiate.

Contrast the sense generation account of "land yacht" with that offered in Fauconnier & Turner's (1998) account of conceptual integration. The processes of meaning construction for "land yacht" involve the construction of a frame prompted by "land" in Input1 of the integration network, and the construction of a frame for "yacht" in Input2. Because the overt language of a nominal compound such as "land yacht" provides minimal clues to how the integration of input frames is to proceed, the language user is forced to rely on contextual information and background knowledge. Consequently, there need be no assumption that a

Table 5.11: Metonymic Type Coercion with Rebuttal: Yacht

Yacht	**MTC (rebuttal)**
Central	**Central**
Function: sailing	Not[Function: sailing]
Solid: +	Not[Solid: +]
Natural: −	Not[Natural: −]
Diagnostic	
Price: expensive	
Prow: sharp	
Floats:	

Table 5.12: After Priority Union: Land Yacht

Land Yacht
Central
Solid: +
Natural: +
Not[Function: sailing]
Diagnostic
Price: expensive
Prow: sharp
Floats: +
Surface: +
Dry: +

"land yacht" will be a type of "yacht," nor that "land" and "yacht" are necessarily conceptual counterparts.

Rather, it is assumed that "land" and "yacht" can be used in the discourse context (which employs default models in the so-called "null"

Table 5.13: Land Yacht Car

Car	**Unified Land Yacht Car**
Central	**Central**
Solid: +	Solid: +
Natural: −	Natural: + Natural: −
Function: driving	Function: driving Function: sailing
Diagnostic	**Diagnostic**
Wheels: 4	Wheels: 4
Engine: +	Engine: + Price: expensive Prow: sharp Floats: + Surface: + Dry: +

context) to evoke a useful frame in their respective input spaces in the conceptual integration network. An important aspect of conceptual blending theory, absent from the feature-list approach promoted by Franks (1995), is the use of structured representations that facilitate cross-domain mappings. Because the attribute value structures evoked in mental spaces are hierarchically structured such that causal and relational information is readily available, they enable the establishment of mappings between elements and relations in the two inputs that help constrain the projection of structure to the blend. An illustrative list of the corresponding elements and relations in the inputs and the blend from the network for "land yacht" can be found in table 5.14

In their discussion of "land yacht," Turner & Fauconnier (1995) point to salient differences between the information in the blend and the input domains. For example, in contrast to a luxury car driven by a single person, yachts are usually operated by a multimember crew. Moreover,

Table 5.14: Conceptual Integration of Land Yacht

Input1	**Input2**	**Blend**
Land	Yacht	Land Yacht
Elements	**Elements**	**Elements**
land	water	land
driver	skipper	driver
road	course	highway
car	yacht	luxury car
owner	tycoon	rich owner
Relations	**Relations**	**Relations**
Drives (driver, car, road)	Sails (skipper, yacht, course)	Drives (driver, car, highway)
	Yacht	Luxury Car
	Function: sails	Function: drives
	Sign-of: upper-class	Sign-of: wealth
	Owner: tycoon	Owner: rich person

the "land yacht" has many features not associated with the generic car frame. Turner & Fauconnier list passenger-side air bags, electric windows, leather upholstery, and a suspension built for comfort. Although the blended space contains frames whose slots are filled by elements from the domain in the Land input, more identifiable structure in the blend is inherited from the Yacht input. Consequently, the blend provides a way of understanding how the meaning of "land yacht" differs from the meaning of "luxury car." By using the term "land yacht" to refer to a luxury car, the speaker can profile the status associated with owning such a car.

The "yacht" domain provides a useful frame for this blend because of some differences between ownership of a yacht (versus sailboats in general) and ownership of a luxury car (versus cars in general). In the United States, at least, car ownership is quite common and unremarkable, while ownership of a sailboat is more notable. Moreover, the status associated with owning an expensive sailboat is much greater than the status associated with owning an expensive car. Because yacht ownership far supersedes the status of car ownership, sailboats provide a domain that can accentuate the status that goes with ownership. That is, the sailboat domain provides frames that naturally accentuate the status of yacht ownership and inspire awe, which is not as readily available in the domain of cars.

5.3.5 Summary

While it contains many useful insights about privative constructions, the sense generation model seeks to characterize concept combination at the wrong level. Its main shortcoming is its reliance on the ascription of objective features rather than on the higher-level mapping operations that define such features. Conceptual blending is offered as a superior approach to concept combination, because of its appeal to causal and relational information represented in cognitive models, the use of contextually motivated models rather than representations that have been artificially divided into central and diagnostic features, and the role of cross-space mappings constrained by concepts for deception and various kinds of representation.

For humans, representation is not just an internal mental activity, but a physical and a social one in which our intentions and stipulations can make the difference between a real gun and a fake one. Basic human activities such as play, pretense, and representation involve the projection

of partial structure from domain to domain. Because the mappings are not always based on objective features, counterpart structure can't be apprehended with featural information alone. Rather, it involves the integration of perceived features with contextually relevant frames.

5.4 CONCLUSIONS

Blended spaces are mental spaces structured by cognitive models resulting from conceptual blending, in which elements from different domains are combined. Conceptual blending is suggested by the flexibility that people evidence in their actual use of compounds, the fact that people can understand compounds in context, and the fact that they coin new compounds. Conceptual blending occurs via three processes: composition, completion, and elaboration. General characteristics of conceptual blending include the existence of cross-space mappings, the loss of properties from the input domains, the often asymmetric effect of one of the inputs over the other, and the presence of emergent properties.

Contrary to the predictions of a compositional account, data suggest emergent properties of utterances are available *before* the properties of their components. For example, Potter & Faulconer (1979) asked subjects to listen to spoken sentences and to respond affirmatively to pictures if they contained a pictorial representation of a noun named in the sentence. They were told that while some of the sentences might include distracting information about the target word, they were to disregard such information. Nonetheless, upon hearing "burning house," people responded "yes" faster to a picture of burning house – a rather atypical house – than they did to a picture of a typical (nonburning) house.

More recently, Springer & Murphy (1992) asked subjects to read noun phrases such as "peeled apples" and to verify whether various properties – e.g., "white" or "peeled" – were true of the named object. Some of the stimuli were properties of "apple" whereas others were emergent properties of the combined concept "peeled apple." Subjects found emergent properties such as "white" easier to verify than noun properties such as "round," even at fast presentation rates. These data suggest that people quickly integrate lexical information associated with burning things and houses to yield the resultant meaning of the noun phrase.

Analyzing computational-level aspects of meaning construction in these blends suggests a close interrelationship between semantic and

episodic memory whereby information is rapidly abstracted from episodic memories in order to serve as a template for various processes of conceptual blending. Downing (1977) describes one particular deictic compound that arose in a group of graduate students. The term was "bike girl," which referred to a woman who once parked her bike in the vestibule of the department. Although the term was originally coined to perform an identification function, its meaning was eventually extended to refer to any inconsiderate action. This extension of the meaning of "bike girl" seems to appeal to shared background knowledge about the original incident, which involved a negative evaluation of the scenario: "bike girl" being so inconsiderate as to leave a bicycle where people were bound to trip on it.

An important assumption of compositional approaches is that categories have necessary and sufficient conditions that can be combined with algorithmic operations. However, as discussed in section 3, since categories are not normally defined by necessary and sufficient conditions, there is no reason to believe that complex categories result from the combination of such conditions. The response of dual theorists such as Osherson & Smith has been to drop the task of concept combination from the set of the core's responsibilities. In their latest formulation (Smith, Ratterman, & Sera, 1988), the identification procedures are used in categorization, typicality, and conceptual combination, while the core is reserved for membership.

However, Barsalou (1987; 1989; 1993) has argued forcefully against the theoretical utility of positing a stable set of core features in long-term memory for explaining behavior. The instability of categorization behavior as well as people's ability to construct ad hoc, goal-directed categories suggest human behavior is determined by temporary concepts constructed in working memory. Further, examples discussed in this chapter suggest that categories are defined by appeal to idealized cognitive models, and that their combination reflects an appeal to structured background knowledge. For example, there are many noun phrases, such as "hot lid," whose meaning does not involve the properties of the denoted object itself but rather associated frames available from metonymic mappings.

Moreover, Turner & Fauconnier (1996) have previously argued that a number of noun phrases (e.g., "cruelty-free make-up" and "dolphin-safe tuna") rely on speakers' knowledge of seemingly peripheral aspects of the input domains, and mark distinctions that are relevant only because of a salient counterfactual scenario. Further, the extension

of compounds such as "bike girl" seems to exploit aspects of a concept that are important neither for definition nor for identification. Because seemingly peripheral aspects of a concept can be recruited in blends when demanded by current goals, these observations suggest the conceptual core is of little relevance outside of restricted default contexts.

The assumption in the literature influenced by objectivist semantics is that concepts provide the speaker with procedures for identifying an X when confronted with one. However, in conceptual blending theory, this sort of referential success is a trivial consequence of the function of language as evoking and manipulating frames in the context of social activity. The examples discussed in this chapter all support Barsalou's (1993) suggestion that what are combined in particular cases of concept combination are frames created temporarily for the service of meaning construction in a particular context. The role of the knowledge base in long-term memory is one of constraining the topology of the mappings and completing unfilled slots in frames at various levels. The productivity of our concepts seems to be centrally rooted in the ability (i) to accommodate frames at varying levels of abstraction to suit varying contextual conditions; and, (ii) to reabstract aspects of accommodated frames for further application.

5.4.1 Analyzability and Compositionality

One criterion for the adequacy of a theory of how speakers understand each other is how well it accounts for temporary semantic change – or how speakers extend customary meanings for terms to apply in alternative domains or contexts. Alternative domains imply the metaphorical use of terms, generalizations based on abstraction (as in "bike girl"), and novel meanings, as in many of the jokes discussed in Chapter 2. Also, given that one motivation for the original commitment to compositionality was to account for the productive character of language use, any account that can handle novel uses of language should receive serious consideration. Conceptual blending is a noncompositional method of concept combination. Although there are procedures for combining frames, the truth conditions of the resultant frames are not a logical function of the inputs. However, given the noncompositional meanings for many noun phrases, perhaps the proper criteria for judging the adequacy of competing accounts of concept combination should be *analyzability* rather than compositionality.

Langacker (1987) defines analyzability as the extent to which the components of a complex expression can be seen to contribute to the meaning of the whole. Because language users have an awareness of how various components of language contribute to the overall meaning, there are many regularities in the way in which linguistic expressions prompt the speaker to construct meaning. However, it is simply not the case that the meaning of a given expression can be algorithmically derived from a preexisting set of meanings for its component parts. If analyzability rather than compositionality is the requirement, then speakers need not deterministically derive the meanings of an expression from its component parts. Blended models fulfill the requirement of analyzability because they can be sensibly related to the input frames. More importantly, construals afforded by frame networks can support action and interaction in the world.

Compositional approaches to noun modification have often incorporated the assumption that predicating adjectives specify a particular value of an attribute. However, people's behavior on tasks thought to reflect conceptual knowledge suggests that understanding the difference between "red apple" and "green apple" goes beyond their differing values for the attribute Color. Green apples are more tart than red apples, and are more commonly used for cooking (Shoben, 1991). Of course, simply given the data on predicating adjectives, it is unlikely that the semanticist would relinquish a compositional approach to noun modification. One might argue (quite reasonably) that the differences between red and green apples are a fact about the world and not a fact about the meaning of "red apple" and "green apple." However, the idea that the combined concept "green apple" emerges from a conceptual integration network allows a unified approach to both predicative and nonpredicative cases of noun modification.

This is not to say that there is no difference between blends that are built up on-line in the course of meaning construction and conventionalized blends. Conventionalized blends such as those invoked in established nominal compounds are the result of blending frames from two or more domains that have been coordinated by established cross-space mappings. However, the language user need not repeat this blending operation anew every time she uses an established compound. Once a phrase such as "ground ball" has been associated with the appropriate frame, there is no need to blend the input frames on-line. However, the productive way in which people are able to extend these conventionalized blends suggests that the blended model for "ground ball" maintains links to its original inputs.

The account of conceptual blending given here applies not only for the obviously noncompositional examples, but also for the more compositional-looking phenomena. While blending theory is motivated by the need to explain emergent properties, it does not entail that emergent properties will *always* arise. For instance, classic examples like "brown cow" result from the composition of "brown" and "cow" with no alteration in the blended space, and are easily explained by Fauconnier & Turner's (1994;1998) model of conceptual blending. Rather than positing one mechanism for well-behaved phenomena and another for less predictable phenomena, we posit one mechanism that accounts for all the data.

But even the boring case of "brown cow" is not as compositional as it might seem. After all, a brown cow is not brown through and through: Only her fur is brown. The process of composition applies quite straightforwardly because the mapping of color predicates has been conventionalized to produce default mapping strategies. For instance, colors often predicate salient external surfaces. Moreover, a cow need not be *entirely* brown to be dubbed a "brown cow." She might, for example, have one white leg. In fact, one can legitimately apply the phrase "brown cow" to objects that are neither brown nor cows: Imagine a pencil drawing of some cows, labeled with English words such as "brown" and "black and white." While compositional accounts would need to recruit new theoretical machinery to explain this sort of a usage of "brown cow," conceptual blending theory can appeal to the same processes of cross-domain mapping as invoked to explain meaning construction for "small lids," "hot lids," "caffeine headache," "fake gun," "trashcan basketball," and "land yacht."

In this chapter, we have seen that conceptual blending in integration networks is a very general set of processes that operate on encyclopedic information represented in frames constructed in response to contextual needs. Conceptual combination indexed by complex noun phrases involves coordinating frames from input domains on the basis of cross-space mappings in which the establishment of counterpart structure can be based on identity, analogy, and even connections to salient counterfactual scenarios. These observations in turn suggest that the same processes of conceptual integration that underlie meaning construction in noun phrases also operate for statements of analogy, metaphors, and identity in counterfactual utterances.

CHAPTER SIX

Conceptual Blending in Metaphor and Analogy

Metaphor has historically been portrayed as colorful language – aesthetically pleasing but without cognitive import (Hobbes, 1965; Quintillian, 1921–1933). However, in recent years, cognitive semanticists such as Lakoff & Johnson (1980), Sweetser (1990), and Turner (1991) have argued that metaphor is, in fact, a pervasive phenomenon in everyday language and, moreover, that it represents the output of a cognitive process by which we understand one domain in terms of another. Cognitive linguists define metaphor as reference to one domain (known as the *target*, theme, or base domain) with vocabulary more commonly associated with another domain (known as the *source* domain, phoros, or vehicle). On this construal, metaphoric language is the manifestation of conceptual structure organized by a *cross-domain mapping*: a systematic set of correspondences between the source and target that result from mapping frames or cognitive models across domains.

On this view, known as *conceptual metaphor theory*, a speaker invokes a metaphor whenever she refers to one domain, such as verbal argumentation, with vocabulary from another domain, such as physical combat. Conceptual metaphor theory is motivated by the existence of linguistic data like the following (from Lakoff & Johnson, 1980: 4), in which argument is discussed in terms that might just as well be applied to war:

ARGUMENT IS WAR
Your claims are *indefensible*.
He *attacked* every *weak* point in my argument.
His criticisms were right *on target*.
I *demolished* his argument.
I've never *won* an argument with him.
You disagree? Okay, *shoot*!

If you use that *strategy*, he'll *wipe you out*.
He *shot down* all of my arguments.

In the "argument is war" metaphor, argument (the target domain) is conceptualized in terms of war (the source domain) such that physical conflict in the source corresponds to a conflict of opinions in the target. Lakoff & Johnson (1980) point out that the discussion of argument in terms of war is not confined to a few isolated expressions, but can be seen in a number of different examples. For instance, in both sorts of conflict we employ strategies and attack positions. Similarly, just as we talk about battles in terms of a series of attacks, defenses, and counterattacks, we can talk about arguments in an analogous way.

Moreover, the transfer of vocabulary is not arbitrary, but seems to maintain some of the logic of the source domain. For instance, a position in a battle corresponds to a position in an argument, and the spatial configuration of participants in a physical battle maps onto the ideological configuration of participants in a debate. Occupying a position is having an opinion, occupying the same position corresponds to agreement, and occupying different positions corresponds to disagreement. Consequently, *allies* are people who occupy the same position, while *adversaries* occupy different, opposing positions. So, not only do we use some of the same terms to talk about argument and war, the two domains share inferential structure.

In fact, metaphoric expressions seem to reflect the way we think and feel about arguments. While the objective features of physical battles and verbal arguments are quite different, the two domains share abstract similarities. Of course, perceiving these similarities depends crucially on the ability to make a connection between a military position (a piece of land), and an opinion (a mental construct). However, given the initial mapping, simple frame-based inferences about positions in the source domain are understood to obtain for target domain counterparts. Consequently, the systematicity of the shared vocabulary can be seen to reflect a transfer of inference schemas from the source domain into the target.

In fact, over the last twenty years, researchers in cognitive semantics have identified scores of metaphors in every language examined, giving empirical credence to Lakoff & Johnson's (1980) claim that metaphor is pervasive, and systematic, and plays a crucial role in the organization of conceptual structure (for review, see Gibbs & Steen, 1999). One important outcome of this research is the observation that systematic

correspondences between domains themselves reflect cross-domain mappings at a more abstract level. For example, the metaphor *"love is a journey"* is motivated by clusters of expressions like the following (from Lakoff, 1993: 206):

> Our relationship has hit a *dead-end street*. . . . Look *how far we've come*. It's been *a long, bumpy road*. We can't *turn back* now. We're at a *crossroads*. We may have to *go our separate ways*. The relationship isn't *going anywhere*. We're *spinning our wheels*. Our relationship is *off the track*. The marriage is *on the rocks*. We may have to *bail out* of this relationship.

As in "argument is war," these expressions reflect a systematic set of correspondences between travelers and lovers, vehicles and the relationship, and the speed of the vehicle and the intensity of the relationship. Moreover, these mappings are shared by a number of other metaphors that pair long-term, purposeful activities with journeys. So, not only do we construe relationships as journeys, our careers and even our lives can also be conceptualized as journeys. In all of these metaphors, states tend to be mapped onto locations, changes onto movements, causes onto forces, and purposes onto destinations. Viewing metaphorical language as a manifestation of the conceptual system thus explains why the correspondence between elements and relations in the domains of a metaphor is systematic rather than random.

Conceptual metaphor theory (Lakoff & Johnson, 1980) represents a departure from traditional approaches in ascribing to metaphor a central role in a theory of language and conceptual structure. Cognitive linguists have employed the theory's tools to reveal the influence of metaphoric mappings on a wide range of cognitive and linguistic phenomena. Besides their productive role in lexical semantics, cross-domain mappings have attracted the attention of cognitive scientists interested in the role of analogical reasoning in problem solving and decision making (see Holyoak & Thagard, 1995, for review). Gentner & Gentner (1983), for example, argue that physics students use cross-domain mappings of the properties of macro-level physical phenomena to construct mental models that enable them to reason about electricity. They found that participants in their study used a mapping from gates to switches in order to construct a mental model of a switch. Moreover, the students who achieved an adequate understanding of the target domain were those who began with an analogy that matched causal structure in the everyday source domain to that in the target domain of electricity.

A key emphasis in research on metaphor and analogy has involved how conceptual structure in the source domain influences the development and construal of the target domain (see, e.g., Gentner, 1983, 1989; Gick & Holyoak, 1980, 1983; Holland, Holyoak, Nisbett, & Thagard, 1986; Lakoff & Johnson, 1980; Lakoff, 1987). On such accounts, mapping involves, first, noticing a shared system of relations that hold in both the source and target domains, and, second, placing objects from the two domains into correspondence based on common roles in the shared relational structure. In this way, reasoners can begin with a partial mapping of components that play similar roles and later extend the mapping to import novel inferences from the source domain to the target.

Indeed, analogical and metaphorical mapping have been characterized as importing schemas from a well-understood, concrete source domain in order to structure an abstract and less understood target. However, in the following sections, we will see how these traditional accounts of cross-domain mapping have both over- and underestimated the role of the source domain in metaphor and analogy. Section 1 suggests that metaphoric language does not reflect projections between *two* domains, a source and a target, but rather mappings between four or more spaces in a conceptual integration network. Section 2 examines cases of analogy where the point seems to be to project inferences from the blended space in the network to the input that would traditionally be considered the source domain. Finally, in section 3, I argue for a modification in the treatment of metaphoric and analogical mapping that acknowledges both the import of construal and the dynamic character of the cognitive models that speakers construct to understand metaphoric utterances.

6.1 BLENDING IN METAPHORS

Originally, in mental space theory (Fauconnier, 1994), metaphor was handled much like other cases of indirect reference. On Fauconnier's (1994) account, a metaphor such as (1) is handled by setting up two mental spaces: one for the source domain (anatomy) and one for the target (geography).

(1) Paris is the heart of France.

The "heart" is linked to "Paris," and the "body" is linked to "France" by analogical connectors. Once these spaces are linked, models that detail the importance of the heart to sustaining the body are cognitively

accessible to the target domain and can be mapped onto target space counterparts.

However, the two-space model suggests a straightforward correspondence between the two domains that is not always supported by the data. In particular, the two-space model cannot explain examples of cross-space mapping in which source domain models undergo accommodation so as to be more compatible with target domain models. For example, Holyoak & Thagard (1995) point out that the source domain in metaphoric language frequently undergoes a radical change. In (2), for instance, while the target domain "job" is not fundamentally altered, important attributes of the source domain of "jails" – such as having iron bars and being inhabited by criminals – are not mapped.

(2) My job is a jail.

Holyoak & Thagard note that the meaning of "jail" in this metaphor is not the conventional one, but a more abstract meaning of a confining situation.

Moreover, metaphoric expressions can also display emergent structure – that is, implications that don't appear to originate in either the source or the target domain. For example, while neither surgeons nor butchers are customarily considered to be incompetent, a surgeon described metaphorically as a "butcher" by his colleagues does not have a good reputation. Grady, Oakley, & Coulson (1999) argue that incompetence is not a feature of the source domain, and show how the inference of incompetence actually arises from the composition of the butcher's means and the surgeon's ends in a blended space. Because the goals of surgeons and butchers differ so drastically, the butcher's tools and methods are profoundly ill suited for the surgeon's task.

To account for the influence of both the source and target domains in examples such as these, Fauconnier & Turner (1998) have proposed that cross-space mappings in metaphor and analogy recruit the processes of blending. The meaning of metaphoric utterances can be represented in a conceptual integration network. In such a network, the source and target domain each structure one input space. The generic space represents abstract commonalities in the other spaces, and the blended space inherits structure from its inputs, and contains emergent structure of its own. In blending theory, the nature and extent of accommodation of the source domain in a metaphor are related to the conceptual integration network in which the mapping occurs: one-sided, symmetric two-sided, and asymmetric two-sided networks.

6.1.1 Cross-Domain Mapping in One-Sided Networks

In a one-sided network, relations from the source domain are projected into the blended space and bound to elements projected from the target domain. Such a mapping results in little or no accommodation of the source models at the frame level, although, naturally, some accommodation must occur at the specific level. One-sided frame networks capture some of the same data as traditional accounts – cases where the target domain is conceptualized with source domain models. One example of a concept that results from an analogical blend is that of the "computer virus" (see Fauconnier, 1997), computer programs written for the express purpose of damaging other people's computational resources.

Although a computer virus is a very different kind of entity from a biological virus, the two domains share structure and vocabulary. For instance, a computer virus works by attaching itself to a program where it *replicates* and *spreads* to other programs. Fauconnier (1997) notes the generic frame common to both the computer virus and its biological counterpart. This is represented schematically in figure 6.1. As in conventional accounts of analogical mapping, counterparts map to distinct elements whose attributes are licensed by the structure of the inputs. For example, in the case of the computer virus, particular programs in the domain of computers are picked out as counterparts to viruses in the health domain. Moreover, the choice of counterparts is not arbitrary, but motivated by each element's role in shared relational structures.

The array of concepts associated with "computer virus" can be captured in conceptual blending theory with a one-sided frame network. Frames from the domain of biological viruses are evoked in the source

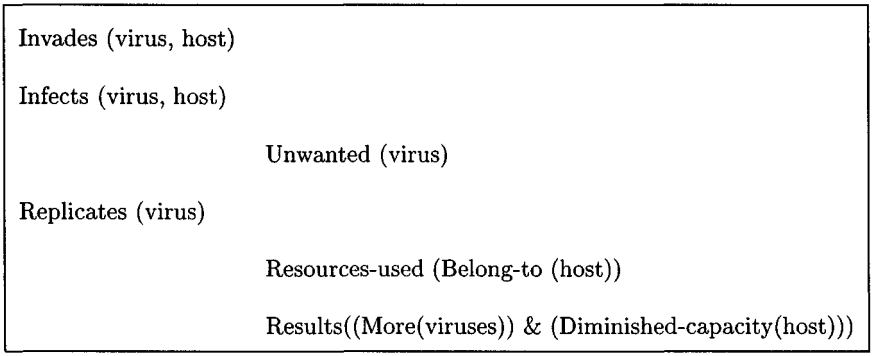

Figure 6.1 Virus Frame. Generic frame common to the computer virus and its biological counterpart.

input, and elements from the target domain of computer viruses are evoked in the target input. Frames from the source are projected into the blended space and filled with elements from the target. The process of composition then allows accommodation at the frame-specific level. Thus we do not expect our computer to pick up a virus from a sick *person*, though we might be wary of using a floppy disk that has been in an infected computer.

6.1.2 Two-Sided Networks

However, more interesting for current purposes are analogical blends that require asymmetric two-sided networks. As noted in Chapter 5, two-sided networks project frame-level information from both input domains into the blended space. Consequently, the blend involves hybrid models that couldn't obtain in either domain, and sometimes models that couldn't obtain at all. The mapping problem is more difficult in two-sided networks than in one-sided ones, and processes of completion and elaboration can yield construals that differ substantially from those available in either input domain. The following examples suggest that speakers strategically exploit these hybrid models to evoke desired target domain construals. Because blends are often exploited for local, rhetorical purposes, the source domain in a metaphor can be chosen for purposes of hyperbole or understatement rather than its abstract similarity to the target domain.

Digging Your Own Grave

Take, for example, the use of the idiom "digging your own grave" in (3):

(3) He's digging his own grave.

In conceptual metaphor theory, idiomatic expressions such as this one have been shown to exhibit systematicity that results from the projection of structured inference schemas from the source domain to the target. In (3), the source domain is that of Grave Digging, and the target is responsibility for a bad personal situation, or Trouble for short. In the source domain, the elements are the digger and the grave; in the target, the elements are the agent and the situation. The digger maps onto the agent, and the grave maps onto the situation. The digger's action in the source (digging) produces a grave; similarly, the agent's action in the target produces a bad situation. Consequently, a person digging his own grave corresponds to an agent actively contributing to his own plight.

At first blush, the mapping appears straightforward. However, there are in fact a number of discrepancies between the structure in the source and the target domains that argue against a two-space account of meaning construction in (3). For example, the grave has a different significance in the input domain of Grave Digging than it does in the idiom. In the idiom, of course, the grave represents the digger's plight. However, in the Grave Digging input, the grave represents the grave digger's goal. A further mismatch concerns the role of the digger in the idiom and his role in the Grave Digging domain. In the idiom, digging your own grave implies actively contributing to your own plight. However, in the source domain, digging is not the *cause* of death, but the consequence. Similarly, the grave digger plays no agentive role in the death of the deceased party.

One might object here, arguing that these discrepancies derive from the fact that the source domain is an atypical case of grave digging: Most people don't dig their own graves. However, note that in the idiom, "digging deeper" maps onto the target as worsening the situation. In the source domain, by contrast, regardless of whose grave you are digging, digging deeper does not in and of itself cause *anyone* to die. I can go out in my backyard and dig a ten-foot grave for my nemesis, but that won't bring him any closer to death. I might even dig a grave for *myself* on Sunday morning and spend the afternoon playing basketball. The basketball might kill me, but the existence of the grave will not. It is, after all, just a hole in the ground.

Nonetheless, some people may still object that the model being recruited in the idiom "digging your own grave" is an even more abnormal case. Rather than the normal case of a grave digger, the idiom recruits a gruesome scenario in which prisoners of war are forced by sadistic guards to dig their own graves. In such a case, it might be argued, the deepness of the grave is correlated with the nearness of death. However, if this scenario is indeed serving as the source domain in this mapping, it is an excellent example of a source domain whose structure is altered in the blend. This outrageous Gestapo scenario does not map well at all onto the target domain in which a person is thought to be blithely muddling her own affairs with no outside provocation.

These discrepancies between source and target domain logic suggest that the mapping between Grave Digging and Trouble is not straightforward, but mediated by a blended space. In this example, the source domain of grave digging contributes elements and relations. However, the causal structure of the blend is a function of the way in which a

cognitive model built with building blocks from the source domain can be mapped onto target domain concepts via conventional mappings. In fact, it's the grave qua hole that accounts for much of the twisted logic of this idiom. A grave is a hole in the ground, and holes are generally bad. A grave is also metonymically associated with death, itself conventionally associated with failure (see Lakoff, 1993). Further, digging a grave involves digging downward, and down is bad (Lakoff & Johnson, 1980; Lakoff, 1987). Digging deeper causes you to be deeper in the hole so that a deep hole in the blend can be mapped onto a worse situation in the target. Someone contributing to the hole by digging is thereby responsible for the bad situation.

Many approaches to cross-domain mapping in cognitive science stress the importance of using a source domain whose causal structure maps appropriately onto the target. And, while this might be an important constraint on analogical reasoning for purposes of problem solving, it is less so for understanding expressions such as (3). The choice of source domain in this idiom does not seem to be guided by the way frames from Grave Digging can be mapped onto Trouble. Rather, the domain of Grave Digging is exploited for other reasons. First, it provides a domain which can – via conventional mappings in the conceptual system – reduce to a simple model what is potentially a long, complex series of actions in the target domain. Second, it hyperbolically conveys the seriousness of the target domain situation. Finally, it makes certain contrast classes available that might not have been otherwise available in the target domain.

6.1.3 Digging Deeper

The blending framework can also account for productive extensions of the "digging your own grave" idiom. For example, picture a discussion between three friends, Larry, Moe, and Curly. Larry, talking about a meal he's recently had, comments, "And this meat was so tough – God, it was worse than your cooking, Moe." When Moe begins to take offense, Larry attempts to backpedal. "Well, not that you can't cook. I mean your frozen pizzas..." At this point, Curly might interject, "Keep digging, Larry," to signal the failure of Larry's attempt to cover for his offensive remark.

As in the conventional idiom, the agent maps onto the digger, the digging maps onto the cause of trouble (in this case, speaking), and the implication is that continuing to dig will only make the situation worse. However, as argued previously, the significance of "digging" in

the expression, "Keep digging," differs from its significance in the source domain. For the real grave digger, digging is just part of the job, but, for Larry, digging will only worsen his situation. The productive character of the expression, then, does not draw upon the inherent structure of the source domain of grave digging, but from the emergent structure of the blend.

The productive character of the Grave Digging/Trouble blend is also apparent in the attested example in (4).

(4) It's not too late to exhume ourselves from the shallow grave we've dug for ourselves.*

This blend is quite spectacular in some respects – especially because the digger and the corpse both map onto the collective agent. In the original blend, Trouble maps onto the Grave, which metonymically evokes Death; however, in the blend constructed to understand (4), Trouble maps onto both the Corpse and the Grave. Further, the two blends differ slightly in how the agent can best avoid Trouble. While, in the conventional blend the cessation of digging maps onto the aversion of Trouble, in (4) averting Trouble involves exhuming the corpse. Moreover, if exhuming the corpse seems more drastic than merely stopping digging, it may be because actors in (4) have already *dug* their graves, while the actor in (3) is still *digging*.

Nonetheless, the blend in (4) resembles the conventional blend in at least one important respect. In both the idiom and the more elaborated blend in (4), the degree of Trouble can be mapped onto the depth of the grave. If someone is digging her own grave, the best course is to stop before getting in "too deep." Further, while it's not too late to get out of our "shallow" grave, if the grave were any "deeper," it might pose problems. Once again, this inference schema does not originate in the realistic domain of graves and corpses, but from the unrealistic, simplified model constructed in the blended space.

Examples such as (3) and (4) suggest that the cross-space mapping in metaphor is not a straightforward one from source to target, but rather is mediated by mappings into the blended space in a conceptual integration network. These examples show that the inferences suggested by metaphoric utterances need not result from projections based on shared relational structure. In this respect, the source domain in a metaphor is

* This statement originally appeared in an editorial in the University of California, San Diego, newspaper *The Guardian* about the plight of the American educational system.

less important than previously thought, as causal structure in the source can be quite irrelevant for the resultant construal of the target domain. Rather, the import of the source input for the resultant target domain construal is more complex.

6.1.4 Born on Third Base

The source domain in a metaphor is properly understood as an *input* to the conceptual integration network. So, while blending processes certainly can and frequently do capitalize on causal and relational structure in the source domain, they need not. In the last examples, we saw how source domain elements are often projected into the blend and linked to the target via conventional mappings (such as the connection between Down and Bad) that were irrelevant in the source domain of Grave Digging. In these examples, speakers recruit blending to capitalize on inferences afforded by conventional mappings between the simplified blended models and the target. Similarly, source domain models are often recruited and transformed by blending in order to evoke culturally shared conceptualizations of the target domain known as *cultural models*.

For example, during the 1992 presidential primary campaign, candidate Tom Harkin criticized the incumbent, President George Bush, with the comment in (5).

(5) He's a guy who was born on third base and thinks he hit a triple.

Most Americans will easily understand Harkin's analogy as pointing out the contrast between the perceived and the actual cause of Bush's success as a politician. While Bush perceives his success to be a function of his intrinsic abilities and hard work, Harkin points to Bush's privileged heritage and implies that his success is undeserved. Harkin thus uses concepts and terminology from the domain of baseball to discuss Bush's success. Moreover, the mapping between the sports domain and the social domain is not one-sided, but combines models from each. Consequently, understanding Harkin's blend requires a certain amount of American cultural knowledge about both input domains.

For one thing, we have to understand the terms "triple" and "third base" in relationship to the institution of baseball. Baseball, of course, requires the batter to hit the ball and run around a series of four bags (bases) arranged in a diamond. Each base is labeled according to the order in which base-runners are required to touch them. The base where

the batter hits the ball and begins base-running is called home plate. The first base that the runner runs toward after hitting the ball is called first base, the next is second, and the next is third base. Upon hitting the ball, the batter will run toward first base while the defensive team attempts to field the ball so as to get the batter out, or, at the very least, to stop his progress around the bases. If the batter is able to get back to home plate without being put out, his team receives one run.

It is customary in baseball to label a batter's hit by the number of bases he was able to tag before his initial progress was stopped. If he gets to first base, it is called a single; if he gets to second base, it is called a double; if he gets to third base, the hit is called a triple; and, if he gets all the way around the bases, his hit is called a home run. Thus a player who hits a triple will end up on third base. While there are many different ways to arrive at third base (e.g., a batter hits a single and is advanced to third by other batters), hitting a triple, by definition, results in the batter standing on third base. Moreover, a batter who hits a triple winds up on third base largely as a result of his own efforts.

As mentioned previously, complete understanding of the meaning of Harkin's statement depends not only on knowledge of baseball, but also on knowledge of the American cultural model of success. As children, Americans are indoctrinated with the notion that any deserving American-born citizen can become president. Another product of American cultural transmission is the Horatio Alger story. This is a story in which the protagonist is born into a poor family, but manages, through hard work and ingenuity, to become a successful businessman in his adulthood. The Horatio Alger story is meant to highlight the virtues of the protagonist as well as those of the American system, which affords his rise from rags to riches.

The vocabulary of success is reminiscent of the vocabulary invoked to describe paid labor. You earn success through hard work, just as you earn money through hard work. Moreover, the two domains share inference schemas. You deserve money that you earn; you don't deserve money that you don't earn. Similarly, you deserve success that you earn; you don't deserve success that you don't earn. Because wealth is one of the main properties of American success (power is the other), there is a metonymic connection between success and money that motivates the extensive array of entrenched mappings between domains of success and paid labor. This sort of metonymic structure mapping has been referred to as a pragmatic function mapping, and often results (as it has here) in shared vocabulary between the two domains.

Further, the familiarity and structure of the frames for paid labor can be fruitfully employed to reason about the more abstract concept of success. In the idealized model of paid labor, a person who works earns money for her labor, and the amount of money earned is proportional to the effort exerted and skills deployed. This model can be used to discuss why a particular person might or might not deserve the money that she has. Moreover, this same model can be mapped into the domain of success, via a mapping between money and success. A person who works achieves success as a result of her labor, and the amount of success achieved is proportional to the effort exerted and the skills deployed.

According to the Idealized Success model, the equality of opportunity reputably engendered by the American politicoeconomic system allows Americans to succeed at a level proportional to their effort and ability. The level of one's success is, then, an indicator of the level of one's effort and ability. It is a function of this cultural model that we evaluate successful people favorably and unsuccessful people unfavorably. Successful people are successful because they earn success through hard work. On the other hand, unsuccessful people are unsuccessful because they are unable or unwilling to earn success.

Of course, this is not the only American cultural model of success. We realize that success is more difficult to achieve for some than for others. Part of the appeal of the Horatio Alger story is that the protagonist is able to advance from the bottom rung of society to the top in the course of a lifetime. The achievement is portrayed as admirable, not only because it indicates hard work on the part of the protagonist, but because the climb is all the more difficult from the bottom than it is from the top.

For this reason, we supplement the American cultural model of success with an alternative model that takes into account the circumstances into which a person is born. A person who comes from a wealthy family has advantages that a person from a poor family does not. Moreover, these advantages allow the children of rich parents to start their pursuit of success from a higher level than the children of poor parents. So, on the Idealized Success model, one's level of success is a prima facie indicator of one's abilities and hard work, on the Silver Spoon model, an individual is not given credit for success gleaned from advantages of birth but only for advancement beyond that level.

Harkin's seemingly simple statement thus appeals to frames from the domain of Baseball, both of the idealized models of success, and a frame that blends the social and the sports domains. The phrase "born on third

Table 6.1: Initial Frame Network

Input	**Input**	**Blend**	**Generic**
Baseball	Success	Baseball-Success	Path
Elements	**Elements**	**Elements**	**Elements**
George Bush	George Bush	George Bush	Agent
Third Base	Wealthy Family	Third Base	Point-a
Relations	**Relations**	**Relations**	**Relations**
On(Bush, third base)	Born(Bush, wealthy family)	Born(Bush, third base)	Start(Agent, Point-a)
			Near(Point-a, Goal)

base" evokes a two-sided frame network with Baseball and Success as input domains. In the domain of Baseball, one can be *on* third base, but one cannot be *born* on third base. Of course, it's physically possible for a baby to be born on third base; however, such an event is not covered by the rules of baseball. If it did occur, it would not be considered part of the baseball game proper, but a bizarre event that occurred during the course of a game. Similarly, in the target domain of Success, one can be born into a wealthy family, but not onto a base. It is only in the blended space of the frame network where it is permissible for Bush to be born on third base. The frame network is outlined in table 6.1.

Further, the connective "and" in (5) links "born on third base" with "thinks he hit a triple" and signals that the two are connected descriptions. The verb "thinks," because it produces an opaque context, sets up a parallel frame network to represent Bush's beliefs. Elements and relations in the Thinks network can be linked to counterparts in the initial frame network in a way that allows speakers to understand contrasts between the actual and the counterfactual state of affairs (see table 6.2). In the generic domain, for example, the agent thinks he's advanced to Point-a, when in fact he started out at Point-a. In the Baseball domain, Bush is on third base and thinks he hit a triple. Note that in the baseball input, there's nothing particularly amiss in being on third base and thinking you've hit a triple. In the Baseball input, Bush is not necessarily deluded as to how he got to third base.

However, in the blended space, Bush most assuredly *is* deluded as to how he got to third. Similarly, in the Success input, Bush is deluded as

Table 6.2: Thinks Network

Input	**Input**	**Blend**	**Generic**
Thinks	Thinks	Thinks	Thinks
Baseball	Success	Baseball-Success	Path
Elements	**Elements**	**Elements**	**Elements**
George Bush	George Bush	George Bush	Agent
Triple	Success	Triple	Point-a
Relations	**Relations**	**Relations**	**Relations**
Hit(Bush, Triple)	Earned(Bush, Success)	Hit(Bush, Triple)	Advanced-to(Agent, Point-a)

to how he earned his success. The rhetorical force of Harkin's statement comes from the contrasting causes of being on third base in the blended space in the two frame networks, and the way in which those causes map into the target domain of Success. The blended domain involves an effect (being on third base) and two competing causes – one which is an actual cause in which Bush was a passive participant, and one which is the counterfactual cause where Bush is an active participant.

The structural relationship between the elements in the blended spaces parallels that constructed between the elements in the Success input spaces. Moreover, it parallels the relationship that exists between the two success models, Idealized Success and Silver Spoon. Bush thinks his success is due to his own hard work and ingenuity. Therefore, he deserves the success he has earned. However, the actual cause of Bush's success derives from the fact that he was born into wealth; his success, then, is neither earned nor deserved. The actual cause of Bush's success derives from the fact that he was born into wealth and appeals to the Silver Spoon model. However, Bush's belief about the cause of his success is governed by the Idealized Success model.

The import of the contrast between the perceived and the actual cause of Bush being on third base derives largely from the contrast between the two success models. We seem to fault Bush, not only for unearned success but for his use of the Idealized Success rather than the Silver Spoon model. In the domain of Baseball, it matters little whether or not the player knows how he got to a particular base. In fact, it's rather implausible for a baseball player *not* to know how he got to a particular base. So while Baseball provides a domain with which Harkin

can exaggerate Bush's delusions of grandeur, the rhetorical force of (5) derives from the sociocultural import of the target domain frames. Nonetheless, the particular configuration of success frames that the listener constructs on hearing (5) is mapped from the models built in the blended spaces.

The models evoked in the blended spaces thus play a critical role in the construal ultimately constructed by the culturally competent listener. The model built in the blended space makes for a better contrast between the actual and the counterfactual spaces than contrasts that are available in the realistic Baseball domain. Moreover, the use of the word "born" in this context evokes the Silver Spoon model and helps to constrain the mappings between the blend and the target. Although (6), (7), and (8) all maintain Harkin's tension between earned and unearned success, none has the same critical force as (5).

(6) He stole third base and thinks he hit a triple.
(7) He was airlifted onto third base and thinks he hit a triple.
(8) They put third base where he happened to be standing and he thinks he hit a triple.

While the model in the blended space is unique, the interpretation of (5) is considerably constrained by conventional mappings. Identification of the target domain is aided in this case because success is often characterized metaphorically as progress along a path. Moreover, political success in particular is often discussed with sports metaphors. Existing pervasive metaphoric connections between source and target domains aid the native speaker in the identification of the appropriate mappings between spaces. Without knowledge of the source input (baseball), it is difficult for the listener to formulate a coherent mapping to the blend or the target domain. Further, without knowledge of the Success models, the relevance and implications of this structural relationship are simply lost on the listener.

The meaning of (5) is constructed through the recognition of particular relationships that exist between culturally constituted and shared cognitive models rather than referencing the outside world directly. Moreover, understanding the statement involves more than a mere allusion to particular cultural models. It involves exploiting particular structural relationships that exist both within components of models and between the different models invoked. It is the apprehension of these relationships and our ability to map them across disparate domains that result in the transfer of inference schemas from one domain to

another.* Importantly, Harkin does not simply exploit preexisting concepts from the domain of Baseball to evoke disparate cultural models of success. Rather he prompts the listener to construct a novel frame in the blended space that is, in turn, mapped onto target domain frames so as to suggest a particular construal of Bush's success.

This represents an important difference between conceptual blending theory and the theory of conceptual metaphor on which it builds. In conceptual metaphor theory, metaphors are seen as instantiations of entrenched mappings between cognitive domains, while in blending, the meaning of a metaphor is constructed on-line in conceptual integration networks. The two-space explanation is extremely useful in explicating commonalities in mappings over a wide range of entrenched metaphorical expressions. However, to account for the implications of particular metaphoric utterances, it's frequently necessary to point to blending in conceptual integration networks.

6.2 RETROSPECTIVE PROJECTIONS

Another difference between conceptual metaphor theory and blending concerns the direction of projections. Because most cognitive scientists have taken problem solving as the paradigmatic case of cross-domain mapping, there has been an undue emphasis on examples in which we have a better understanding of one domain than another. This, in turn, has led to a focus on how inference schemas from the source are exploited in the target domain. Similarly, cognitive semanticists have focused on the computational role of entrenched metaphor, enabling abstract ideas to be conceptualized with easily manipulated image-schematic representations (Johnson, 1987).

However, nothing inherent to cross-domain mapping mandates mapping from the source to the target. The ability to link counterparts across disparate domains based on shared relational structure is a basic human ability. Moreover, the productive character of the integration that occurs in blending is a similarly basic ability that we actively exploit to suit our needs. In many cases, especially humorous ones, people set up blends

* However, the reader should not infer from this that speakers are consciously aware of the complex set of mappings that lurk beneath their understanding of statements such as Harkin's. In fact, the oddity of many exotic blends passes unnoticed by the majority of speakers because we efficiently derive the information we need from the blends.

where the rhetorical goal is to project inferences from the blend to the input that would traditionally be construed as the source domain. This section discusses several examples that suggest it is the *purpose* of the mapping that ultimately determines the direction of the inference transfer (Holyoak et al., 1986).

6.2.1 The Menendez Brothers Virus

Consider the following joke, which alludes to the trial of Erik and Lyle Menendez, two young men who confessed to murdering their parents:

> Menendez Brothers Virus: Eliminates your files, takes the disk space they previously occupied, and then claims it was a victim of physical and sexual abuse on the part of the files it erased.

Understanding this joke depends upon the reader's knowledge of the case, and the controversy surrounding the brothers' legal defense. The Menendez brothers admitted to killing their parents, but claimed their actions were motivated by self-defense against parents who had repeatedly abused them, both physically and sexually, since childhood. While many Americans found the Menendez brothers' case persuasive, others cited it as an example of how the U.S. justice system attributes too much weight to mitigating factors in its treatment of the perpetrators of violent crimes. Skeptics pointed to the multimillion dollar estate the brothers inherited from their parents, and noted that the night Mr. and Mrs. Menendez were gunned down in their living room, they were armed only with bowls of strawberry ice cream.

One way of explaining the process of meaning construction in this joke involves cross-domain mapping between the event scenario associated with the Menendez brothers' trial and the domain of computer viruses. The target domain of computer viruses shares a system of relations with the source domain of the Menendez brothers. Moreover, elements and relations have been placed into correspondence with objects in the source, based on similar roles in their shared relational structure. However, the virus described in the joke has some features – such as claiming to be an abuse victim – that appear to result from blending in a two-sided frame network.

Table 6.3 lists the named elements in the source domain of the Menendez brothers, and their corresponding elements in the target domain of computer viruses. In the source domain of the Menendez

Table 6.3: Elements in Source and Target Spaces

Source	Target
Elements	**Elements**
a′ Menendez Brothers	a Computer Virus
b′ Mr. & Mrs. Menendez	b Affected Files
c′ Property	c Disk Space

brothers, the elements include the brothers, their parents, and the Menendez family property. Similarly, in the target domain of computer viruses, common elements include viruses, affected files, and disk space. The generic space contains the more abstract agents, patients, and an inanimate entity. Finally, a fourth space is set up in which blending of the two input domains may occur.

Table 6.4 lists the corresponding relations that exist between elements in each of the four spaces. I have tried to employ the predicates that are most appropriate to the domain in question. For example, the term "kill" is used to refer to the social source domain, while the term "erase" is used to refer to the technical target. However, the language user is not similarly constrained. As noted previously, once elements in two or more domains have been linked, the Access Principle (Fauconnier, 1994) can be invoked to refer to corresponding elements in either of those spaces. Similarly, predicates that are customarily associated with one particular domain can be applied to refer to counterpart relations in

Table 6.4: Counterpart Relations in Source, Blend, Generic, and Target

Source	Blend	Generic	Target
Relations	**Relations**	**Relations**	**Relations**
Kill (a′,b′)	Erases (aa,bb)	Eliminates (a″,b″)	Erases (a,b)
Acquire (a′,c′)	Takes (aa,cc)	Takes (a″,c″)	Takes (a′,c′)
Owned-previously	Occupied-previously	Occupied-previously	Occupied-previously
(c′,b′)	(cc,bb)	(c″,b″)	(c,b)

linked spaces. Consequently, the text of the joke utilizes terms from the generic domain (e.g., "eliminates," "takes," and "occupied"), perhaps to facilitate mappings to the input spaces.

Interestingly, out of the context of the joke, the predicates that have been equated involve quite disparate concepts. Objectively, killing people and erasing files have almost nothing in common. Acquiring property and taking disk space are markedly distinct actions. Moreover, the concept of property ownership in the social source is a very different thing from occupying disk space in the technical target. However, once the context of the joke has been set up, linking the disparate properties in the two domains proceeds quite naturally.

The blended space itself concerns the Menendez Brothers computer virus and is structured by frames from the target domain (viz., erasing files and taking over disk space) that are shared by counterparts in other spaces. Because the target domain contains a counterpart relation for "killed," "killed" maps onto "erased" in the blend. Although the initial structuring of the blended space is quite congruent with knowledge about computer viruses, the social input projects some structure that has no counterparts in the target domain. Viruses often delete files, occupy disk space, and even have colorful names such as the Menendez Brothers Virus. However, the suggestion that a computer virus could be the victim of physical or sexual abuse is patently absurd.

The notion of an abused computer virus (in the sense intended) is incoherent in the target domain and chimerical in the blend. The possibility of an abused virus arises only in the blended space where it enjoys a short-lived existence, confined to local purposes. Although the Abused (virus) does not get mapped to the target domain of real computer viruses, the links between the source and the blend enable the projection of structure between these two spaces. In particular, the inference that the virus' claim is ridiculous and false gets transferred back to the source domain, where it triggers a similar inference for the Menendez Brothers Virus' source counterparts.

In table 6.5, we can see the projection of an entire event scenario from the source input of the Menendez brothers' murder trial into the imaginary scenario associated with the blended space. Because the event is extended in time, the representation of the source domain includes three mental spaces: one for each of the relevant time periods (TIME 0, TIME 1, and TIME 2). The base space in the source domain is the time at which the Menendez brothers have been accused of murder (TIME 2).

Table 6.5: Parallel Event Scenarios in the Source and the Blend Domains. Each domain contains three mental spaces that correspond to three events, occurring at three different times. The earliest event (at TIME 0) concerns the Menendez brothers' childhood, and is represented in the elements and relations of the Claim space. The latest event (at TIME 2) is in the courtroom, and is represented in the Base space. The event represented in the Crime space occurred between TIME 0 and TIME 2.

Source Event Scenarios	**Blend Event Scenarios**
Base TIME 2	**Base TIME 2**
Elements	**Elements**
a' Menendez Brothers	aa Menendez Brothers Virus
Accused (a')	Accused (aa)
Claim TIME 0	**Claim TIME 0**
Elements	**Elements**
a' Menendez Brothers	aa Menendez Brothers Virus
b' Mr. & Mrs. Menendez	bb Affected Files
Relations	**Relations**
Abused (b',a')	Abused (bb,aa)
Victims (a')	Victim (a)
Crime TIME 1	**Crime TIME 1**
Elements	**Elements**
a' Menendez Brothers	aa Menendez Brothers Virus
b' Mr. & Mrs. Menendez	bb Affected Files
Relations	**Relations**
Killed (a,b)	Erased (aa,bb)

At this time, they offer a claim about the events of their childhood. The content of this claim is represented in the structure of the claim space that occurs at TIME 0. Moreover, the content of the claim space is offered as an explanation of the murder of Mr. and Mrs. Menendez,

which occurred at TIME 1, after the purported child abuse and before the accusation of murder. The entirety of this event scenario gets mapped onto the blended domain.

Partial structures from the source and from the target have been integrated into a single structure in the blended domain, and integration results in the local creation of a novel sort of computer virus. In the blended joke space, computer viruses can make excuses for their behavior – something that neither computer nor biological viruses ever do! Moreover, the links between the component spaces in a frame network allow us to capture an important aspect of the meaning of this joke that would be overlooked by conventional cognitive science approaches to cross-domain mapping. The rhetorical topic of this joke is not the target domain of computer viruses, but the social source. Although the elements and relations of the target domain have been placed into correspondence with elements and relations in the source, the mapping does not result in productive inferences about computer viruses. Rather, the mappings employed in the Menendez Brothers Virus joke serve to highlight a controversial construal of the *source* domain of the Menendez brothers' trial.

6.2.2 The Menendez Brothers Virus II

In fact, metaphoric and analogical blends are often set up to promote or accentuate particular construals of the input spaces. For instance, in the Menendez Brothers Virus joke, the blended space recruits a framing of the source domain that accentuates the degree to which the Menendez brothers profited from their actions. Recall that the framing of the (real) Menendez brothers' actions in the source space was quite controversial. Besides being televised on Court TV, the brothers' trial was the subject of countless news reports, several books, and three made-for-TV movies. Moreover, the outcome of the first trial was a hung jury. At the center of the controversy was the issue of who played the role of the evil agents, and who the role of the innocent victims. The following are two contrasting framings of the Menendez brothers' affair: the Conspiracy Framing and the Victim Framing.

Conspiracy Framing
Brothers plan to acquire parents' possessions by committing homicide.
Brothers shoot parents and inherit their possessions.
When arrested, the brothers claim long-term physical and sexual abuse as a mitigating factor.

Victim Framing
Kitty and Jose Menendez physically, sexually, and psychologically abuse their two sons.
Brothers sense imminent escalation of parents' abuse.
Brothers shoot parents in self-defense.

To appreciate the extent to which the Menendez Brothers Virus (MBV) joke involves a controversial framing of the source domain, one need only contrast the sequence of events stated explicitly in the joke with other possible sequences, given the same sort of mapping. The sequence in the MBV joke involves (1) eliminating files; (2) taking disk space; and (3) claiming physical and sexual abuse. However, consider the following alternative joke, which employs the same mappings as the Menendez Brothers Virus joke presented in a different order:

> Menendez Brothers Virus II: Suffers many years of physical and sexual abuse from the .exe files on your hard drive; finally decides to get revenge and escape abuse by deleting the offending files.

In this example, the elements in the two input domains have been mapped in the same way as in the original: The brothers map onto the virus, the parents map onto the files, and the physical and sexual abuse have no counterparts in the computer virus domain. Moreover, as in the original joke, the murder maps onto the deletion of files. However, the alteration of the narrative structure of the joke results in a fairly dramatic change in its implications. In the original scenario, the claim is an excuse whose validity was questionable. However, the same claim in the latter scenario serves as a justification. In the original scenario, the brothers are framed as agents motivated by greed, but in this scenario the brothers are victims of their circumstances, and motivated by fear, revenge, and self-defense.

In both scenarios, the structure in the blended space seems to influence construal of the social input domain. In the original joke, the structure in the blend reinforces the conspiracy framing from the source with background knowledge imported from the target domain of computer viruses. Because the coherence of the blended domain ultimately depends on congruity between the particular framing imported from the source and the logic of the target domain, the construal of computer viruses as *agentive* reinforces the Conspiracy Framing while undermining the Victim Framing. So, while the structure of the blended space built to understand the original joke serves to promote mappings from

the blend to the source input, the structure of the blend built to understand the second joke provides a context that makes the imported Victim Framing appear implausible.

6.2.3 The Right to Life Virus

Unfortunately, the cognitive import of cuing the construal of events such as the Menendez brothers' trial has received little attention from cognitive scientists. However, as we will see in Part III, cognitive models play an important role in determining the character of experience due to the way they structure our expectations and cue the appropriate affective response. While this aspect of comprehension differs from the traditional notion of the computational role of metaphor in cognitive science, it is nonetheless important for explaining human behavior. The Menendez Brothers Virus discussed in the last section concerns a case where the blend is used to help settle a controversial issue. Similarly, in the following example, the blend exploits noncontroversial aspects of one domain to alter construal of another, more controversial, domain.

> Right to Life Virus: Won't allow you to delete a file, regardless of how old it is. If you attempt to erase a file, it requires that you first see a counselor about possible alternatives.

Meaning construction in this example involves conceptual blending of elements from the social domain of reproductive rights and the technological domain of computer use. In particular, frames concerning women's reproductive rights have been blended with frames for the relationship between a computer user and her electronic files. This section demonstrates how conceptual blending has been used to capitalize on the availability of an uncontroversial framing in the technological domain to map it onto its more controversial social counterpart. The model built in the blended space serves as an intermediary.

In fact, blending in this example begins with the title that couples the term "Virus" with "Right to Life," the name of the American antiabortion movement. Because the name of a biological virus is often the name of the disease that the virus transmits, the title serves to frame Right to Life activism negatively as a disease caused by the virus. Moreover, the name of a computer virus often serves as a cue to do conceptual blending between a physical or social source domain and analogous features of the technical target. Consequently, the reader is prepared to map frames associated with Right to Life activism onto frames associated

Table 6.6: Elements in the Right to Life Virus Mapping

Source	Generic	Blend	Target
Elements	**Elements**	**Elements**	**Elements**
Right-to-Lifers	Antagonist	RLV	Computer Virus
Woman	Protagonist	You	User
Fetus	Disputed-Object	File	Files

with computer viruses. The precise nature of the mapping is made explicit in the description following the title, which evokes elements and relations from both inputs.

Tables 6.6 and 6.7 list the elements and relations evoked by the first phrase, "won't allow you to delete a file." The first mapping is a case of simple composition, in which elements and relations in the social input have been matched with counterpart elements and relations in the technical target. That is, -Allow (Right-to-Lifers, (Abort (Pregnant-Woman, Fetus))) has been matched with Prevent (Computer-Virus, (Delete (User, Files))). The match between Right-to-Lifers and Computer Virus is cued by the name Right to Life Virus (RLV). The match between aborting a fetus and deleting files is facilitated by the conventional mapping between killing and deleting and/or erasing. This conventional mapping is evident as gangsters talk about "rubbing out" their rivals or "taking them out of the picture." Similarly, computer users frequently refer to erasing files as "killing" them.

The Right to Life Virus is another example in which meaning construction results in a concept that is not congruent with target domain

Table 6.7: Relations in the Right to Life Virus Mapping

Elements	Relations
Source	-Allow(Right-to-Lifers, (Abort (Woman, Fetus))) Irrelevant(How-Young (Fetus))
Generic	Prevent(Antagonist, (Eliminate (Protagonist, Disputed-Object))) Irrelevant(Age-of (Disputed-Object))
Blend	-Allow(RLV, (Delete (You, File))) Irrelevant(How-Old (File))
Target	Prevent(ComputerVirus (Delete (User,Files))) Irrelevant(Age-of (Files))

Table 6.8: Elements in the Right to Life Virus Mapping

Source	Generic	Blend	Target
Elements	**Elements**	**Elements**	**Elements**
Woman	Protagonist	You	User
RtL-Legislation	Antagonist	RLV	Computer Virus
Fetus	Object	File	Files
Counselor	—	Counselor	—
Alternatives-to-Abortion	Alternatives-to-Elimination	Alternatives-to-Deletion	Alternatives-to-Deletion

expectations about computer viruses. For instance, it is generally the role of the virus – not the user – to delete files. Moreover, the default assumption in the virus domain, that the user is *worried* about the virus deleting her files, conflicts with the assumption in the joke that the user actively wants to delete her files. From the perspective of the technical target, it is quite odd for a virus to prevent a user from deleting her files. But, as noted previously, the Right to Life Virus' activities can be readily interpreted as the result of mapping structure from the source domain of reproductive rights. Mappings between elements from the various spaces in the Right to Life Virus blend are outlined in table 6.8.

Although superficially the mappings in the social and technical domains are quite symmetric, this symmetry obscures important differences in the way each predicate is embedded in background knowledge. These include differences in the frames to which each predicate connects, differences in the relative prominence of connected topics, and differences in the controversy surrounding frames in the social source versus their counterparts in the technical target. The importance of the location of a particular model with respect to background knowledge can be seen in the contrast between the implications of seemingly similar frames for Delete(User, Files) and Delete(Virus, Files).

The default preconditions for the Delete(User, Files) frame involve a scenario in which the user *wants* to delete the files and has a sensible, beneficial reason for doing so. However, Delete(Virus, Files) evokes a scenario in which the deletion of files is harmful and contrary to the wishes of the user.

(17) That virus overwrote dozens of my files – luckily, no important ones.

(18) # That virus overwrote dozens of my files – luckily, some important ones.
(19) That virus overwrote dozens of my files, and it upset me.
(20) # That virus overwrote dozens of my files, but it upset me.

Thus the difference in the preconditions for Delete(User, Files) and Delete(Virus, Files) naturally leads to the differences in the implications of each, including the construal of the virus as harmful.

Similarly, the apparent symmetry in the mapping between Abort (Woman, Fetus) and Delete(User, Files) obscures important differences in the way each predicate is embedded in background knowledge. In the technical domain, it is generally agreed that the relationship between a user and her files is one of ownership. Files exist to be used by the user and they may be created and destroyed as the user sees fit. It can be immoral to delete someone else's files, but not (ceteris paribus) to delete one's own. However, in the social domain, the nature of the relationship between a woman and her fetus is a controversial question that bears on the morality of abortion. While Right to Life activists endorse a construal that profiles the natural end point of pregnancy and emphasizes the mother's role as caretaker for the child, Pro-Choice activists endorse a construal that profiles early stages of pregnancy and emphasizes a woman's right to control her own body.

Several aspects of meaning construction in the RLV joke rely crucially on frames evoked in the blended space. As noted previously, the source domain of reproductive rights includes competing frames for construing the relationship between a woman and her body. However, in the technical domain of computer use, the relationship between a user and her files is far less controversial and less important with respect to the morality of deleting files. Although the user's right to remove her own files is not normally a salient aspect of the technical target, the structure recruited from the reproductive rights domain allows it to assume enhanced prominence in the blend. Completion of the blended space evokes the user's right to remove her own files, which can be projected back onto the source domain to evoke the woman's right to abort. Thus blending the controversial Pro-Choice construal of the normative relationship between a woman and her body with the target domain construal of the relationship between a user and her files serves to resolve (temporarily) some of the controversy inherent to the reproductive rights domain.

Another interesting facet of meaning construction in the RLV joke involves the qualification on when it is most permissible to delete a file: "Won't allow you to delete a file, regardless of how old it is." In the domain of reproductive rights, it is fairly clear that the less developed the fetus, the better the case for abortion. For example, compare (21) to (22), and (23) to (24).

(21) Those Right-to-Lifers disapprove of abortion, even in the first trimester.
(22) # Those Right-to-Lifers disapprove of abortion, even in the last trimester.
(23) Charlie approves of abortion in the first trimester, but not in the last.
(24) # Charlie approves of abortion in the last trimester, but not in the first.

However, in the domain of computer use, the older the files, the better the case for deletion (c.f. (25) and (26)).

(25) His boss won't let him delete anything, even the oldest files.
(26) # His boss won't let him delete anything, even the newest files.

Although we can envision situations that legitimate the deletion of either old or new files, the pattern of data in (25) and (26) suggests the operation of an idealized cognitive model in which files have a lifespan of usefulness. While most of the structure in the Right to Life Virus is projected from the reproductive rights domain, the blend inherits the lifespan model from the domain of computer use. This suggests that while the structure of the blended space is largely determined by the mappings from the reproductive rights domain, the emergent structure is determined with respect to the logic of the technical domain.

Nonetheless, the Right to Life Virus displays some properties that are decidedly unlike a virus. The term "allow," for example, suggests a certain intentionality appropriate for the social domain, inappropriate for the technical target, yet seemingly permissible in the blend. The Right to Life Virus also seems to have inherited motives from its counterpart in the reproductive rights domain. In the source domain, the proscription on abortion is motivated by moral misgivings about the termination of the fetus. In the technical domain, computer viruses do not entertain moral misgivings. However, there is a sense in which this particular computer virus actually does have moral misgivings regarding the

Table 6.9: Relations in the Right to Life Virus Mapping

Source	**Generic**	**Blend**	**Target**
Conditional	**Conditional**	**Conditional**	**Conditional**
Attempt-to-Abort (Woman, Fetus)	Attempt (Protagonist, Eliminate (Protagonist, Object))	Attempt (You, Delete (You, File))	Attempt (User, Delete (User, File))
Expansion	**Expansion**	**Expansion**	**Expansion**
Requires (RtL-Legislation, Consult (Woman, Counselor, Alternatives-to-Abortion)	Requires (Antagonist, Action (Protagonist))	Requires (RLV, Ask (User, Counselor Alternatives-to-File-Deletion)	Requires (RLV, Seek (User, Alternatives-to-File-Deletion))

user's deletion of files. The anthropomorphic features of the blended virus are rhetorically exploited to negatively construe the Right to Life movement.

As the joke progresses, the asymmetric nature of the projections becomes more pronounced. For example, take the phrase, "If you attempt to erase a file, it requires that you first see a counselor about possible alternatives." This is not at all congruent with target domain expectations, although it can be understood as a blend of a frame from the reproductive rights domain with the structure that has already been set up in the blended space. In particular, it alludes to controversial state laws that require pregnant women who request an abortion to seek the advice of a counselor. The mappings are outlined in table 6.9. If we elaborate the blend to evaluate the moral status of deleting a file, we generate the inference that the RLV is interfering unnecessarily in the affairs of the user and project an analogous framing back onto the source domain Right-to-Lifer.

Overall, the way this blend works is to recruit structure from the target domain concerning the relationship between a user and her computer files. Moreover, given the way the mappings have been established, this structure can be projected onto a preexisting framing available from the source domain, namely the Pro-Choice construal of the normative relationship between a woman and her body. So while retrospective projection in the Menendez Brothers Virus largely involves incongruencies

between input domains, the Right to Life Virus more directly exploits the framings available uncontroversially in the technical target and projects those framings into the source.

6.2.4 Heidi Fleiss and Her Counterpart

But why do such asymmetries in the direction of projections arise in the first place? One suggestion, and one with considerable empirical support, is that models are more likely to be projected from a well-understood domain onto a less-understood domain. However, given the rhetorical motivation for examples in which inferences are projected to the apparent source, another determinant may concern the ambiguity of the culturally available framings for elements in each input domain. This is certainly the case in the following example concerning the trial of Heidi Fleiss for tax evasion.

Fleiss, better known to tabloid readers as the Hollywood madam, was arrested in 1993 for running an upscale prostitution ring that catered to Hollywood celebrities. The case caught the attention of the press (and the public) when Fleiss threatened to publish the names of her male clients. When Fleiss' case finally came to trial in 1994, it received nightly coverage on the evening news. Trial coverage at one cable TV station even included expert commentary by a former madam, Sydney Biddle Barrows. Because Barrows had been convicted of running a prostitution ring that catered to East Coast socialites, the parallels between Fleiss' 1994 predicament and Barrows' past provided fodder for a host of jokes, including one by Barrows herself. In the following example, Barrows explains the logic behind the cable station's decision to choose her as an expert commentator.

> How does the Heidi Fleiss trial resemble the Gulf War? Ask former celebrity madam, Sydney Biddle Barrows. The socialite and sometime lecturer on relationships will cover the trial of her reputed counterpart from L.A. "If you want an expert on war, you get a retired general," says Barrows who will provide commentary for the cable network America's Talking, "I'm not exactly a general, but I am retired."*

As the columnist's introductory question hints, Barrows' statement sets up an analogy between television coverage of the Gulf War and coverage of the Heidi Fleiss trial. Table 6.10 outlines the elements in

* Gina Bellafante. (1994). Seen and heard. *Time* 144 (21): 125.

Table 6.10: Elements in the General to Madam Mapping

Source	Generic	Blend	Target
Network	Network	Network	Network
General	Expert	Madam	Barrows
War	Topic	Trial	Fleiss Trial

the four domains. The source input is the glut of expert commentators employed by the networks to provide commentary on the Gulf War. In an eight-month study of media coverage of the Gulf War, Steel (1992) counted over 188 different experts, a large majority of whom were retired military officials. Perhaps more telling as to public perception of the prerequisites for the job of expert commentator was a *New Yorker* cartoon depicting a retired general's wife saying, "Oh he's pretty depressed – he's the only retired general the networks haven't called."

Presumably, Barrows' rhetorical goal is to legitimate her role as expert commentator for the Fleiss trial. At one level, Barrows' statement maps the qualifications of an expert commentator for war onto the qualifications of an expert commentator for the Fleiss trial by evoking frames in the Generic space (see table 6.11). However, the ultimate effect of Barrows' analogy is to underscore differences in the social construal of frames evoked in the source and target domains.

Barrows, in (9), evokes a symmetric two-sided frame network in which the source input is military expertise, and the target input is expertise on the Heidi Fleiss trial.

(9) If you want an expert on war, you get a retired general.

Table 6.11: Relations in the General to Madam Mapping

Elements	Conditional	Expansion
Source	Want (Network, Expert (War))	Get (Network, Retired (General))
Generic	Want (Network, Expert (Topic))	Get (Network, Retired (Leader))
Blend	Want (Network, Expert (Soliciting))	Get (Network, Retired (Madam))
Target	Want (Network, Expert (Fleiss-Trial))	Get (Network, Barrows)

Table 6.12: Conditional Spaces in Barrows' Analogy

Input1	**Blend**	**Input2**	**Generic**
Gulf War	Blend	Fleiss Trial	Generic
Condition	Condition	Condition	Condition
Elements	**Elements**	**Elements**	**Elements**
Network	Network	Network	Network
Expert	Expert	Expert	Expert
War	Prostitution	Prostitution	Topic
Relations	**Relations**	**Relations**	**Relations**
Want (Network, Expert)	Want (Network, Expert)	Want (Network, Expert)	Want (Network, Expert)
Expert-on (Expert, War)	Expert-on (Expert, Prostitution)	Expert-on (Expert, Fleiss Trial)	Expert-on (Expert, Topic)

The source input projects structural relations into the blend, while the target input projects elements needed to fill the slots. Because Barrows' statement employs a conditional, its representation includes two spaces for each domain in the network: a condition space and an expansion space. Table 6.12 outlines the mappings between elements and relations in each of the conditional spaces, while table 6.13 outlines mappings between elements and relations in the expansion spaces. Example (9) prompts the listener to structure the Gulf War input, and invites the listener to evoke similar structure in the other spaces of the network.

Barrows' goal is to evoke structure in the generic space and indicate abstract commonalities between her own qualifications and those of a salient type of expert commentator, the retired military official. This occurs via completion of the expertise frame, where the relationship between the leader's experience and the expert's topic is made explicit. The conceptual structure added by completion is outlined in table 6.14. Part of the humor in this example is due to the fact that the frames evoked in the source and in the target are structurally equivalent while their sociocultural significance is quite different. Despite their formal similarity, the social implications accruing to actions that yield expertise in the military domain differ from those that yield expertise in the domain of prostitution.

Table 6.13: Expansion Spaces in Barrows' Analogy

Input1	**Blend**	**Input2**	**Generic**
Gulf War	Blend	Fleiss Trial	Generic
Expansion	Expansion	Expansion	Expansion
Elements	**Elements**	**Elements**	**Elements**
Network	Network	Network	Network
General	Madam	Madam	Leader
Relations	**Relations**	**Relations**	**Relations**
Get (Network, General)	Get (Network, Madam)	Get (Network, Barrows)	Get (Network, Leader)
Retired (General)	Retired (Madam)	Retired (Barrows)	Retired (Leader)

Barrows acknowledges this tension with the statement in (10):

(10) I'm not exactly a general, but I am retired.

Clearly Barrows' statement in (10) does more than negate the proposition that she's a high-ranking officer in the military. After all, this is an analogy and not meant to imply that Barrows could provide expert commentary on military issues. Rather, Barrows' comment suggests that her analogy has somewhat of an undesired effect in that it serves to insult the generals of the source domain.

This brings up an aspect of analogical reasoning that is often ignored by cognitive scientists, perhaps because it is not easily formalized: value. Barrows, in this example, is attempting to project the respectability associated with being a retired general. While cognitive modelers are well aware that the use of analogy is often motivated by ephemeral things like value (see, e.g., eloquent discussion in Holyoak & Thagard, 1995), there are no computational models of analogy that address the issue seriously.

Table 6.14: Conceptual Structure from Completion of Barrows' Analogy

Gulf War	**Blend**	**Fleiss Trial**	**Generic**
Leader-in (General, War)	Leader-in (Madam, Prostitution)	Leader-in (Barrows, Prostitution)	Leader-in (Leader, Topic)

By contrast, rhetorical theorists have seen the transfer of value as a principal component of analogical reasoning. Perelman & Olbrechts-Tyteca (1969: 381) write:

> Although analogy is reasoning that deals with relations existing within the phoros and within the theme, what distinguishes analogy fundamentally from simple mathematical proportion is that in analogy the nature of the terms is never a matter of indifference. For the effect of an analogy is to bring the terms A and C and B and D closer together, which leads to an interaction and, more specifically, to increasing or decreasing the value of the terms of the theme.

Perelman & Olbrechts-Tyteca (1969) note that the change in value due to the interaction of domains in an analogical comparison can sometimes produce a comic effect. Perelman & Olbrechts-Tyteca's example is quite similar to Barrows' in that it involves a comparison between a domain construed in a positive light (dessert of royalty) and one culturally construed in a negative light (dessert of punishment):

> "Brave! brave, by heaven!" cried my uncle Toby, "he [King William] deserves a crown."
>
> "As richly, as a thief a halter," shouted Trim [the faithful corporal].*

In this example, the punishment frame that has been evoked in the source domain has the contemptibility of the recipient as a salient aspect. The humor is motivated by the contrast between dessert based on William's regal nature versus the dessert rooted in a thief's contemptibility.

Similarly, in the Barrows' case, one domain is positively construed and another is negatively construed. Interestingly, these construals stem not so much from judgment of the actions themselves, but involve the whole web of cultural models that surround them. Prima facie, killing people is less acceptable than having sex; but, it is far more respectable to be a retired general than a former madam. Moreover, these differences in the cultural construal of the various frames create the possibility of projecting structure from the blended space back onto the source. Negative cultural framing of prostitution causes the positive properties surrounding being a general to slip into the more negative properties surrounding being a madam.

Barrows' use of the phrase "not exactly" does a lot of the rhetorical work here. The term "exactly" is used here as a negative polarity item,

* Sterne, *The Life and Opinions of Tristram Shandy*, Book VIII, chapter 19, p. 517. New York: Modern Library, Inc., cited in Perelman & Olbrechts-Tyteca, 1969: 378.

as evidenced by the fact that it is good under sentential negation (11), but not in simple affirmative sentences such as (12) (following Israel, 1996).

(11) I'm not exactly a general.
(12) * I'm exactly a general.

Further, it occurs in monotone decreasing contexts* such as (13):

(13) Few of the students were exactly intelligent.
(14) * A few of the students were exactly intelligent.

When used as a negative polarity item, the term "exactly" is what Israel (1998) calls a *high scalar understater*. As discussed in Chapter 9, Israel defines understatement with respect to a scalar model: A sentence is understating just in case it expresses a proposition that is weaker than another proposition available in the context. However, in this particular context, it's unclear what is being understated – Barrows' qualifications, Barrows' social status, or a variety of other things.

Moreover, because the General frame is being recruited for its positive evaluative properties, hedging in such a case tends to invite a contrary inference as in (15).

(15) He didn't exactly love the movie. (hated it)
(16) He didn't exactly hate the movie. (just disliked it)

Thus Barrows' denial that she is a general invites an underspecified contrary inference.

6.2.5 Summary

Where conventional accounts of analogical mapping emphasize the unidirectional mapping of structure from the source domain to the target, conceptual blending theory attempts to provide a more comprehensive account of the links between models that are retrieved from long-term memory and models that are locally constructed (and subsequently discarded) in the course of meaning construction. Importantly, the meaning

* A monotone decreasing expression is a quantified expression in which assertions about the superset entail assertion about its subsets. Expressions that promote monotone decreasing contexts include "no," "neither," "not one," "at most 1, 2, 3, etc.," "only 1, 2, 3, etc.," and "few." So, if at most half the team showed up for the race, at most half the team ran the race in under an hour.

of an utterance that involves conceptual blending should not be thought of as residing in the blended space. Rather, it involves understanding the relationships within and between the various spaces in the frame network. From the perspective of conceptual blending theory, the source and the target spaces are both inputs to the blended space, and neither is inherently more important.

Moreover, the metaphors and analogies discussed in this section reinforce the constructivist approach to language, in which utterances do not simply encode referential relationships but cue the construction of models with which to construe the ongoing communicative act. Metaphoric blends serve rhetorical goals in a number of ways. For instance, blending provides a way of changing the relative prominence of a particular frame, of making a connected frame relevant in a way that it was not relevant in the source domain, and of skirting the need to address controversy in a particular domain. Although the hyperbolic remark does not convey an objective truth, it might lead to an appropriately motivating construal. Insofar as this is the rhetorical goal, this utterance would be said to succeed.

6.3 CONCLUSION

Historically, metaphor has been marginalized and treated as a peripheral phenomenon of little significance either to semantics or for cognition more generally. In traditional approaches to language, the meaning of a metaphor is reduced to a literal statement of the resemblance between source and target. This view grows out of the more general assumptions of truth-conditional and objectivist semantics discussed briefly in Chapter 1, that sentences have meaning in virtue of correspondence to the world via their truth conditions. Because metaphors often express things that could not possibly be true, their meaning must be reducible to literal statements that are either true or false.

In truth-conditional approaches, the peculiar character of some metaphors is attributed to the use of language in an uncharacteristic nonliteral, and consequently surprising and pleasing, manner. However, such an account of metaphor vastly underestimates the subtlety of this phenomenon. It provides no explanation of the variability in people's judgments of the aptness of metaphors, nor for why different metaphors that point to the same literal similarity can have slightly different meanings and implications. What's more, traditional approaches provide no account of the systematicity evidenced in metaphorical

expressions, both at a coarse level of metaphorical clusters (e.g., "argument is war") and at the finer level of systematic mappings that linguists have noted between source and target domain elements (see Lakoff, 1993, for review).

In contrast to truth-conditional approaches, I have drawn from the account of metaphor offered by the theoretical perspective of cognitive semantics. In cognitive accounts (e.g., Fauconnier, 1994; Lakoff, 1987), metaphoric language reflects the operation of mapping mechanisms rooted in the underlying conceptual system. By subsuming metaphor under the rubric of cross-domain mappings, the cognitive approach can explain data that truth-conditional approaches cannot. While on truth-conditional approaches the systematicity of metaphoric language is a happy coincidence, on cognitive approaches it arises naturally out of regularities in the conceptual system. Moreover, the cognitive semantics approach allows us to unify our understanding of metaphor with that of meaning and inference more generally. Understood as a subset of phenomena that exhibit cross-domain mapping, metaphoric language can be seen as demonstrating the productive character of mapping and blending as central cognitive processes.

The cognitive semantics approach to language suggests linguistic data are an important source of information about the way in which cognitive models are organized, accessed, and constructed in the course of meaning construction. However, in their zeal to demonstrate the relevance of metaphoric language to general issues in cognitive science, cognitive semanticists may have unwittingly adopted some unwarranted biases. In particular, existing accounts of cross-domain mapping in cognitive science emphasize its computational aspects and its role in problem solving. Although computational efficiency is certainly a concern in cross-domain mapping, it is not the only pressure that constrains the conceptual system. Cognitive semanticists need not make the same assumptions that have limited other cognitive scientists interested in conceptual structure.

Rather than dismissing metaphor as mere rhetoric, we might do well to approach issues in cross-domain mapping from the perspective of the rhetorical theorist. As I have argued elsewhere (e.g., Coulson & Flor, 1993), rhetorical theory is based largely on empirical evidence as to the persuasive effect of various argumentative techniques. Although we speak of "mere" rhetoric, and consider tropes to be ornamental language, tropes are maintained because of their efficacy for altering other

people's belief structures and motivating various courses of action. As such, the observations of the rhetorical theorist are quite relevant to the cognitive scientist interested in meaning construction. Moreover, because the rhetorical theorist does not share the same computational biases as the cognitive scientist, we can look to her to provide some insights into cross-domain mapping that may have received little notice in cognitive science.

In fact, one of the central ideas of conceptual blending theory was prefigured by rhetorical theorists' treatment of analogy. Perelman & Olbrechts-Tyteca (1969), for example, note that the phoros (source) is often modified in the course of an analogy in order to facilitate parallels between it and the target. Similarly, Black (1962) notes that the concept of "wolf" employed in (27) is an anthropomorphized wolf.

(27) Man is a wolf.

Observations of this sort led Black to argue that the two domains in a metaphor "interact." Although Black's interactionist theory has been criticized for being too vague, even his critics admit his treatment of metaphor (Black, 1962, 1991) contains many insightful observations (Indurkjya, 1991).

Perelman & Olbrechts-Tyteca (1969) note that elements in the phoros often get their significance not from the phoros itself, but by virtue of their relationship to the theme. Similarly, many of the implications of the idiom "digging his own grave" derive from the significance of elements that the blended model inherits from the target domain rather than the source. Likewise, although the blend "born on third base" was implausible from a baseball standpoint, it mapped nicely onto the target domain cultural models of success.

Perelman & Olbrechts-Tyteca also point to the nontrivial modification undergone by phoros elements in the comparison process, discussing how terms in the phoros can be endowed with imaginary properties that emphasize its parallels with the theme. Moreover, they even note the bidirectional possibilities of these sorts of meaning construction, which I have dubbed retrospective projections (Perelman & Olbrechts-Tyteca, 1969: 378):

> In the interaction which occurs in analogy, action on the theme is the more pronounced, but, as we have seen, action in the opposite direction is by no means unimportant. Interaction – in either direction – manifests itself structurally and through transfers of value deriving from the structure.

Although cognitive semanticists have always acknowledged the import of experiential aspects of cognitive (especially metaphorical) models, overemphasis on representational and computational aspects of metaphoric models may have led us to overlook their functional role.

6.3.1 Functions of a Blend

Cross-domain mappings are often touted by cognitive scientists as generating novel inferences about a target domain. However, the point of the mappings in many of the examples discussed in this chapter is not to provide structure for the target domain – as in most cognitive accounts of analogical mapping – but rather to highlight one particular construal of the source domain, perhaps over competing construals. The current perspective suggests a new way of thinking about what constitutes a novel inference. Because the mapping operation involves integrated frames rather than isolated predicates, the choice of one particular framing over another necessarily results in a different set of attendant inferences. Besides the acquisition of unknown facts, a novel inference might involve a new construal of a well-understood phenomenon, a change in prominence of a particular element, or simply the availability of connected frames.

Metaphor and analogy are often motivated by rhetorical goals such as the need to suggest correct inferences, elicit affective responses, or elicit a construal that is appropriately motivating. One way to enhance a particular construal is to employ a blend that changes the salience of an element in the original input domain, as in the Harkin "triple" example. The change in salience can occur because of the way in which the elements in the blend map onto target domain models. Alternatively, salience changes can occur because the blend affords the use of a frame that profiles another aspect of the scene.

Another function of a blend can be to make certain contrast classes available, either from inherent properties of input domains or from emergent properties of the blend. For example, in the "digging your own grave" blend, the depth of the grave is mapped onto the degree of trouble. Moreover, as we saw in the Right to Life Virus example, if one of the inputs connects to different frames, it enables speakers to bring new structure into the blend and project it to the other input. Finally, the event integration that occurs in conceptual blending can also be put to rhetorical use, since presenting something as an integrated scenario can often enhance a particular construal of some event.

Blends that recruit cultural models are particularly effective because of their capacity to evoke an associated affective response, and because they seem to have motivational potential. For example, in the "argument is war" metaphor, war provides a culturally framed domain that can convey the significance of various events in an argument. Although "war" is hyperbolic as a source domain, its function in the blend might be largely for its affective capacities. Similarly, while "death" is hyperbolic in the "digging your own grave" idiom, this hyperbole might plausibly subserve motivational ends.

Blends can also be used to capitalize on culturally framed domains so as to make a controversial frame appear less controversial by blending. In the Right to Life Virus example, analogical connectors were used to exploit an unambiguous framing of the relationship between a computer user and her files in order to project an analogous framing onto the more controversial relationship between a woman and her unborn child. Similarly, in the Sydney Biddle Barrows example, we saw how the amount of framing ambiguity can be a major determinant in the direction of structure generated in the blend. In this example, the unambiguous cultural framing of prostitution was inadvertently projected retrospectively onto the retired general.

The rhetorical perspective on cross-domain mapping suggests that, besides creating an understanding of a novel domain with familiar concepts from another, conceptual blending in metaphor and analogy is driven by the need to evoke an appropriately motivating construal of the target. Consequently, a given source domain might be chosen simply because it provides frames that make the desired aspect of the target salient. The nature of the correspondence between the source domain and the target is less relevant than in traditional cognitive approaches for two reasons: First, source domain frames are often altered in the blend; and, second, commonalities in the abstract or causal structure are neither necessary nor sufficient to evoke a motivating construal. Rhetorically motivated examples of metaphoric language examined in this chapter suggest a role for metaphoric blends that goes beyond the identification of cross-domain mappings.

Rather than emphasizing the extent to which metaphorical utterances instantiate entrenched mappings between source and target domains, conceptual integration networks represent only those cognitive models that are particularly relevant to the mapping supported by the utterance. What's projected in metaphor is not static information in long-term memory, but dynamically constructed entities in working memory.

Conceptual blending theory allows us to drop the old idea of concepts as static structures in long-term memory in favor of dynamically constructed models constrained both by information in long-term memory and by local, contextual cues. More interestingly, even though many of the models used in these processes are idealized, inaccurate, and even fantastic, they are invoked in such a way as to enable necessary inference, evoke appropriate affective responses, and motivate action.

CHAPTER SEVEN

Counterfactual Conditionals

If Xena Warrior Princess went head to head with Wonder Woman, who would win?

Each week the weblords at www.electricferret.com post in-depth discussions of issues just like this. They call their on-line magazine *Comic Book Universe Battles* and post a weekly play-by-play account of a battle between two fictional characters. For example, besides the wrestling match between Wonder Woman and Xena, other noteworthy battles have included Count Dracula versus Buffy the Vampire Slayer, and the aliens (from *Alien*, *Aliens*, and *Alien 3*) versus the Borg of *Star Trek* fame. While the topics inevitably involve imaginary characters in scenarios bizarre even from the standpoint of fiction, the debates can be quite spirited. Take, for example, the following excerpt from a discussion about a competition between cartoon sleuth Scooby Doo and the *X-Files'* FBI agents Mulder and Scully to uncover the cause of some mysterious deaths at an amusement park.

> Brian: [suppose] the mysterious occurrences at the amusement park are due to some paranormal, super-human, extra-terrestrial, and/or mystic force or entity. If that's the case, Scooby, et al., are completely out of their league. They can't handle real ghosts: Scoob and Shaggy would run, Daphne would scream and ask Fred to hold her close.... Mulder & Scully, however, are old hat at stuff like this....
>
> Steve: Your whole premise ... is totally ridiculous. It's not an issue of whether the monsters are paranormal or not; it's an issue of what the Scooby-Doo gang *thinks* they are.... In fact, Scooby and Shaggy always inadvertently capture the baddies *with the full belief that they are genuine monsters!*

While this discussion between Brian and Steve is whimsical and its outcome irrelevant, their reasoning proceeds much as if they were discussing the relative merits of real investigators. Further, while any

number of factors can be brought into the debate, it's not the case that just anything goes. No one could argue, for example, that one of the characters could use his or her superpowers to solve the case. Although Wonder Woman has superpowers, Scooby does not. Even fictional characters come with frames that can be used to predict their behavior. Brian and Steve's knowledge of the fictional characters can be seen to constrain the sorts of inferences they make about what Scooby, Shaggy, Scully, and Mulder might do in various hypothetical scenarios.

When reasoning about things that did not or could not happen, people appeal to the same sorts of operations as they do when reasoning about events that can and do occur. In all cases, people recruit partial structure from different domains, establish mappings between structure evoked in different input spaces, and integrate information via the creative mechanisms of conceptual blending. While the meaning construction in counterfactuals is often invoked for planning and evaluation of future events, the same acrobatic processes of meaning construction are also used to promote particular construals, to highlight shared experiences, and, as in Brian and Steve's case, for the sheer fun of it.

7.1 TRUTH, ACCEPTABILITY, AND CONDITIONALS

Curiously, logical treatment of conditional constructions like (1) often seems even less constrained than discussions of comic book battles. Because such constructions resemble the if-then of first order logic, many philosophers have considered conditionals to be instances of material implication.

(1) If the mysterious occurrence is due to paranormal activity, Mulder solves the case.

In the case of (1), both speakers and logicians agree that if (1a) and (1b) are both true, then (1) is a true statement.

(1a) The mysterious occurrence is due to paranormal activity.
(1b) Mulder solves the case.

However, if the "if" in counterfactual conditionals is understood as material implication, there are whole classes of conditionals that are true from a logical standpoint but unacceptable to speakers. For example, in material implication, conditionals are true whenever the consequent is true. This suggests that the truth of (1) depends on the truth of (1b) – and the existence of paranormal activity is irrelevant. Moreover, conditionals

are true whenever the antecedent is false. Consequently, (1) is true whenever (1a) is false, regardless of whether Mulder solves the case. In fact, as far as first order logic goes, (2) is perfectly acceptable.

(2) If coffee is red, Mulder solves the case.

This discrepancy between the truth and acceptability is particularly pronounced for *counterfactuals*, statements about the consequences of things that happen to be false. In a case where proposition P is contrary to fact, counterfactual reasoning involves drawing inferences about what P *would* imply if it *were* true. Because a counterfactual is by definition a conditional in which the antecedent is false, logically all counterfactuals are true statements. Apparently, the conditions that license the utterance of a counterfactual (that is, the case where the antecedent is false) render it trivially true. Such concerns have led some philosophers to argue that counterfactuals are *not* in fact cases of material implication, and prompted the construction of formalisms with which to represent differences in the acceptability of counterfactual utterances (e.g., Lewis, 1973).

However, it is precisely because the demands of meaning construction are somewhat orthogonal to the need to give truth conditions that counterfactual reasoning can be a productive activity. Because most counterfactuals stand in contrast to some actual world, understanding a counterfactual involves projecting partial structure pertaining to the actual state of affairs into a blended space where it can be integrated with more general information represented in other inputs to the integration network. For example, imagine I'm rummaging through my refrigerator for something to eat for lunch. Upon finding turkey, cheese, and mustard, I might utter (3).

(3) If I had bread, I could make a sandwich.

In this case, the inputs are the structure in the Actual space (in which I have turkey, cheese, and mustard) and a Sandwich frame with slots for Bread, Condiments, Meat, and Cheese. By composing the Bread from the Sandwich frame with the structure in the Actual space, I can construct my counterfactual sandwich-making scenario in the Blended space (see table 7.1). Though simplified, the integration network depicted in the table 7.1 captures the systematic correspondence between the current context and the counterfactual scenario in the blended space. The projection of structure from the inputs and the emergent structure that

Table 7.1: Elements in the Counterfactual Sandwich Blend

Input1	Blend	Input2
Actual	**Counterfactual**	**Sandwich**
Seana	Seana′	Agent″
Fridge	Fridge′	Fridge″
Turkey	Turkey′	Turkey″
Cheese	Cheese′	Cheese″
Mustard	Mustard′	Mustard″
	Bread′	Bread″

arises in the blend are both constrained by structural relations between the elements in each of the spaces.

The same sorts of blending operations can be seen to underlie meaning construction in more exotic counterfactuals. For example, meaning construction for (4) involves blending structure from the Actual space with information about Scooby and his sandwich-making abilities.

(4) If Scooby were here, he could make me a sandwich.

The Scooby space includes a sandwich-making scenario, abstracted from years of watching Scooby Doo, in which Scooby makes a sandwich. Both the Actual space and the Scooby space are Input spaces in the network, and the counterfactual scenario is represented in the Blended space in which Scooby makes a sandwich in my kitchen. The conceptual integration is fundamentally similar in (3) and (4), although Input2 in (3) draws on generic knowledge of sandwich making, while Input2 in (4) recruits memories of sandwich making in Scooby Doo cartoons. In both cases, the counterfactual scenario is conceptualized by integrating knowledge of the actual situation with more general knowledge by establishing mappings and completing the blend via imaginative processes of blending.

However, there are also some interesting contrasts between (3) and (4). For example, while both (3) and (4) make claims that can be evaluated for their acceptability, most people would agree that, of the two, the claim in (3) is a more straightforward, objective one. The blend in (4) differs from the blend in (3), not only in the incorporation of a cartoon character, but in its motivation. While (3) probably alludes to the absence of bread, (4) highlights Scooby's sandwich-making abilities. Differences

such as this led Turner (1996) to distinguish between *lab rat* counterfactuals and more exotic *spotlight* counterfactuals. The sort of counterfactual that has most concerned philosophers of science (e.g., Goodman, 1973) are lab rats such as (5) in which all the properties of the actual world are assumed to obtain in the counterfactual world, save the one being connected with a particular outcome.

(5) If this match had been struck, it would have lit.

Turner calls counterfactuals such as (5) lab rats because of their importance for scientific reasoning. Because lab rats posit minimal alteration in the blend, scientists can use them to run thought experiments much like biologists use lab rats.

However, Turner also points to the value of another sort of counterfactual that he calls the *spotlight*. Unlike the carefully controlled thought experiment that occurs when we entertain a lab rat counterfactual, the antecedent in a spotlight can involve substantially different assumptions than those that obtain in the actual world. However, the point of a spotlight is not its predictive power, but the way in which these constructions "spotlight" some aspect of the actual world that supports the line of reasoning expressed in the counterfactual. Spotlights are not deep metaphysical statements, but rather serve rhetorical goals by underscoring particular aspects of experience or perception. The next section considers meaning construction in a number of spotlight counterfactuals.

7.2 ANALOGY AND IDENTITY

The key to the rhetorical efficacy of the spotlight counterfactual can be found in its relationship to analogical reasoning. Suggesting commonalities between the way we understand analogy and the way we understand counterfactuals, Fauconnier (1996) points to (6) as an example of a counterfactual that clearly employs analogical mapping.

(6) In France, Watergate wouldn't have harmed Nixon.

The scenario is counterfactual because the Watergate scandal didn't happen in France, it happened in the United States. However, moving Watergate to France is a lot more complicated than, for example, imagining bread in my kitchen. Moreover, the import of (6) is as much the contrast between French and American politics as it is a prediction about Richard Nixon's would-be fate in France.

From a truth-conditional perspective, (6) is a nightmare. It would be quite difficult to treat this example as a lab rat, copying over all true propositions, and positing that Watergate happened in France. The problem is (as philosophers such as Goodman, 1973 and 1978, are well aware) that facts are related to one another, so that changing one thing can produce changes in the truth values of other propositions. The reason we can understand (6) is that we *don't* copy over all the true propositions. We construct a space with minimal representational structure to yield the desired inference.

Because understanding (6) involves imagining what would happen if the Watergate scandal occurred in France, it also involves analogical mapping. Meaning construction for (6) thus involves the establishment of mappings between elements in French and American politics and the construction of parallel structure. In particular, (6) requires the construction of a blended space in which an event from American history is blended with knowledge of the French political system. Mappings in one such blend are listed in table 7.2.

In the American politics space, "Watergate" serves to evoke an event scenario from American history in which then-President Nixon ordered operatives to break into Democratic headquarters at the Watergate hotel in Washington, D.C. When the president's criminal activities were uncovered by the Justice Department, he was forced to resign from office.

Table 7.2: In France, Watergate Wouldn't Have Harmed Nixon

Input1	**Blend**	**Input2**
French Politics	**Counterfactual**	**American Politics**
Elements	***Elements***	***Elements***
Jospin	Nixon'	Nixon
France	France'	US
	WG'	WG
Relations	***Relations***	***Relations***
Prime Minister	Prime Minister	President
(Jospin, France)	(Nixon', France)	(Nixon, US)
	Scandal (WG')	Scandal (WG)
	Implicated (Nixon', WG')	Implicated (Nixon, WG)
	Not-Harmed (Nixon', WG')	Resigned-Because (Nixon, WG)

While a speaker could not understand (6) without *some* knowledge of the Watergate incident, she need not retrieve all the details. As in the analogical mappings discussed in the last chapter, only partial structure is projected into the blend. Elements in different spaces are put into correspondence with one another, and the Access and Projection Principles apply.

One factor that contributes to (6)'s spotlight status is the way in which the Watergate frame requires accommodation to be mapped onto French politics. The French counterpart for the Watergate scandal might involve breaking into a French hotel in Paris, or it might involve a more schematic criminal activity frame. Although the blend is accessed from the American Politics space (i.e., we use terms such as "Nixon" and "Watergate," which come from American politics), its elaboration is constrained by knowledge of French politics. The contrast between Nixon's plight in the American Politics space and that of his blended counterpart triggers construction of a French Politics frame that can relate the scandal to the lack of ramifications.

Consequently, (6) suggests inferences about the French political system and the temperament of the French populace. This example suggests that analogy and counterfactuals both involve analogical connectors and the projection of partial structure into a blended space. Further, this type of blending is often rhetorically motivated because of the structure that it suggests in the input spaces. In (6), for example, elaborating the blend suggests that the French are more cynical than Americans about the character of their political leaders. Although the structure in the blended space itself is implausible, the structure that emerges from its elaboration might prove to be useful for reasoning about French politics.

Fauconnier's (1996) observation that counterfactuals involve analogical mapping is quite insightful, and follows naturally from certain aspects of mental space theory (Fauconnier, 1994). The connection between analogy and identity lies in the fact that, just as different individuals have different properties, a single individual's properties can change over time. Consequently, the formal requirements of structuring spaces and linking counterparts can be very similar in cases that involve analogical mappings between different individuals in similar scenarios, and cases that involve mappings between the same individuals in actual and counterfactual scenarios.

Contrast (7) and (8):

(7) A better quarterback would have gotten that pass off.
(8) Two years ago he would have gotten that pass off.

The blending that goes on in (7) and (8) is quite similar except that in (7) the ability of *another* quarterback is blended with the current scenario, while in (8) the current scenario is blended with the ability the *same* quarterback had two years ago. Of course, (7) and (8) are not identical: While (7) profiles a stative property of the quarterback, (8) profiles the result of a change in state. However, in both cases, the current scenario is compared to a model evoked in a blended space that takes the current scenario as one of its inputs.

These ability-plus-situation counterfactuals are set up because of the framing consequences they have for the input domains. The construal prompted by meaning construction in these examples can accentuate the contrast between actual and counterfactual, or make the linking models between actual and counterfactual more prominent. In (7), for example, explicitly considering how a better quarterback would perform emphasizes the actual quarterback's suboptimal status. Moreover, because ability has been construed as an inherent facet of the input quarterback, it doesn't make sense to compare the quarterback to himself, as in (9).

(9)# A better Steve Young would have gotten that pass off.

In contrast, if the lack of ability is considered a contingent facet of the quarterback, then comparing a quarterback to himself does make sense, as in (10).

(10) A healthier Steve Young would have gotten that pass off.

If a model exists that can transform the actual to the counterfactual scenario, then raising the counterfactual will serve to highlight the linking model. Thus, in (8), the contrast between the quarterback's performance in the Current space and the Two-Years-Ago space prompts the construction of a frame that can transform one into the other. For example, in (11), the player's presumed future ability is blended with the current situation.

(11) Two weeks from now he would have caught that pass.

In this blend, one input is a space structured with a schematic representation of the play that just occurred, while the other is a hypothetical future space in which the player's skills have improved. The structure in the blended space is identical to that in the first input, except that the player catches the pass. Interestingly, the prediction about the player's future ability has been assumed as fact.

Although the rhetorical goal of the utterance may involve construal of the player's future ability, the remark also serves to frame his current performance. By highlighting the linking Practice Makes Perfect model, the remark attributes his failure to a factor that is presumed to be reversible. Further, the statement frames his current failure in a fairly positive light. Although the unorthodox access properties of this blend are quite striking, formally it differs little from (12).

(12) Two years ago Favre would have been dropped for a loss on that play.

In (12) we have the less spectacular case where the current scenario is blended with past ability. A statement such as this one about Favre simultaneously construes his past ability negatively, construes his current ability positively, and gives him credit for improvement. Comparison of past to present is easier if the two scenarios are identical on all dimensions but the one that is currently relevant (much like the independent variable in an experiment). Although the construction of the counterfactual space relies extensively on imaginative ability, the rhetorical goal is to present the frame in a scientifically persuasive manner.

Fauconnier (1996) describes another example of a blend in which an agent's ability is changed and blended with aspects of current reality.

(13) When I was in my prime, I could have easily pinned Hulk Hogan.

The speaker here posits a wrestling match between a younger version of himself and Hulk Hogan. To do so, he constructs a counterfactual scenario in which information about the speaker's past is blended with information about Hulk Hogan's role in the modern wrestling world. The counterfactual space is used here to create an arena in which the speaker can compete with Hulk Hogan. Further, because the speaker easily wins the match with Hogan in the counterfactual scenario, the blend provides a context that evokes the Champion frame. Like many of the analogical blends discussed in Chapter 6, this blend results in retrospective projections to one of the input frames, suggesting that, in his day, the speaker was a champion wrestler.

Presumably one motivation for using a counterfactual statement is to evoke a frame that is unavailable in the current context. A good way to convey ability in a sport like wrestling is the caliber of the opponents you can beat. However, if the speaker is much older than his listeners, naming former opponents might not have the desired effect. Consequently, construction of the counterfactual can be motivated by facts about his

audience's knowledge base as well as by facts about Hulk Hogan. In positing the counterfactual, then, the speaker doesn't seek to transform the actual into the counterfactual, but to present a context in which his point will be readily evident.

Understanding the relationship between the input and the counterfactual involves appreciation of the mappings between the elements and relations in the two spaces. In some cases (e.g., "If this match had been lit . . ."), appreciating the mappings involves constructing a model of causal relationships that can simulate the transformation from the actual into the counterfactual. However, in other cases, appreciating the mappings involves accommodating the structure in the network to maximize the number of cross-space mappings. Because causal information is an important sort of higher-order information represented in frames, knowledge of causal relationships often structures counterfactual blends. However, other sorts of higher-order relations can also sustain counterfactual reasoning. The scientist's reasoning about matches being struck is, then, a special case of a more general blending process.

7.3 FRAME-SHIFTING AND SCALAR REASONING

7.3.1 Frame-Shifting

Examples of analogy and counterfactuals can both result in the abstraction of a schema that can be used to reason about one or both of the input domains. For example, in (13), the blend was used to create a scenario in which Hulk Hogan and the speaker could compete (albeit counterfactually) on equal terms. The mappings in this blend are listed in table 7.3. However, this example was taken from a longer exchange* that occurs as part of the comic strip *Drabble*, in which middle-aged couch potato Drabble brags to his son about his wrestling ability, saying, "In my prime, I could have pinned Hulk Hogan in seconds." When Drabble's son questions his claim, the boy asks his mother for confirmation. She responds as follows:

> Mom: Of course, dear. But when your father was in his prime, Hulk Hogan was in kindergarten.

So, while Drabble constructs a blend in order to evoke a Champion frame, Mom explores an alternative blend with a different rhetorical

* Discussed in Fauconnier (1996).

Table 7.3: Drabble's Blend

Input1	**Blend**	**Input2**
Drabble's Prime	Counterfactual	Hogan's Present
Elements	**Elements**	**Elements**
Drabble	Drabble′	Drabble″
Hogan	Hogan′	Hogan″
Relations	**Relations**	**Relations**
Wrestler (Drabble)	Wrestler (Drabble′)	Former-Wrestler (Drabble″)
	Champion-Wrestler (Hogan′)	Champion-Wrestler (Hogan″)
	Pins (Drabble′, Hogan)	
Champion (Drabble)	Champion (Drabble′)	Former-Champ (Drabble″)

goal. Mom's joke capitalizes on a metonymic reading of "In my prime," in order to trigger frame-shifting. Although rhetorically motivated, the blend she constructs is actually more like a lab rat than a spotlight. Having interpreted "in my prime" as designating the point in history when Drabble was fit, Mom assumes facts that obtained at that time also obtain in the counterfactual scenario. The exception, of course, is that a wrestling match occurs between Drabble and Hogan. The mappings in Mom's blend are listed in table 7.4.

Table 7.4: Mom's Blend

Input1	**Blend**	**Input2**
Drabble's Prime	Counterfactual	Hogan's Present
Elements	**Elements**	**Elements**
Drabble	Drabble′	Drabble″
Hogan	Hogan′	Hogan″
Relations	**Relations**	**Relations**
Wrestler (Drabble)	Wrestler (Drabble′)	Former-Wrestler (Drabble″)
Child (Hogan)	Child (Hogan′)	Champion-Wrestler (Hogan″)
	Pins (Drabble′, Hogan)	
Older (Drabble, Hogan)	Older (Drabble′, Hogan′)	Older (Drabble″, Hogan″)

The contrast in the implications of the counterfactual scenarios constructed by Mom and Drabble shows that Drabble's counterfactual is not constructed to explore alternative possibilities, but rather to convey his former wrestling prowess. To do so, he constructs a blended scenario that couldn't possibly have happened. Speakers don't usually motivate *why* particular frames are projected into the blend, and listeners are usually quite willing to suspend belief about certain things in order to generate emergent structure that gets projected back to the inputs. In unpacking the Hulk Hogan blend, the listener is not meant to consider how Drabble and Hogan might come to fight, but only who is the better wrestler.

Jokes often capitalize on the nondeterministic nature of structuring a counterfactual space. In this example, Astor blends her current self with Churchill's wife to yield emergent properties of motive and opportunity to murder Churchill.

Lady Nancy Astor, who entered the House of Commons in 1919 as its first woman member, had few betters in the art of vicious repartee. But she did have some, and Winston Churchill was one of them. After one hot and heavy Parliamentary set-to, Lady Astor is reported to have said to Churchill, "Winston, if you were my husband I would poison your coffee." "If you were my wife, Nancy," replied Churchill suavely, "I would drink it."*

Churchill elaborates on Astor's blend from his own perspective by imagining his reaction to the scenario in which his wife has been replaced by Astor. However, in shifting the blend to his perspective, Churchill attributes his counterpart in the counterfactual space with the knowledge that Lady Astor has poisoned his coffee. Although in doing so Churchill has only used the information at his disposal to construct a frame, one gets the feeling that he has somehow cheated. The joke violates the expectation that his counterpart in the blended space wouldn't know about the poison in the coffee. Moreover, the problem (and the humor) arises because the frame that Churchill constructs conflicts with the presumed rhetorical goal of Astor's utterance.

The nature of Churchill's alternative structuring procedures suggests that the use of goal hierarchies can help constrain interpretation. If using alternative goals can yield different interpretations, attempting to infer speaker's goals might help rule out spurious interpretations. Although speakers tend to use these goals, they're not required to do so. Moreover,

* Taken from Isaac Asimov's *Treasury of Humor*.

alternative interpretations are not completely random, but involve frequently encountered alternatives such as slot-filler ambiguities.

Once again the counterfactual blend has been used to set up frames so as to promote the speaker's desired construal of the events. In this example, the blend is invoked to capitalize on inferences supported by the marriage frame. Besides providing a plausible scenario for the murder, it provides a normative backdrop for evaluating the feelings that Nancy and Winston have for each other. Astor's desire to murder Churchill is worse coming from his (counterfactual) wife than from a stranger. Moreover, Churchill's desire to die rather than be married to Astor inverts the default preference hierarchy for death and marriage. Like analogy, counterfactuals are often rhetorically motivated.

7.3.2 Scalar Reasoning

Another reason to set up a counterfactual is to promote scalar reasoning, discussed more extensively in Chapter 9. Briefly, scalar reasoning involves inferences that arise from a *pragmatic scale*, a set of propositions ordered along some dimension such that elements are related by material implication. For example, in (15), the speaker sets up a pragmatic scale of distances ordered to produce inferences about Tony's ability to run them.

(15) Tony can't run five kilometers, much less ten.

Once we know that Tony can't run five kilometers, then a fortiori we can infer that he can't run distances greater than five kilometers. Often, speakers use counterfactual constructions that describe what might transpire in a situation that invokes an end point of a pragmatic scale.

For example, in (16), the speaker posits a scenario in which it would be highly likely that she become romantically involved with the addressee.

(16) I wouldn't go out with you if you were the last man on earth.

However, she does so by constructing a counterfactual blend between an end-of-the-world scenario and her current feelings for the addressee. Because the end-of-the-world scenario is an end point on a scale of favorable dating scenarios, it entails she won't date him in the less favorable real-world dating scenario either. The fantastic setting of this blend ultimately serves to emphasize the feelings that license the scalar implicature.

The extent to which counterfactual blends afford hyperbole makes them ideal for scalar implicature. Turner & Fauconnier (1995) discuss blending in (17), which recruits frames for Environmental Responsibility and blends them with Foreign Policy frames associated with the Yugoslavian civil war of the 1990s.

(17) If Bosnian Muslims were bottle-nosed dolphins, the West Europeans and Americans would never have allowed the slaughter of innocents to go on as long as they did.

Because the Great Chain of Being model (see Lakoff & Turner, 1990) places dolphins below humans on the scale of beings, the effect of the Environmental Responsibility frame in the blend suggests that an analogous Bosnian Responsibility frame ought to apply in the rhetorical target.

7.4 COUNTERFACTUAL SELVES

Calvin: I got another letter from my past self.
Hobbes: What's it say?
Calvin: "Dear Future Calvin, I wrote this several days before you will receive it. You've done things I haven't done. You've seen things I haven't seen. You know things I don't know. You Lucky Dog! Your pal, Calvin."
Calvin: I feel so sorry for myself two days ago.
Hobbes: Poor him. He wasn't you.

As Calvin's humorous letter points out, people change over the course of time, even to the point where they become alienated from their former selves. This facet of our concept of personal identity (or, in this case, disidentity) is exaggerated by Calvin writing to himself two days hence. Calvin's letter is funny precisely because we do tend to think of Calvin on Monday and Calvin on Wednesday as the very same Calvin. As discussed in Chapter 1, mental space theory is motivated by the need to represent different attributes and relations that can obtain for the "same" element. When Calvin's properties on Monday and Calvin's properties on Wednesday differ significantly from one another, the mental space theorist uses spaces to partition information about Calvin, and identity links to represent the connection between Calvin counterparts.

In spite of many changes over the course of a lifetime, our concept of personal identity affords links between identity counterparts that participate in radically different frames. Elaboration of various spaces that

contain a linked element can thus be wildly different. Because important concepts like personal responsibility rely on making (or breaking) identity connections, discussion of these issues produces a tension that leads to some interesting counterfactuals.

7.4.1 Knowledge Blends

In (18), the speaker begins by opening a space to represent the events of his past, accessed from his adult perspective.

(18) ... there was no freedom of choice for me as to whether I should be aborted or allowed to live. Had there been, I would not have chosen to be aborted.*

By noting that, as a fetus, he had no freedom of choice, he sets the scenario up as a relevant possibility. Further, he alludes to Pro-Choice rhetoric of his opponents and invites the listener to construct a frame in which the pregnant woman chooses whether or not to have an abortion. Having set up a disanalogy between the freedom enjoyed by the fetus and that enjoyed by the mother, the speaker moves to consider the counterfactual scenario in which the mother's choice is blended with the fetus' situation (that is, the fetus is asked to make the mother's choice).

The mother's choice presumably involves the impact of pregnancy and child rearing on her own and the life of a future child. However, in the blend, the fetus' choice is between being aborted or allowed to live. The Choice frame in the blend differs from the mother's choice, reflecting accommodation toward the perspective of the fetus. Further, the speaker employs structure imported from the adult space imbuing the fetus with agent status necessary for freedom of choice, as well as conferring the adult cognitive ability the fetus would need to make such a choice. Admittedly, though, the choice between being aborted and being allowed to live should not require much deliberation.

As noted in previous examples, the blend provides the speaker with an opportunity to entertain scenarios that would be ruled out by plausibility constraints operating in the input domains. By attributing the fetus with rational abilities of his adult state, and positing a scenario in

* Letter writer to the *Wall Street Journal*. Cited in Lemmons (1994). The transitive-ergative interplay in a polysemic complex: The issue of abortion. In Carlon, Davidse, & Rudzka-Ostyn (eds.) (1994).

which the fetus is treated like a rational agent and given a choice, the speaker creates a context that gives voice to the concerns of the anthropomorphized fetus while it silences the concerns of the mother.

Although (18) is an exotic case in which the blend combines the features of two disparate entities, counterfactual blending is employed quite routinely to combine the intelligence, experience, or disposition of one person with the situational needs of another. The construction in (19), for example, is an everyday variant of (18), in which the speaker's knowledge and the addressee's situation are combined.

(19) If I were you . . .

The blended scenario can then be elaborated so that the addressee can judge the speaker's advice.

Turner & Fauconnier (1998) note a similar blend that came up in discussion of a bizarre moral dilemma. In a case that provides strong support for the adage that truth is stranger than fiction, a comatose woman was raped by a hospital worker and subsequently became pregnant. The controversy concerned whether the pregnancy should be terminated on the grounds that the woman was in a persistent vegetative state (PVS).

Goldberg (a law professor): Even if everyone agrees she (the comatose woman) was pro-life at 19, she is now 29 and has lived in PVS for 10 years. Do we ask: "Was she pro-life?" Or do we ask more appropriately: "Would she be pro-life as a rape victim in a persistent vegetative state at 29 years of age?"

To determine the normative course of action, Goldberg recommends blending the rational abilities of a healthy individual with the experience of an individual who is brain dead. There is a folk model that experiences, especially traumatic experiences, change people's cognitive abilities in a way that increases their understanding of similar experiences, improves their decision-making competence, and confers authority on their beliefs. In this case, the law professor appears to have applied this model to a case in which the relevant experience precludes the possibility of benefiting from it. Goldberg's argument is that her perspective at nineteen is not relevant because she could not (or would not) have envisioned the current scenario. Rather than merely marking the continuity of reference, identity links are exploited to evoke proper construals of agency, innocence, and moral authority as necessitated by the speaker's rhetorical goals.

7.5 CONCLUSIONS

In all of these examples, partial structure from one situation is combined with partial structure from another, often in ways that are quite implausible. For example, Hulk Hogan is sent back in time to wrestle a young Drabble, two bitter enemies are joined in marriage, and the rational abilities of an adult are combined with the innocence of a fetus. Remarkably, these implausible blends do seem to promote the speakers' communicative ends. Drabble suggests he was once a wrestling champion, Churchill suggests his dislike for his distinguished colleague, and the speaker in (18) conveys his opinions about the validity of abortion rights. These implausible blends in counterfactual spaces are often done to change the prominence of some aspect of the scenario in order to make it consistent with another high-level frame.

As in analogy, structuring the blended space to represent counterfactual content involves the establishment of abstract relational correspondences that can result in productive inference. Like analogy, meaning construction for counterfactuals involves the projection of partial structure into a blended space and makes extensive use of cross-space mappings. As the nature of the slot-filler ambiguity in the Hulk Hogan example suggests, the metamorphoses that an individual undergoes over the course of a lifetime present the same projection problem for identity connections as analogical connections. Consequently, counterfactual constructions can also recruit analogical links, as in Fauconnier's (1996) Watergate-in-France example.

In fact, we might think of counterfactuals as *really* good analogies because the alternate scenario is identical to the target in all but one key respect. Like the target space in an analogy, the counterfactual space inherits most of its structure from the source input. But while the point of an analogy is usually to evoke structure in the target, the point of a counterfactual is usually to project structure back onto the source input.

The conceptual integration that occurs in counterfactual spaces is yet another instance in which projection mappings are used to evoke different patterns of inference, affective responses, and action frames. In everyday life, counterfactuals are almost always set up for rhetorical purposes in which the aim is to profile a particular construal. Since the rhetorical efficacy of a counterfactual does not depend on the plausibility of the counterfactual scenario arising in the first place, speakers employ counterfactual blending to promote scalar implicature and to trigger frame-shifting. In Turner's parlance, the counterfactuals of everyday

speech are spotlights, intended to underscore some aspect of experience or perception.

Contrasting traditional philosophical interests with the common usage of counterfactuals, Turner (1996) argues eloquently for the irrelevance of truth values in determining the meaning of a spotlight:

> Distinctions of actuality, probability, and possibility are not principally at issue. We see in such cases little anxiety about clearly specifying antecedents, consequents, and principles of causal connection; about demonstrating consistency with well-established historical facts, well-established theoretical laws, and well-established statistical generalizations that are projectable; or about analyzing degree of proximity to the actual world.

However, in some ways the spotlight and the lab rat are not as different as one might think. Recall that philosophers are mainly interested in counterfactuals because scientists employ them to draw causal inferences such as the relationship between the striking of a match and the subsequent flame in (5).

(5) If this match had been struck, it would have lit.

However, Goodman has noted that even this trusty lab rat requires nontrivial faith in background assumptions for its validity. Imagine, for example, that the match was wet, or struck in an airless environment (Goodman, 1973). In fact, Goodman's problems with matches that do and do not light bear more than a passing similarity to Searle's (1979) question about when a weightless cat can be said to be "on" a weightless mat in outer space. Both accentuate the bond between meaning and background that arises because words designate elements in partial frames whose elaboration is constrained by background knowledge.

In the absence of other information, people appeal to default assumptions such as the existence of oxygen and gravitational fields. Furthermore, default assumptions display a sort of context-sensitivity we might expect from a data structure whose slots are interdependent. In the case of physical events such as the implications of striking a match, the background assumptions themselves are not particularly revealing. Goodman's original argument was insightful precisely because imagining lighting a match does not involve imagining getting it wet. However, imagining that a red apple is a green apple changes its taste as well as its color. All of the same issues that arise in simple cases of conceptual combination in noun modification arise again for counterfactuals and other sorts of conditionals. This is because operations that appear to

involve a single predicate in fact involve an elaborate web of mappings between elements and relations in disparate domains.

In many ways, the puzzle of counterfactuals arises out of a discrepancy between our ideas about meaning construction and the way that these processes actually proceed. We have seen that meaning always arises out of a limited set of background assumptions. Moreover, every time we set up a new space, a new set of background assumptions operate. In fact, perhaps the difference between the lab rat and the spotlight involves the extent to which the background assumptions deviate from the norm. Because invoking background assumptions is ultimately a matter of practice and convention, some frames are more entrenched than others. While lab rats appeal to these entrenched frames for interpreting relationships between physical elements, spotlights appeal to models that apply more controversially. If a model is an entrenched predictive model, the statement is a lab rat. However, in more complicated and frame-ambiguous cases, the statement will seem more like a spotlight directing attention to something that had gone unnoticed.

The curious fact about counterfactuals has been that while the computational requirements of explicitly representing all the consequences of change are quite formidable, people seem to do it remarkably well. However, perhaps the reason people have so much apparent success with this problem is that they do *not* represent all the consequences of change. Blending theory suggests people use simplified, partial models, and use rhetorical goals to restrict their inferences. The flexibility of meaning construction is such that the spaces we link can be as close as a match before and after being lit, or as distant as those inhabited by Agent Scully and Scooby Doo.

PART THREE

APPLICATIONS: BLENDING, FRAMING, AND BLAMING

Schank & Abelson's (1977) formulation of scripts represented a major advance in cognitive science in its acknowledgment of the constructive nature of comprehension. However, in some ways scripts and even the more recent, technologically sophisticated implementations do not go far enough in explaining the constitutive role of cognitive models in the structure and interpretation of experience. The latter has been almost exclusively the province of cognitive anthropologists who treat scripts as culturally shared representations that people use to give meaning to their actions in the world. Anthropologists Quinn & Holland (1987), for example, argue that the uniformity in culture members' description of particular scripts suggests that these data structures have cultural origins. Besides providing descriptions of action sequences in typical events, scripts can be used by people to structure plans, expectations, and actions. The organizing capacity of scripts extends well beyond the brain and into the set of practices that help constitute the social world. The restaurant script is a useful representational structure precisely because practices in real restaurants conform (more or less) to the script.

Another term for these culturally shared frames is *cultural models*, taken-for-granted models shared by members of a given social group, discussed briefly in Chapter 6. As data structures, cultural models consist of standardized sequences of events in a pared-down world. The definition of cultural models retains essential elements of scripts, frames, and schemas, while emphasizing their intersubjective nature and cultural origins. But although cultural models are widely shared among culture members, they need not correspond to anything in the external world. A good example is Sweetser's (1987) discussion of what she calls the simplified speech act world. In this world, people speak

to communicate helpful information; people's beliefs are adequately justified (and, as a result, are true); and, finally, people say what they believe.

Obviously, the simplified speech act world is an idealized construct. People do not always say what they believe, nor do they always speak with the intention of providing helpful information. Nonetheless, the simplified speech act world is an efficient representation for reasoners to use to define speech acts that *don't* fit the model. By showing how speakers might employ the speech act model to note deviations from its component parts, Sweetser (1987) provides a unified account of the semantics of the word "lies" and related terms "social lies," "white lies," and "mistakes." Moreover, one can see from its resemblance to Grice's (1975) maxims that the model is used implicitly by speakers to understand the pragmatic implications of each other's statements. Once again we see how the data structures that linguists invoke to understand lexical semantics can help explain performance on other cognitive tasks such as reasoning, problem solving, and judging the behavior of others.

In Part III, background knowledge moves into the foreground as we examine the ways in which speakers negotiate how culturally shared models should be used in the social world. In the past, cognitive scientists have studied cultural models via discourse analysis (e.g., Holland & Skinner, 1987; D'Andrade, 1987; Hutchins, 1980), as well as the systematic use of native intuitions (see Lakoff & Johnson, 1980; Lakoff & Kovecses, 1987). The analysis in the next two chapters results from a combination of these techniques to spoken and written discourse on the morality of abortion. The abortion topic was chosen chiefly because it is familiar to most Americans and prone to provoke animated discussion from informants. Despite interviewers' efforts to convince them otherwise, informants invariably feel that their responses to questions are being judged as correct or incorrect (Cicourel, 1988). It was hoped that the controversial nature of this topic would mitigate the problem.

In their examination of longer stretches of naturalistic discourse, the next two chapters illustrate real-world examples of many of the semantic leaps discussed in Parts I and II. Chapter 8 examines pragmatic features of informants' discourse to provide an account of the inferential machinery that people utilize when reasoning about the morality of abortion in the case of a pregnant rape victim. Chapter 9 demonstrates the use of frame-shifting in serious contexts, and considers how it interacts with

pragmatic scales. The creative character of the conceptual blending that supports the various framings, reframings, and frame shifts in these examples suggests that the way people use cultural models to construe the social world is far more complex than one might expect from the classic account of knowledge instantiation.

CHAPTER EIGHT

Framing in Moral Discourse

We laugh every time a person gives us the impression of being a thing.

– Henri Bergson, *Laughter**

In pointing out that individuals often laugh when confronted by a person who does not sustain in every way an image of human guidedness, Bergson only fails to go on and draw the implied conclusion, namely, that if individuals are ready to laugh during occurrences of ineffectively guided behavior, then all along they apparently must have been fully assessing the conformance of the normally behaved, finding it to be no laughing matter.

– Erving Goffman, *Frame Analysis*†

In the passage that follows this quotation, Goffman (1974) argues that the perception of everyday activity relies extensively on the projection of frames – so extensively, in fact, that it is often necessary to point to extreme cases of frame-shifting to reveal the constructive nature of social perception. Although their social function was central to Bateson's (1976; 1956) original concept of frames, this facet of background knowledge has been less prominent in contemporary cognitive science. One exception to this trend can be found in the work of cognitive anthropologists who study *cultural models*, or culturally shared frames. While previous chapters have focused on the role of frames in meaning construction prompted by language, many of those same frames are used to structure our actions and our expectations as we interact with each other in the world. Consequently, this chapter explores the social dimension

* Bergson, H. (1911). *Laughter*. Translated by Brereton, C., & Rothwell, F. London: Macmillan & Co., 58.

† Goffman (1974: 39).

of frames, examining real-life uses of cultural models in informants' discourse about the morality of abortion.

Section 1 lays out the cultural models of action, responsibility, and punishment relevant to the moral arguments analyzed; section 2 examines discourse data extracted from a series of interviews in which informants legitimated their views on the morality of abortion (see Coulson, 1992); and, section 3 discusses commonalities in the interpretation of language and activity. Analysis confirms the intuition that moral reasoning, at least about this issue, largely involves the presentation of alternative *framings* for various aspects of a morally ambiguous scenario. A central cognitive activity, framing *a* as *b* involves integrating portions of the frames or models in which *a* and *b* participate. We will see the nontrivial nature of the blendings that support framing activity, as well as the nontrivial nature of its consequences.

8.1 CULTURAL MODELS OF ACTION, RESPONSIBILITY, AND PUNISHMENT

D'Andrade (1987) provides an example of a cultural model in his exploration of the American folk model of mind. Describing the American commonsense view of the mind, D'Andrade details what sorts of states and processes occur in the mind (e.g., beliefs, desires, perceptions, intentions, resolutions, and emotions); the syntax of the relevant terms; the relationships that can exist between internal states and external actions; and also the relationships that can exist among the internal states themselves.

The idealized model of action, for example, concerns an agent with an intention, who performs an action that has consequences. On the model, the intentions cause actions, and actions cause their intended consequences. An agent's belief that his action A will produce consequence C and the desire to produce consequence C via action A are both important components of the intention to produce consequence C. The action model is what underlies Americans' (and D'Andrade suggests the model holds for many Europeans as well) interpretation of their own and each other's actions. Our default assumption is that all human acts are caused by intentions. If we observe a person performing an action (e.g., opening a door), we assume that the act is caused by the person's intention (the intention to open the door).

We don't mention intentions in our description of actions because the intention is implicit to the very notion of an action. Because all actions are assumed to stem from intentions, this information is seen as trivial. As

Bergson notes, unintentional action is the exception, provoking laughter and ridicule. While we don't usually mention that acts are caused by intentions, we *do* sometimes consider it relevant to specify the cause of the intentions. Because intentions can be caused by desires, wishes, needs, beliefs, and combinations of these, one might explain another person's actions by recourse to the desire that gave rise to the intention, which, in turn, gave rise to the action. Alternatively, one might simply make reference to the external cause that gave rise to the desire (wish, need, etc.) that in turn gave rise to the intention and, consequently, to the action.

The concept of responsibility is based on this idealized model of action. An agent is responsible for her actions because they are caused by her intentions. Similarly, an agent is responsible for the consequences of her actions, first, because of the causal link between the agent and the action on the idealized model of action, and, second, because of the causal link between actions and consequences. The attribution of praise, blame, and punishment is based on the concept of responsibility, and consequently involves a similar use of the idealized model of mind. One receives praise, blame, or punishment as a consequence of responsibility for an action. If the action is judged as having good consequences, the actor responsible receives praise; if the action is judged as having bad consequences, the actor responsible receives blame and possibly punishment.

The cultural model of just punishment is evoked when the action model is instantiated by an agent in such a way as to evoke the negative evaluation of others. In such a case, the model dictates that the agent responsible requires blame and punishment. The crucial aspect of this model is that the recipient of the punishment be coextensive with the agent responsible for the action and consequence. Otherwise, the cultural model of unjust punishment applies. Note that there is no a priori reason why an agent should be responsible for his actions. The attribution of responsibility to the agent results from the use of the models of mind and action. Agents are accountable for the consequences of their actions because, on the model, acts are causally linked to consequences by the agent's knowledge of the outcome and her desire for that outcome.

8.2 FRAMING AND THE MORALITY OF ABORTION

Of course, it is one thing to talk about cultural models in the abstract, and quite another to claim that people actually use them in interpreting

and evaluating events. It may seem unlikely that these simple models could capture the complexities of social life. However, this section argues that people do indeed use cultural models in their reasoning about issues such as the morality of abortion. We shall see that cultural models are effective only because of speakers' ability to adapt their models dynamically to meet contextual demands. Consequently, while the models themselves don't capture the intricacies of the issues, the people who use them can, and often do.

Because talk is central to social life, and because meaning construction works via the retrieval and elaboration of knowledge structures, we turn to discourse data for a glimpse of cultural models at work. The following excerpts are taken from a series of interviews in which informants were asked to discuss their views on the acceptability of abortion (Coulson, 1992). Though informants were asked about a number of hypothetical scenarios, the focus here is on the acceptability of abortion in the cases of rape and incest. In particular, these excerpts concern why the acceptability of abortion is somewhat less contentious for rape victims than for other women with unwanted pregnancies.

Asked to justify abortion in the case of rape, informants point to two sorts of things. The first is the undesirable effect that pregnancy and child rearing will have on the rape victim. This includes the idea that the mother will not want a child that is the result of a rape; that the child will remind the mother of the rape; and that bringing up the child will be painful for the mother. Moreover, the trauma caused by pregnancy and child rearing in such a case is linked to undesirable ramifications for the quality of the child's life – that is, the mother will not do a good job of bringing up such a child, and both the child and the mother will suffer from such an arrangement. As we will see in this section, informants also pointed to the unintended nature of the pregnancy caused by rape.

8.2.1 Choice, Trauma, and the Rape Exception

An example of the latter sort of response can be found in the following excerpt from an interview with a young woman in her early twenties. The interviewer begins by posing a hypothetical situation regarding a pregnant rape victim, and asks P whether it would be wrong for this girl to obtain an abortion. The question and P's response follow:

(1) I: Okay, how about, um, J is seventeen years old and is pregnant as a result of a rape.

(2) Would it be wrong for her to have an abortion?
(3) P: No. Because it wasn't her choice.
(4) It was something she couldn't prevent.
(5) And there's no use traumatizing her with pregnancy
(6) and, um, being young and being pregnant.
(7) Having this child, um,
(8) I mean I would have a very hard time.
(9) I would probably have an abortion if I were raped.
(10) No matter what the age I would support abortion.

In response to the interviewer's question, P denies that it would be wrong for a rape victim to obtain an abortion, and supports her claim with the statements in (3)–(9). However, for us to see the claims made by P as a justification for the rape exception, they must be understood as functioning against the background of a set of taken-for-granted beliefs (Hutchins, 1980). For example, P's denial in (3) that the rape victim chose to get pregnant relies on a cultural model of rape in which sex results from the man's intentions and contrary to those of the woman. But, given that the *choice* to get pregnant is neither necessary nor sufficient for pregnancy to ensue, what is it that makes P consider the woman's choice a morally relevant factor? Similarly, what prompts P in (5) to deny the *use* of *traumatizing* the rape victim with pregnancy?

The answers lie in the cultural models of action and responsibility that speakers use to make moral judgments. On the cultural model of action, actions are caused by the agent's intentions. By denying the rape victim's intention to get pregnant, P's move in line (3) is to deny that sex that occurs in rape constitutes an action on the part of the woman. Because the model for responsibility attribution is based on the action model, pointing to the unintended nature of the pregnancy can function as an argument that rape is not an action on the woman's part and that the woman is not responsible for the rape and its consequences. Similarly, P's statement in (4), denying that the rape victim could have prevented the pregnancy, also argues against her responsibility for it.

The way in which P cancels the Responsibility frame suggests that the application of cultural models has much in common with the more well-known process of setting up speech act frames (see Levinson, 1983, for review). Just as P cancels the Responsibility frame by denying the preconditions for its application, naming the preconditions for an event frame seems to be a good way to access the rest. For example, to explain how he got to the party, Bill might felicitously offer (a) to suggest that

he drove, or offer (b) to explain why he *didn't* drive.

(a) I have a car.
(b) I don't have a car.

Further, he might use (b) to deny Bob's request for a ride by denying the preconditions for fulfilling the request. Indeed, indirect speech acts are often made by inquiring about the preconditions for an action, such as the requestee's ability to comply with the request (as in (c)).

(c) Can you pass the salt?

Given that canceling the application of a cultural model involves denying one of its preconditions, it's interesting that P did not merely note that the rape victim might not want children. Rather she points to the trauma she imagines the rape victim might experience in carrying the pregnancy to term. One coherent interpretation of (5), then, is as an expansion of the claim that the attribution of responsibility is unfounded. Moreover, the presupposition that pregnancy will be a traumatic experience for the rape victim prompts a construal of pregnancy as punishment that results when the Responsibility frames that normally apply to pregnancy are blended with the model of pregnancy as a traumatic experience for the rape victim. The resultant blend is a transformation of the default reproduction scenario with significantly different moral implications.

In the reproduction scenario, pregnancy is the consequence of an action (sexual intercourse) for which the woman is responsible, and that engenders obligations to both the fetus and the eventual child. In the rape scenario, pregnancy is the consequence of an action (rape) for which the rapist is responsible, and that engenders trauma for the rape victim. In the blend, pregnancy is construed both as a consequence of an action that engenders obligations to the fetus and the eventual child, *and* as the cause of trauma. The composition of these two conceptualizations of pregnancy in the blended space is what enables the framing of pregnancy as a punishment of the rape victim. Moreover, because the agent of pregnancy in the blend (as well as in the rape space) is the rapist, and the recipient of the "punishment" is the rape victim, the scenario is completed with the unjust punishment model.

While P's remaining move, from the construal of pregnancy as unjust punishment to support for abortion, might seem obvious, it is far from trivial. For one thing, it recruits a framing of abortion that fits into the punishment scenario. Moreover, it relies in part on the elaboration of a

blend in which P projects herself into the pregnant rape victim's scenario. Citing the difficulty associated with such a pregnancy ((8), "a very hard time"), P asserts in (9) that *she* would have an abortion in this situation. The empathy P expresses toward the hypothetical pregnant rape victim is generalized in (10) as support for abortion. To understand this move as coherent, we must frame abortion as a means for preventing ineffective and unjust punishment. Construed in this way, abortion is the morally acceptable course of action.

8.2.2 Blaming the Victim

One might expect the success or failure of alternative framings to depend on some measure of similarity between the various pairings of representational structures – the number of shared features or shared relational structure, for example. However, on structural grounds, the blend that affords the framing of the rape victim's pregnancy as punishment is far from optimal. Punishment is not typically a causally related consequence of the crime, but rather something dictated by an external authority. Moreover, even in cases of *unjust* punishment, the recipient of the punishment is not typically the victim of the crime, nor is the criminal the agent of punishment. More likely, this blend arises because the cultural framing of pregnancy as entailing responsibility is so entrenched that it is readily available as a counterfactual input.

As in the previous example, the following excerpt is taken from a longer interview about the morality of abortion. The interviewer is the same, but the informant, S, is a woman in her late twenties. During the course of a discussion concerning a young rape victim who was denied the right to obtain an abortion, the interviewer asked S why rape is considered a compelling case for the right to obtain an abortion. S's argument, like P's, also involves an appeal to the prototypical operation of the responsibility model. But, while P gave reasons the responsibility model should not apply to the rape victim, S's answer suggests she thinks the responsibility model should not apply to any cases of unwanted pregnancy. Moreover, S cancels the responsibility model by appealing to an alternative model embedded in the rape scenario: the Victim frame.

(1) I: Do you think the fact that she was raped
(2) makes the case more compelling?
(3) S: Well, for some people. I don't think, well, I mean for me,

(4) I feel like, yes, that's very important.
(5) She should have the right to do that anyway.
(6) But it seems like, well, at least some people,
(7) that would convince them.
(8) I: Yeah. Why, why do you think that rape is such a compelling reason?
(9) S: Um, because I think people may tend to blame,
(10) they may tend to blame the woman for getting pregnant.
(11) But yet if she were raped, then they can start to slightly see her as a victim.
(12) I: Uh huh.
(13) S: And not taking the blame for her, um, pregnancy as much.
(14) I: Uh huh.
(15) S: I mean, I don't, yeah, I don't, um, ascribe to those feelings,
(16) but I can see where people do.
(17) That if the woman gets pregnant, it's completely her fault
(18) and she must bear the outcome of that.
(19) I: Uh huh.
(20) S: If she were raped, it's not as much her fault.

S goes to some effort in lines (3)–(5) to express the opinion that she sees nothing exceptional about the rape victim. S considers abortion morally unproblematic, making the rape victim's case no different from any other unwanted pregnancy. Nonetheless, she is quite able to articulate the reasoning behind the rape exception, identifying a contrast between the agentive role of the pregnant rape victim with that of a woman whose pregnancy results from a more conventional encounter. In the normal case, a woman is assigned responsibility for pregnancy (lines (9) and (10)), while the woman who becomes pregnant as the result of a rape is construed as a *victim* and, consequently, is not subject to the same sorts of attributions of responsibility (lines (11)–(13)).

Like P's argument, S's explanation is directed at how the generic unwanted pregnancy scenario is transformed via integration with the rape scenario. However, the particular articulation of the cultural models differs. For instance, in S's blend, both victimhood and blame (or at least their perception) admit of degrees. In the normal case of pregnancy, it's "completely her fault", while in the rape scenario, people can "slightly see her as a victim," in which case it's "not as much her fault." Moreover, the two are inversely proportional: As the perception of victimhood increases, blame goes down.

Although the connection between pregnancy and abortion is never made explicit, it emerges from S's statements in (17) and (18). In (17), S sets up a conditional space in which a woman gets pregnant. The expansion of the conditional involves the attribution of responsibility for the pregnancy. Because S is simulating the reasoning of others, the sequence from the antecedent to the consequent does not follow the temporal order of events in w's life, but rather the temporal ordering of the moral attribution process. Marked by *"and,"* (18) is connected to the scenario described in (17). Further, the use of the term "must" signals the use of a modal space that represents the way things should be. The normative course of events is thus to "bear the outcome" (in this case, quite literally) of pregnancy. Since abortion is not consistent with the events in the modal space, it is ruled out (see figure 8.1).

Despite the similarity in the two informants' explanations of the rape exception, S's choice of words differs markedly from P's in suggesting a negative construal of the pregnant woman and her situation. In (9) and (10), for example, she uses the term "blame," which involves a negative evaluation of both the agent and the outcome of the action. She uses the term "fault" in (17) and (20) in connection with the woman's pregnancy.

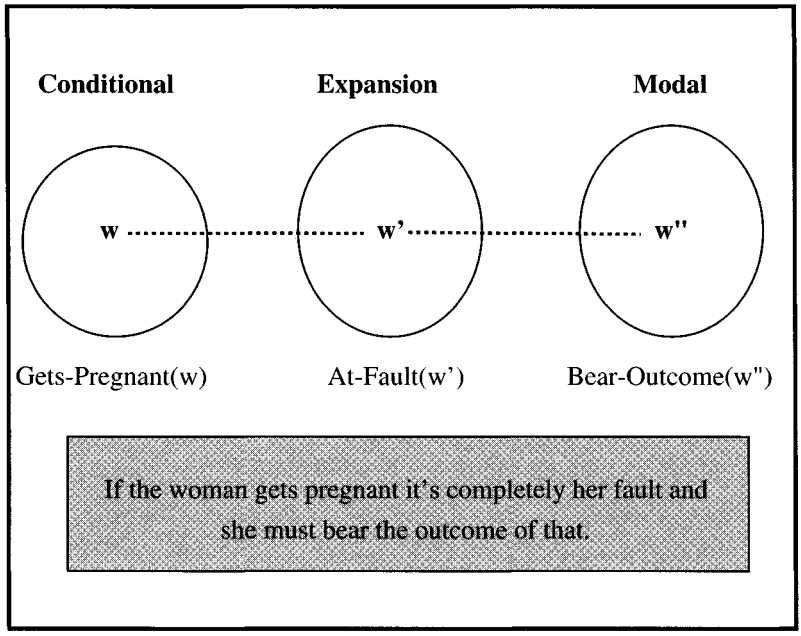

Figure 8.1 Conditional Reasoning in S's Argument.

Further, S uses the term "bear" in (18) to describe the consequences of getting pregnant.

One explanation for S's negative framing is that the scenario in question is not the prototypical (idealized) scenario in which pregnancy is an intermediate step in starting or growing a family. Instead, S is dealing with a particular notion of unwanted pregnancy in which the concept of blame has replaced that of the more neutral Responsibility frame. Her exaggerated negative framing may also be a way of cueing the discrepancy between her own views and the conservative view S has been asked to explicate. S herself does not "ascribe to" the construal she voices, a possibility that arises out of the oft-overlooked distinction between mapping and blending. The partial mappings between structure in cultural models and the social scenarios are necessary but not sufficient for integration to occur in the blend.

8.2.3 From Victims to Criminals and Back

Given that neither of the informants discussed in the previous sections espoused an anti-abortion viewpoint, one might question whether their explanation of the logic of the so-called rape exception was accurate. However, elements of our informants' reasoning can also be seen in the following statement by former Idaho Governor Cecil Andrus. Unlike either of the informants, the anti-abortion Andrus considers the cases of rape and incest as truly exceptional with respect to the morality of abortion. The following statement was issued to justify his veto of a bill that would have made abortion illegal in Idaho.

(1) This bill is drawn so narrowly that it would punitively
(2) and without compassion further harm an Idaho woman
(3) who may find herself in the horrible, unthinkable position
(4) of confronting a pregnancy
(5) that resulted from rape or incest.
(6) This law would force the woman to compound the tragedy of rape.
(7) On the eighth day she ceases to be the victim and becomes a criminal.

Andrus' argument here differs somewhat from that offered by our informants in that he makes no mention of the responsibility model and its possible application to cases of unwanted pregnancy. On the other hand, there is no implication that the Idaho woman he refers to *is* responsible for the pregnancy. In fact, Andrus' use of language seems

designed to highlight her nonagentive status, and to accentuate the extent to which she deserves our sympathy. To this end, he recruits an entrenched metaphorical mapping between situations and locations (see Lakoff, 1993) to frame the pregnant rape victim's situation negatively in (3) as a "horrible, unthinkable position" in which the woman "finds herself." The pregnancy itself is metaphorically framed as an adversary that the woman must "confront" (in (4)). Moreover, she is framed explicitly in (7) as a "victim."

In comparison to our informants' on-line responses to the interviewer's questions, the off-line character of Andrus' prepared statement allows more room for rhetorical flourish. For instance, Andrus begins by recruiting a metonymy between the bill and the (counterfactual) effects of the bill, and blending the fused result with an agent model normally reserved for people. Personification of the bill allows Andrus to frame it in lines (1) and (2) as acting "punitively" and "without compassion" to "harm an Idaho woman" – quite remarkable feats for a document.

For all his rhetorical flair, though, Andrus appeals to the same cultural models as our informants. He invokes the unjust punishment model in (1) with the use of "punitively," noting that pregnancy causes an increase in the trauma caused by rape (lines (2) and (6)). The unwanted character of the pregnancy is emphasized by the blended construal of the whole situation as a horrible place where no woman would even *think* of going. Moreover, in line (7), Andrus contrasts what he considers the woman's proper role (the victim) with the improper role (the criminal) she assumes in the unjust punishment model. Finally, although the ostensive topic of abortion is never mentioned, Andrus implicitly appeals to the same model as did our informants: abortion as a means to prevent unjust punishment.

Overall, the blends in Andrus' argument seem motivated by his rhetorical goal: to justify his veto without supporting the morality of abortion. By personifying the bill, we can evaluate it with cultural models that are usually used to judge and evaluate other people, and more easily judge the bill as immoral. Moreover, he uses the term "law" in (6) to access the counterfactual space in which he signs the bill. In this space, the same cause/effect metonymy obtains as before, and the personified law "forces the woman to compound the tragedy of rape." It is only in this counterfactual space where the victim has the role of criminal, and it is the law in this space that metes out the unjust punishment. The point of Andrus' argument is to prompt inferences in the reality space in which he vetoes the bill to allow the victim to remain a victim.

8.2.4 Testimony Against the Rape Exception

In theory, the abortion issue is wholly distinct from questions of responsibility for pregnancy. However, in practice, responsibility for pregnancy *does* affect attitudes toward abortion and the women who undergo the procedure. The default assumption seems to be that pregnancy signals a moral agency that entails both responsibility and obligation. Factors such as youth, innocence, and victimhood can undermine agency, and serve to mitigate responsibility and obligation engendered by pregnancy, thus making abortion a moral option. In the following example, we hear from a woman whose personal experience makes her an expert on this issue. We might expect this woman, as a victim of rape, to endorse the rape exception. However, the following excerpt from a letter, written by Lee Ezell to the editor of the *Los Angeles Times*, argues otherwise.

(1) I say thanks that no Planned Parenthood Clinic was available to me
(2) in 1963, when, as a virgin teenager, I was raped and became pregnant.
(3) The state of California would have been taking advantage of me
(4) in my crisis state by offering me this seemingly easy out.
(5) As an unwanted child myself,
(6) I decided abortion was too permanent a solution to my temporary problem.

Lee begins her letter with the statement, "My mother cried the day she found out she was pregnant with me," and continues to describe the trials and tribulations of her upbringing, including her experience as a pregnant rape victim. Lee's argument here is directed in part at promoting what I call *epistemic authority*, in which a speaker conveys information that suggests her authority to speak to the issue. Contextually relevant experiences are thought to convey knowledge and license epistemic authority. Since arguments for the rape exception rely on mental emulation of the pregnant rape victim's feelings, personal testimony about the actual experience can serve either to reinforce cultural framing of the scenario or, in this case, to undermine it. Moreover, because Lee's experience includes being a pregnant rape victim *and* an unwanted child, she is particularly well suited to make claims about this issue, typically posed as concerning the conflicting interests of mother and child.

In noting her mother's somewhat analogous experience of unwanted pregnancy, Lee has set up two spaces with parallel structure: a pre-1963

Table 8.1: Parallel Structures

Pre-1963	**1963**
Elements	**Elements**
Mrs-Ezell	Lee′
Lee-F	Julie-F′
Pregnancy	Pregnancy′
Relations	**Relations**
Pregnant (Mrs-Ezell, Lee-F)	Pregnant (Lee′, Julie-F′)

space in which her mother is pregnant, and a 1963 space in which she herself is pregnant with her own daughter, Julie (see table 8.1). The parallel scenario between mother and daughter makes for particularly complicated access properties because (the 1963) Lee is linked to her mother (Mrs-Ezell) analogically, and to her own fetus by identity. Moreover, because the topic is the morality of abortion, the context suggests counterfactual expansions for both spaces in which the mothers terminate their pregnancies (see table 8.2).

With the phrase in (5), Lee blends her past as an unwanted child with her 1963 predicament. The blend is a simple composition of the two input spaces in which the property of being an unwanted child is projected from the Past Outcome space, and provides an integrated unit that makes salient conflicting frames relevant to the current context (see table 8.3). In fact, elaborating the blend leads to a counterfactual scenario in which Lee qua pregnant rape victim aborts herself qua unwanted child.

Though Lee's initial expression of gratitude for the unavailability of abortion runs counter to the cultural framing of the rape exception, her

Table 8.2: Possible Consequences

Pre-1963 Expansion	**1963 Expansion**
Elements	**Elements**
Mrs-Ezell	Lee′
Lee-F	Julie-F′
Relations	**Relations**
Aborts (Mrs-Ezell, Lee-F)	Aborts (Lee′, Julie-F′)

Table 8.3: As an Unwanted Child

Past Outcome	1963 Situation	Blend
Elements	**Elements**	**Elements**
Mrs-Ezell	Lee′	Lee″
Lee	Julie-F′	F″
	Pregnancy′	Pregnancy″
	Rape′	
Relations	**Relations**	**Relations**
Child (Lee)		Child (Lee″)
Mother (Mrs-Ezell, Lee)		
Unwanted (Lee)		Unwanted (Lee″)
	Pregnant (Lee′, Fetus′)	Pregnant (Lee″, F″)
	Unwanted (Pregnancy′)	Unwanted (Pregnancy″)
	Victim (Lee′, Rape′)	

argument is essentially an attempt to reframe the scenario as culturally defined. For instance, she frames herself in (2) as a virgin teenager and a rape victim, thus undermining her own moral agency, not to mention any of the desires that might be associated with the intention to get pregnant. The rape victim's pregnancy is still being framed as a problem, and abortion is still a means of preventing the trauma associated with this sort of a pregnancy. But while informants left the relationship between the traumatic pregnancy and abortion unspoken, Lee contrasts these elements with a metaphoric blend. In (4), her unhappy situation is framed as a container ("in my crisis state"), and the abortion as a "seemingly easy out." Moreover, the "seemingly" in this latter phrase results from blending Lee's teenaged perception (where abortion is merely an "easy out") with her current perspective (in which she regards abortion as an immoral act).

As in Andrus' argument about the adverse effects of the anti-abortion bill he vetoed, Lee uses blending to personify the state of California as a moral agent capable of "taking advantage" of young girls and making deceptive offers ("this seemingly easy out"). Rather than compounding the tragedy of rape by *not* allowing abortion, California would have been "taking advantage" of her by allowing it. By applying models of moral

responsibility to the abstract state of California, it is possible to use the inferential and affective power that these models display in the domain of interpersonal relations in the more abstract domain of sociomoral issues. In this case, the Taking Advantage frame summons up certain emotions in the domain of interpersonal relations, such as sympathy for the victim and scorn for the agent, and frames the act as immoral. In Lee's case, this produces sympathy for her and scorn for California, and frames the availability of abortion as immoral.

In the cultural framing of the scenario, the injury of the rape is compounded by the insult of pregnancy. Here the injury of the initial rape is compounded by the insult of the state in its deceptive offer of a false solution to her problem. Moreover, the sexual connotations of "taking advantage" imply a parallelism between the initial rape and the subsequent deceptive offer. If the content of her letter is any indication of how Lee has constructed a meaningful narrative for the events in her life, cultural models of action, responsibility, punishment, and so forth are clearly present as schematic underpinnings. Moreover, rather than simply negating the cultural framing of the rape victim's pregnancy, she reframes it.

8.2.5 Why Punish Me?

Since it involves questions of agency and responsibility, the abortion debate often leads to discussion of counterfactual scenarios. This is the case with the following example in which a woman considers the morality of the abortion that her mother *could* have had. In the course of a national talk show, Lee Ezell's adult daughter, Julie, made the statement in (18) as an argument against abortion.

(18) Why punish me for my father's crime?

The context in (18) serves to set up a space that contains two elements, Julie's mother and Julie's father. The mother has been raped and is pregnant, while the father has the property of being a rapist. Example (18) concerns the morality of abortion in a counterfactual scenario in which Julie's mother terminated her pregnancy. As noted previously, one framing that is readily available is that pregnancy is unjust punishment of the mother, making abortion a moral possibility. In contrast, Julie frames the rape victim's abortion as unjust punishment of the fetus. In the base space, the fact that Julie's mother is pregnant allows us to set up an element that corresponds to the fetus, and that can be linked to

Julie in the current reality space. The counterfactual space results from conceptual blending of the Abortion scenario with structure in the Past space. Besides an element corresponding to Julie's fetus, the counterfactual space contains one element for Julie's mother and another for her father. Most importantly, the element that corresponds to Julie's fetus is linked both to Julie's fetus in the Past space and to the adult Julie in the Current Reality space.

When the adult Julie says, "Why punish *me*," she designates her counterpart in the counterfactual scenario. However, because the counterfactual structure results from integrating an abstract abortion scenario with the particulars of Julie's situation, the fetus in the blend inherits some of the adult Julie's properties. In particular, conceptual blending in the counterfactual space facilitates the construal of abortion as unjust punishment of the fetus. This construal is controversial in the normal case because the fetus' status as a moral agent is ambiguous. However, in this blend, the counterfactual fetus inherits Julie's status as a moral agent.

Julie's blend resembles one described by Aristotle that concerns a mythical woman who has undergone a transformation into a bird*:

> The address of Gorgias to the swallow, when she had let her droppings fall on him as she flew overhead, is in the best tragic manner. He said, "Nay shame, O Philomela." Considering her as a bird, you could not call her act shameful; considering her as a girl, you could; and so it was a good gibe to address her as what she was once and not what she is.†

While the bird lacks the social framing to perform shameful acts, the woman does not. Gorgias, like Julie, has set up a blend between past and present states of reality.

Blending agentive properties of her adult self with a counterfactual scenario based on her own past, Julie designates an element in a blended space that inherits properties from the adult and fetal Julies that maximizes her rhetorical goals. By employing a blended model, Julie (the speaker) is able to exploit the built-in affective dimension of the unjust punishment model, usually applied only to full-blown moral agents, to generate empathy for her counterpart in the blended space and to implicate the immorality of abortion.

The link between Julie's fetus-woman in the blend and the Julie-fetus in the Past space affords an extension of the sorts of elements to which

* Thanks to Mark Turner for pointing out this example.

† Rhetoric, Book III, chapter 3, ending [1406b].

the unjust punishment model can be applied. Because agency is a precondition for the punishment model, applying the unjust punishment model to an unborn fetus attributes the fetus with the very property that is in question in debates about abortion. Further, because of the parallel structure and the identity links, entertaining the counterfactual suggests consequences for the current reality space in which Julie ceases to exist.

8.3 MERE SEMANTICS

In contrast to many of the made-up examples and decontextualized jokes discussed in the previous chapters, the examples here come from real people talking seriously about things that matter. While the role of background knowledge in the simple examples was not always obvious, the "extra" meaning in moral discourse is more readily apparent. Our familiarity with the abortion debate also enables us to see the baggage carried by many of these terms. Ironically, the abortion debate in particular has been portrayed in the media as a dispute of "mere semantics": If only participants agreed on the definition of "life," all would be settled. However, the preceding analysis suggests that in arguments over the applicability of terms, what's at stake is not whether the *definition* obtains, but whether it is reasonable to adopt the background assumptions that support the terms' application.

Just as the meaning of *bachelor* depends on the applicability of a set of stereotypical assumptions, the meaning of *responsible* also depends on the applicability of a whole set of cultural models. Consequently, terminological wars can be seen as less of a fight over terms per se as over the implications of the models that those terms prompt speakers to construct. Because the culturally shared frames evoked by speakers are so idealized, it is an important function of culture to delimit the cases where models are allowed to apply. Andrus' argument is particularly relevant in this respect because his actions help determine what a *real* criminal is. As governor, his actions play a big role in which models will be institutionalized. Because of his socially defined role, his argument does not just reflect our cultural models, it effects them.

When asked to explain why rape is a compelling reason for abortion, informants responded by apparently rejecting the application of certain responsibility models to rape victims. The dissociation between pregnant rape victims and other women with unwanted pregnancies is made by noting two sorts of things. The first is the extent to which the woman can be held accountable for being pregnant. The second is

the disparity between trauma caused by pregnancy due to rape versus trauma caused by unwanted pregnancy due to consensual sex. The trauma that the unwanted pregnancy may cause the rape victim seems to go above and beyond the normal difficulties associated with pregnancy and child rearing, making it difficult for Americans to treat both types of trauma the same.

Clearly, framing is very important to the way we interpret and experience events, and the way we erect social structure in the world. Besides having different structural characteristics, different models imply different social consequences for the participants, entail different courses of action, and afford a way of changing a situation's experiential character. The framing of an unwanted pregnancy is contested precisely because different cultural models have substantially different moral implications, affective dictates, and physical and social consequences. Construal of the woman as being "at fault" entails moral responsibility, while her construal as a "victim" implicates none. What's more, we are supposed to feel sympathy for the "victim," but contempt for the woman who is "at fault." Finally, differences in framing can have real consequences, as one woman may be forced to carry her pregnancy to term while the other may not.

Pregnancy normally activates cultural models that dictate attitudes, associated actions, and entailed obligations. However, even in cases like the pregnant rape victim, where informants agree that the responsibility model no longer applies, the model remains as an input to the conceptual integration network that ultimately frames the rape victim's situation. In spite of the fact that the scenario differs from the assumptions of the responsibility model, pregnancy is still construed as the result of a moral agent's action, as entailing responsibility, and as engendering obligation. Although the woman no longer fills the agent slot, the pregnancy can still be construed as her responsibility and (consequently) her punishment.

When real-life scenarios diverge too much from the assumptions of the idealized models, speakers can employ processes of conceptual integration to transform them. However, the preceding examples suggest that entrenched models are not easily rejected. Rather, blending is used to reframe the scenario to yield novel inferential, affective, and social consequences. Luckily, while efforts to reframe are substantially constrained by cultural framing of the issues, a little reframing can go a long way. In the case of the pregnant rape victim, for example, the novel application of the responsibility model allows us to understand

her pregnancy as unjust punishment that licenses *her* right to abortion, even when other women might be denied a similar right.

Perhaps it is not surprising, then, that so much ink is spilled and breath drawn in people's efforts to frame the world around them. Framing is contested because it is so central to social experience. Another reason people argue about framing is that framing is *arguable*. The models we use are partial and idealized, and the mappings from models to situations are underdetermined. Moreover, establishing the mappings is only the first step in the more imaginative process of blending. Framing prompts speakers to integrate shared cultural models with conceived scenarios. However, it is quite possible to understand the mappings without truly integrating. This deeper integration is a somewhat voluntary process that requires creative composition, completion, and elaboration in the blended space so that the speaker's conceptualization of the scenario involves actually adopting the attitudes dictated by the model. Consequently, framing is not just knowledge instantiation: Integration can create novel construals, affective responses, and moral implications.

CHAPTER NINE

Frame-Shifting and Scalar Implicature

> Nearly one hundred Internet companies are preparing to entice Wall Street with initial public offerings. Numerous firms delayed their IPOs after many technology stocks that debuted during the summer slid below their initial selling price. . . . But many analysts believe the market may not be able to support so many new tech stocks. Some portfolio managers say it will be a matter of too few dollars for too many deals instead of too much cash for a few IPOs.
>
> – Peter Jon Shuler, for National Public Radio's *Morning Edition* (September 14, 1999)

Anyone with any doubt about the prevalence of natural language quantifiers need only open the morning paper, turn on the television, or tune in a radio broadcast such as the one excerpted above. In this short passage, we have "nearly one hundred Internet companies," "numerous firms," "many analysts," "so many new tech stocks," "some portfolio managers," "too few dollars," "too many deals," "too much cash," and "a few IPOs." Intuitively, one might hope for a straightforward mapping between quantifiers and quantities, perhaps expressed as a percentage point. But, as luck would have it, things are not that simple.

Moxey & Sanford (1993) demonstrate that the comprehension of quantified expressions requires knowledge of the domain in question. For example, "numerous ants" presumably refers to more ants than the number of firms referred to by "numerous firms" in the preceding excerpt. Moreover, Moxey & Sanford point out how inquiries about quantity are frequently directed at eliciting advice about the advisability of pursuing a particular course of action. They suggest that a telephone enquiry to a ticket seller about how many train tickets are currently available might be inspired by the need to know how *soon* the desired

tickets should be booked, as opposed to the exact number of unsold seats. Even utterances that would seem to admit a compositional interpretation require selective bits of background knowledge and enough contextual information to infer the speaker's goals.

In fact, speakers often use quantifiers when they are speaking about entities that occur in the context of a pragmatic scale, a group of objects or scenarios ordered by some relevant semantic dimension. As outlined in section 1, the existence of pragmatic scales enables speakers to use information about one object on the scale to infer information about other items. Examples in section 1 will demonstrate how speakers take advantage of such scales to generate scalar implicatures, stating information about one item on the scale to imply information about the other items. To illustrate the use of pragmatic scales in realistic discourse contexts, section 2 examines scalar implicature in abortion rhetoric. Examples reveal the cultural grounding of these scales, and how speakers strategically link (and sever) scales that allow them to project inferences from the degree of one aspect of a situation to the degree of some other aspect.

9.1 SCALAR IMPLICATURE

A pragmatic scale consists of objects or scenarios ordered along some relevant semantic dimension. Once ordered, statements that concern one member of the scale implicate propositions about its other members. For instance, consider (1):

(1) Al Gore can solve the hardest equations.

To say that Gore can solve the hardest equations is also to implicate that he can solve the less difficult equations. When one sentence (P) *implicates* another (Q), it means that speakers will assume Q on hearing P, unless some qualification is made. Because it can always be canceled, implicature differs from entailment. Linguistic tests for implicature involve the difference in judgments about sentences conjoined with *but* compared to the same sentence conjoined with *and*. If P implicates Q, then "P and Q" or "P but not Q" will sound natural; however, "P but Q" will sound unnatural.

(2) Al Gore can solve the hardest equations and the easy ones, too.
(3) Al Gore can solve the hardest equations, but not (always) the easy ones.
(4) Al Gore can solve the hardest equations, but the easy ones, too.

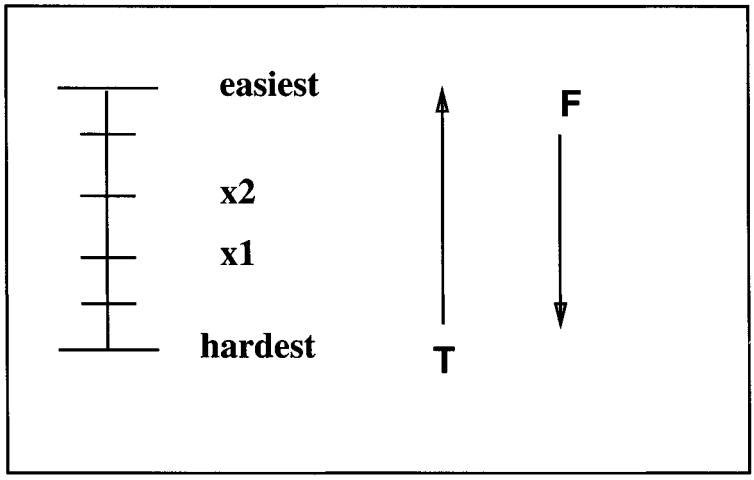

Figure 9.1 Pragmatic Scale of Equations.

These patterns of implicature fall naturally out of interpretation on a pragmatic scale. To interpret (1), one can construct a pragmatic scale of equations ordered by their difficulty. When interpreting a statement that concerns an object on a pragmatic scale, an assertion about the object implicates assertions about objects that fall above it on the scale; to deny something about the object implicates denials about objects that fall below it on the scale. Thus asserting that Gore solved x1 (see figure 9.1) implicates his ability to solve all the less difficult problems that occur above it on the scale ("the easiest"). By contraposition, denying that Gore solved x1 implicates his inability to solve the equations that occur below x1 on the scale ("the hardest").

As a way of accounting for the complex semantics and pragmatics of "let alone," Fillmore, Kay, & O'Connor (1988) generalize the idea of a pragmatic scale to encompass comparisons along more than one dimension. Where a pragmatic scale organizes objects along a single dimension, a scalar model organizes propositions along a variable number of dimensions. Formally, a scalar model is defined as the ordered pair $\langle S, T, D, P \rangle$, where:

S is a set of states of affairs
T is a truth value, either 0 or 1
D is a finite set of semantic dimensions
P is a set of propositions corresponding to the points in D.

For example, a scalar model of my favorite basketball teams might be ordered along two dimensions, one being the teams, the other the number of games I watched that team play last season. The propositions in P all concern the number of games I watched each team play, and T is the set of truth values for each proposition. Assertions interpreted on a pragmatic scale entail the truth of similar assertions about objects located *above* them. The scalar model uses the concept of proximity to the origin in a similar manner: Asserting a proposition P1 in a scalar model entails the truth of all other propositions that are closer to the origin than P1.

This particular model, assumes that my basketball viewing habits are such that the number of games that I watch my favorite team play will always be greater than or equal to the number of times I watch the less favored teams. So, if the teams are ordered ⟨*Bulls, Lakers, Clippers*⟩ (from most to least favorite), and the number of games ranges from one to twenty, the proposition, "Seana watched the Bulls one time," is closest to the origin. "Seana watched the Bulls two times," entails, "Seana watched the Bulls one time." Similarly, "Seana watched the Lakers two times," entails, "Seana watched the Lakers one time," as well as the earlier propositions about the Bulls.

An important feature of the scalar model is that it provides a formal way of defining what is maximally informative. The point in D that is furthest from the origin and whose truth value is 1 corresponds to the maximally informative proposition. Intuitively, we can infer more about the games I watched from the number of times I watched my *least* favorite team than from the number of times I watched the Bulls. Assuming people state the most informative proposition that they can, the statement that I watched the Clippers twice generates the implicature that I watched the Lakers and the Bulls *at least* two times each. Formally, this reflects the fact that all points closer to the origin than our maximally informative point will be true, while all points farther from the origin will be false. By contraposition, the falsity of a maximally informative proposition in P implies the falsity of all other propositions whose ordered pairs are farther from the origin.

Fillmore, Kay, & O'Connor (1988) show how the entailment properties of the scalar model can be used to describe and predict the acceptability of sentences that employ the "*let alone*" construction. For instance, given the context of my favorite basketball teams and the number of games I watched last season, (5) is permissible while (6) is not. Similarly, (7) is permissible while (8) is not; and, (9) is permissible while (10) and (11) are not.

Seana must have seen the Clippers play at least ten times last year, huh?

(5) Seana didn't watch the Clippers play five games last year, let alone ten.

(6) # Seana didn't watch the Clippers play ten games last year, let alone five.

(7) Seana didn't watch the Bulls play ten games last year, let alone the Clippers.

(8) # Seana didn't watch the Clippers play ten games last year, let alone the Bulls.

(9) Seana didn't watch the Bulls play five games last year, let alone the Clippers ten.

(10) # Seana didn't watch the Clippers play five games last year, let alone the Bulls ten.

(11) # Seana didn't watch the Bulls play ten games last year, let alone the Clippers five.

Fillmore, Kay, & O'Connor argue that sentences like (5)–(11) are interpreted as two propositions on the same scalar model, where one proposition is more informative than the other. Further, the acceptability of a sentence that includes "let alone" depends on the entailment relations as defined by the contextually relevant scalar model. In the acceptable sentences, such as (5), (7), and (9), the first proposition is always located closer to the origin than is the second. Because the denial of the first proposition implicates the denial of the second, the first proposition is more informative. For example, in the scalar model set up by (5), watching five Clippers games is closer to the origin than watching ten Clippers games. Because we can infer that I didn't watch ten games from the fact that I didn't watch five games, the statement that I did not watch five Clippers games is more informative. By contrast, in the unacceptable sentences (such as (6), (8), (10), and (11)), the first proposition is farther from the origin than the second, and the same entailment relation does not obtain.

Scalar models have also been used to account for the semantic contribution of "even" (Kay, 1990) and "still" (Michaelis, 1996). Kay's (1990) analysis of "even" is that it marks a clause as more informative than another clause interpreted on the same scalar model. Kay argues that the violation of expectations that is so often encountered in propositions marked by "even" is a side effect of its relative position in the scalar model. If Q is more informative than P, and speakers assume

that *P* is the maximally informative proposition, then the statement "P and even Q" will naturally violate their expectations. Michaelis' (1996) analysis of "still" as a scalar operator accounts for the polysemy of "still" by noting three different types of scales: continuance through time, graded category membership, and the relative likelihood of certain situation/outcome pairings.

Moreover, Israel's (1996) work on the semantics of polarity-sensitive items reveals the operation of scalar models to be even more general. He argues that polarity-sensitive items encode two semantic features that serve basic functions in the meaning construction process: *quantitative value* and *informative value*. Quantitative value concerns an element's exact position in the contextually established scalar model, while informative value concerns an element's *relative* position with respect to the model's other propositions. Israel relates grammatical properties of negative and positive polarity items to pragmatic properties derived from the way they encode quantitative and informative values. For instance, negative polarity items such as "drink a single drop" or "sleep a wink" encode low quantitative values and high informative values. They usually occur in negative contexts because that is where they are strongly informative, allowing speakers to make full use of the inferential possibilities of the scalar model.

Research in semantics and pragmatics discussed in this section indicates that our commonsense reasoning abilities include the capacity for scalar implicature. The fact that this pragmatic capacity is routinely exploited in the lexical semantics of words such as "even," "let alone," "still," and other scalar operators suggests that it may reflect a central cognitive process. However, it is not obvious that these examples are relevant to anyone besides the linguist interested in the acceptability of the sentences in which the items appear. While the linguists' interests are not to be made light of, one might nonetheless wonder about the import of scalar reasoning in other domains. This issue is addressed in the following section, as we see how scalar reasoning is used at the higher discourse levels that are the province of rhetoric.

9.1.1 Pragmatic Scales and Rhetorical Strategies

While the linguist points to the contribution of a particular word to the utterance as a whole, the rhetorical theorist points to the contribution of a particular utterance to the argument as a whole. Similarly, while the cognitive semanticist is good at uncovering and making explicit the abstract

conceptual underpinnings of the linguistic constructions we use in our arguments, the rhetorical theorist is adept at explaining how arguments are themselves structured. The job of the rhetorical theorist is to classify arguments according to the relationship between argumentative goals (the different sorts of things we argue for) and argumentative means.

The inferential properties of statements interpreted on scales lend themselves quite naturally to a variety of argumentative strategies. The most obvious of these is argument a fortiori. For example, if one wanted to argue that Al Gore could solve a particular equation – such as x1 – one could do so by stating (1). Because (1) admits of interpretation on the pragmatic scale previously described, it generates the inference that Gore can solve x1. That is, Gore can solve the hardest equation, so a fortiori he can solve x1. In fact, the "let alone" construction appears to provide English speakers with a ready-made way of expressing argument a fortiori (Fillmore, Kay, & O'Connor, 1988).

From the objectivist perspective, the occurrence of a construction (like the "let alone" construction) that routinely flouts conversational maxims by introducing a proposition that is less than maximally informative is puzzling. The puzzle, of course, is not that deviations from the maxims would be linguistically marked, but why such deviations would ever occur in the first place. However, from the cognitive perspective, the occurrence of the "let alone" construction might be motivated by its rhetorical properties. On such an account, the "let alone" construction provides speakers with a conventional means of expressing a persuasive argument.

Pragmatic scales can also be used in the rhetorical device known as the *double hierarchy* argument described by Perelman & Olbrechts-Tyteca (1969). In a double hierarchy argument, a speaker uses the ordering principle of a well-established pragmatic scale to create a second pragmatic scale whose points are ordered in the same way. Once the two scales are linked, inferences that hold on the first scale also hold on the second. Linking two scales in this manner affords the possibility of moving from inferences generated on the less controversial scale to parallel inferences on the more controversial linked scale. For example, suppose we wanted to argue that Al Gore's academic abilities included spelling as well as equation solving. We might begin our argument with (12) and (13).

(12) In my experience, the best mathematicians are also the best spellers.

(13) Al Gore is one of the best mathematicians.

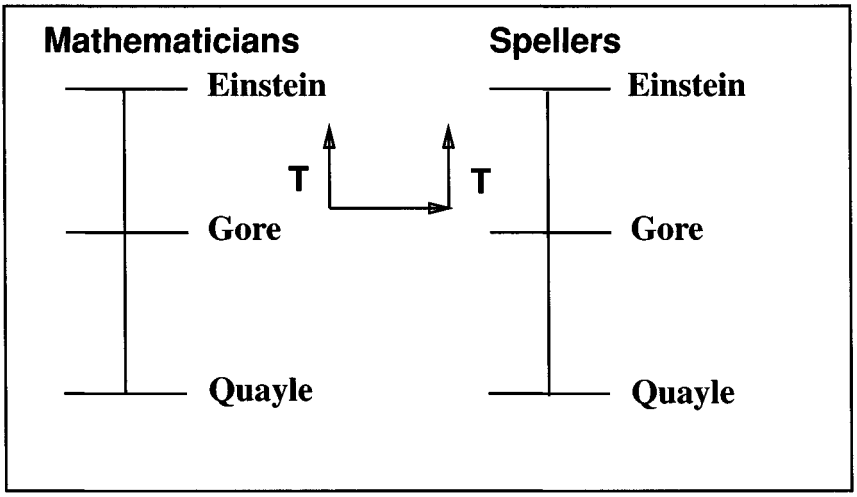

Figure 9.2 Pragmatic Scales: Mathematicians and Spellers.

In this argument, (12) would serve to link the mathematician scale with the speller scale. Because (13) locates Gore near the top of the mathematician scale, it serves as an argument for placing Gore near the top of the speller scale. The argument in (12) and (13) would thus serve to implicate (14) (see figure 9.2).

(14) Al Gore is one of the best spellers.

Because the implications of a given statement interpreted on a scalar model are a function of its ordering principles, the success or failure of a double hierarchy argument will often turn on the way speakers appeal to relevant cultural models. Consequently, the paired scale strategy goes beyond spellers and can be seen in all sorts of discourse contexts.

9.2 RETURN TO THE RAPE EXCEPTION

In a return to the discussion that began in Chapter 8, this section explores the use of double hierarchy arguments in moral discourse on abortion. These attested examples suggest connections between the treatment of scalar implicature and the frame-based constructivist approach to meaning construction. In particular, the focus is on how Americans on all sides of the abortion debate are forced to contend with the existence of a scale of acceptable reasons for abortion. While no single ordering principle is available to organize the acceptability scale, end points such as the rape

exception are clearly demarcated and must be dealt with – even by people who don't accept the validity of the scale. Where proponents of the rape exception capitalize on the organization of background knowledge to link less controversial scales to the acceptability scale, opponents of the rape exception operate by evoking a different set of cultural models to break links between scales.

9.2.1 Linking Scales

As we saw in the discussion of abortion rhetoric in Chapter 8, informants often justify the rape exception by rejecting the applicability of various cultural models typically applied to women with unwanted pregnancies. While informants differed in the extent to which they accepted the application of these models in the "normal" case of unwanted pregnancy, they seemed to agree that the rape exception reflects the particular unsuitability of those models for the pregnant rape victim. In the course of this discussion, one of the speakers uses a double hierarchy argument linking a scale of traumatic pregnancies to a scale of acceptable abortions.

In this case, the double hierarchy emerged relatively naturally from the interviewer's strategy of asking about a series of cases that ranged in how controversial they were. Here, a hypothetical case involving a pregnant rape victim was followed by a question about a hypothetical case involving a pregnant incest victim. Asked about whether abortion is a morally acceptable course of action for a rape victim, P replied in the affirmative, arguing that to require the rape victim to continue her pregnancy would constitute unjust punishment. Asked about the case of the incest victim, P replied as follows:

(1) P: More of a reason probably.
(2) S: More of a reason. Why is it more of a reason?
(3) P: Because then you're dealing with rape and you're dealing with incest.
(4) I have a feeling it would be more traumatizing to be raped by a member of your family.

At the outset of this excerpt, P has already established that rape is traumatic and that it creates a very strong case for abortion. To use the argument from double hierarchy, P shows that incest is a more traumatic cause of pregnancy than is rape. She does this by noting first a qualitative difference between the two cases: The rape victim must deal with the trauma of one sort of event (rape), while the incest victim must deal

with the trauma of two sorts of events (both rape and incest). Following with the statement in (4), P ends the argument with the speculation that incest rape is "more traumatizing" than rape by a stranger.

The specific claims in P's discourse function as an argument for the acceptability of abortion because they set up a double hierarchy that links a scale of pregnancies ordered by how traumatic they are to a similarly ordered scale of the acceptability of abortion in those cases. Having established rape as a case that warrants the acceptability of abortion, P's arguments as to why incest is more traumatic than rape locate incest higher on the scale of traumatic pregnancies. Since inferences that hold on one pragmatic scale in the double hierarchy argument are assumed to hold on the other, locating incest higher on the trauma scale generates the inference that incest, too, warrants the acceptability of abortion.

Here P capitalizes on the relatively uncontroversial hierarchy of traumatic pregnancies to order the more controversial hierarchy of acceptable abortions. Formally, any two scales can be linked together in a double hierarchy. However, in practice, the possibility of linking scales is constrained by the organization and content of background knowledge. The success of the argument depends on the legitimacy of the cultural model that links the scales. For example, the link between the two hierarchies in P's argument is the model of abortion as a means to prevent the trauma caused by unwanted pregnancy. The argument described in the following section attempts to sever scales commonly linked to legitimate abortion, including the double hierarchy that P sets up between trauma and acceptability.

9.2.2 Breaking the Links

Given that arguments *for* the rape exception involve framing the pregnant woman as a nonagentive victim and the pregnancy as a traumatic experience, we might expect arguments against the rape exception to promote alternatives. The speaker here is Bernard Nathanson, a physician who in the 1960s was one of the leading proponents for the legalization of abortion, and in the 1980s became an advocate for the Pro-Life movement. The following statement is excerpted from his writings.*

(1) When pregnancy is confirmed,
(2) one can no longer discuss morality in the framework of

* Quoted in Whitney, C. (1991). *Whose Life?* New York: William & Morrow Co., 167.

(3) rape or incest
(4) (or whether birth control was or was not used,
(5) or whether the child is "wanted" or whatever).
(6) The unwanted pregnancy flows biologically from the sexual act,
(7) but not morally from it.
(8) ... even degradation, shame and emotional disruption
(9) are not the moral equivalent of life.
(10) Only life is.

In this passage, Nathanson acknowledges the existence of a scale of acceptable reasons for abortion, but wants to insulate himself from the models that make the scale possible. Consequently, the argument can be viewed as rebutting at least two double hierarchy arguments often used to legitimize abortion. One argument links the moral evaluation of the cause of pregnancy to the moral evaluation of the pregnancy itself. The other argument, discussed in Chapter 8, links the trauma caused by pregnancy to the acceptability of abortion in that case.

In one line of argumentation, Nathanson attempts to rebut a *hierarchy from cause to effect*. This argument links the accepted hierarchy of the moral status of the cause of pregnancy to a hierarchy of the moral status of pregnancy itself. The link between the hierarchies in this argument is a *liaison of succession* – the relationship between a cause and its effect – and is based on the cultural model of sex as the cause of pregnancy. Thus pregnancy that results from an illicit sexual encounter is morally contemptuous and should be eliminated, whereas pregnancy that results from the desire to provide a loving home for a child has a high value.

To challenge a double hierarchy argument, one can either dispute the placement of one or more terms in the accepted hierarchy or dispute the liaison that connects the two (Perelman & Olbrechts-Tyteca, 1969). Nathanson chooses the latter strategy, arguing in (6) and (7) that while the rapist causes the pregnancy, the morality of the rapist does not cause the morality of the pregnancy. Nathanson does not dispute the causal nature of the link between the two hierarchies, but rather the legitimacy of projecting moral judgments from cause to effect. Because the models that link the scales in this double hierarchy argument are typically backgrounded, explicitly noting their existence can force us to examine their use with a critical eye.

Nathanson's argument is also directed at the blend that enables the construal of the rape victim's pregnancy as punishment. As discussed in Chapter 8, our informants S and P both attempted to reject the framing of

the rape victim's pregnancy as the rape victim's responsibility. So, by arguing in (6) and (7) that the cause of pregnancy is irrelevant, Nathanson attempts to reject the models that afford the construal of pregnancy as punishment. Having rejected the rape frame as irrelevant to the construal of the rape victim's pregnancy, Nathanson goes on to evoke an alternative model of pregnancy in which the moral responsibility for pregnancy originates in the state of being pregnant rather than the action that causes it.

To show the irrelevance of the means of getting pregnant to the acceptability of abortion, Nathanson employs *equalization*, a rhetorical technique used to accentuate the similarity between two different kinds so as to promote the transfer of inferences and moral consequences (Perelman & Olbrechts-Tyteca, 1969). Just as proponents of the rape exception distinguish the rape victim's pregnancy from the idealized reproduction scenario, opponents of the rape exception are obliged to point to their similarities. If we can come to see marginal cases as equivalent in some way to the planned pregnancy, we might be more likely to judge the ambiguous case similarly to the way we judge the unambiguous case in the idealized reproduction scenario.

Moreover, Nathanson also targets the double hierarchy from trauma to acceptability, as exemplified by P's argument. In this double hierarchy, an established scale of traumatic pregnancies is linked to the acceptability of abortion, and the link between the two hierarchies is a cultural model of abortion as the means to preventing the trauma caused by the pregnancy in each case. Marked by "even," the statement in lines (8) and (9) indicates that the stated proposition ("degradation, shame, and emotional disruption are not the moral equivalent of life") is more informative than some contextually established proposition interpreted on the same scale (Kay, 1990). Because the trauma to acceptability argument relies on the model of abortion as the means to prevention of trauma, it establishes "degradation," "shame," and "emotional disruption" as morally relevant factors that abortion will prevent.

With the statement in lines (8)–(10), Nathanson attempts to equalize the marginal cases by invoking the deontic principle of life as the only relevant factor. In presupposing the existence of the double hierarchy from trauma to acceptability, Nathanson's statement displays an understanding of the different factors involved in the abortion model that links those scales. He explicitly acknowledges factors such as the alleviation of trauma and the lessening of obligation, but shifts the focus to what he sees as the most relevant: the termination of the fetus. Rather

than admit that abortion can be more or less acceptable, he presents life as something that cannot be placed on the scale.

In the preceding examples, I have sometimes implied the speakers were engaged in an actual debate. Although this is not the case, this manner of speaking is possible because their arguments implicitly counter the tenets of publicly available positions. Nathanson's argument, in particular, demonstrates the degree to which parties on all sides of the debate understand the arguments of their opponents. In the next example, we turn to arguments constructed by people engaged in a dialogue about the morality of abortion. These observations suggest an important role for scalar implicature in on-line interaction as well as in off-line arguments constructed over the course of time.

9.2.3 Dialogue

The following excerpt concerns a discussion between Ron Reagan, a talk show host, and Joseph Scheidler, the executive director of Pro-Life Action, with an interjection from Randall Terry, director of the anti-abortion activist group Operation Rescue. As the host of the talk show, Reagan primarily seeks to provoke interesting discussion by posing challenging questions to guests and attempting to elicit controversial statements from them. On the other hand, Scheidler is on the show as a representative of a Pro-Life organization. His primary goal is to inform the audience of his position and to persuade them of the morality of that position.

In this excerpt, Reagan asks Scheidler whether abortion is acceptable in the case of rape, then follows with the case of incest. As we have seen previously, in cases of rape and incest, the acceptability of abortion is more ambiguous than in cases in which no rape has occurred. People who are otherwise opposed to abortion are most likely to admit the acceptability of abortion in the case of rape or incest than any other. The existence of the acceptability scale, then, actually makes Reagan's question relevant. Though the people in the audience know Scheidler opposes abortion, they can understand that rape might function as an exceptional case.

(1) Reagan: OK, Joseph, if one of your daughters is raped and becomes pregnant,
(2) should [she] be able to have an abortion?
(3) Scheidler: I become a grandfather then.

(4) Reagan: If, if, God forbid, and I'm not suggesting this is a possibility,
(5) you committed incest with your daughter and she became pregnant–
(6) Terry: The professionalism of this show is just going down the crapper.
(7) Reagan: No, it happens. Incest is a reality
(8) and girls become pregnant because their fathers impregnate them.
(9) Scheidler: Ron, that is the most unnatural thing I can think of.
(10) Reagan: I agree but the,
(11) should the result of that unnatural act be expunged in any way?
(12) Scheidler: It is so repugnant, but to answer your question,
(13) I would never kill a child because of the sin of the father.
(14) I just would not do it.
(15) Reagan: That's all I wanted to know.
(16) Got to go to a commercial. Be right back.

Reagan begins by setting up a hypothetical situation in which Scheidler's daughter is raped and becomes pregnant. This is represented with the mental space configuration depicted in figure 9.3. The first space is

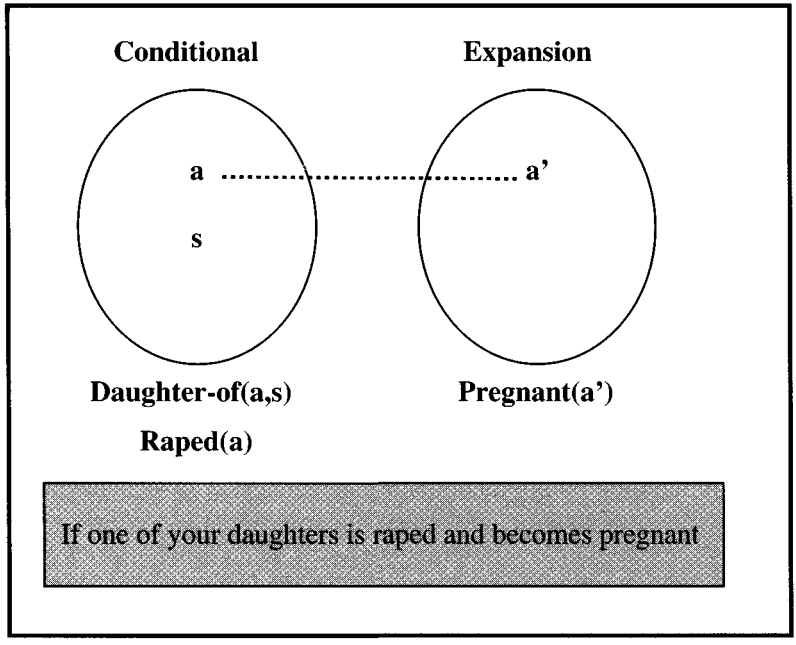

Figure 9.3 Reagan's Hypothetical Scenario.

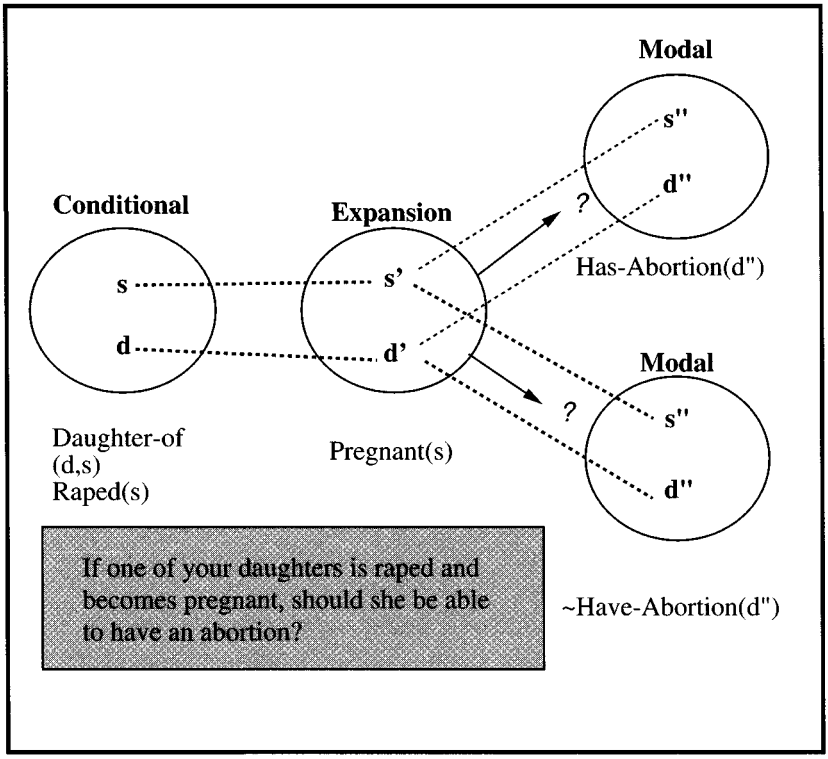

Figure 9.4 Reagan's Question.

a hypothetical space structured with the information about the rape of Scheidler's daughter; the second space represents the expansion of this scenario in which a' has the additional property of being pregnant.

By asking a yes/no question, Reagan poses his question as a choice between two modal expansions of the hypothetical scenario: one in which the daughter should have an abortion, and one in which she should not (see figure 9.4). Moreover, Reagan's question is pertinent because of the existence of a roughly ordered scale of acceptable reasons for abortion. If Scheidler answers this question in the affirmative, he admits the existence of a scale and so compromises his Pro-Life position. By inviting Scheidler to make a denial about a case near the origin, Reagan sets him up to make a potentially emphatic statement.

But while it seems that Reagan has posed the question in a way that forces Scheidler to deal with the continuum of acceptable reasons for abortion, Scheidler manages to avoid the scale altogether. Rather than answering Reagan's question directly, Scheidler expands the

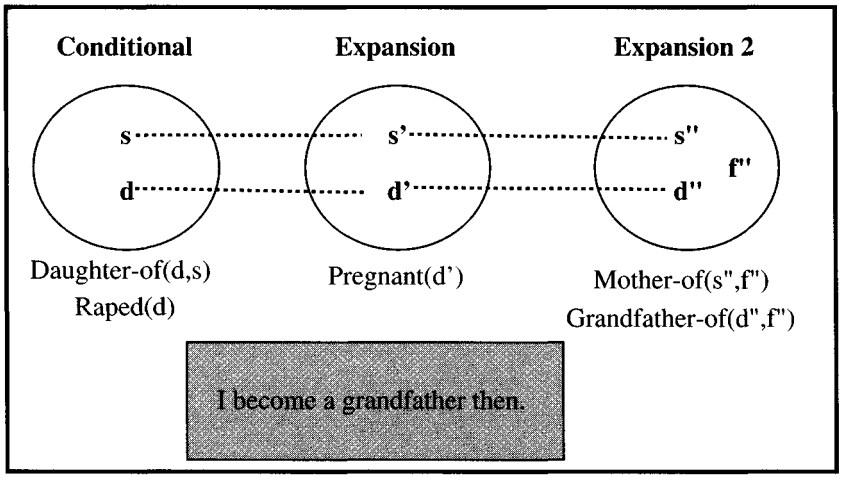

Figure 9.5 Scheidler's Reply.

hypothetical situation that Reagan has set up (see figure 9.5). With this remark, Scheidler has employed completion of the hypothetical scenario set up in Reagan's question. The familial relationship between Scheidler and his daughter, as well as the familial relationship that Scheidler believes to exist between his daughter and her unborn child, is completed with the Grandfather frame.

Understanding Scheidler's response requires listeners to invoke a number of cultural models with which they may or may not agree. First, we know that one becomes a grandfather when his son or daughter has a child. However, Reagan has not brought up a scenario in which Scheidler's daughter has a child, but one in which she is pregnant. For Scheidler's statement to make sense, he must believe that pregnancy is equivalent to having a child. This belief is also consonant with his conviction that it is immoral to terminate a pregnancy. Moreover, it is wrong to kill one's grandchild. So if in this case abortion is killing one's grandchild, the daughter should not be allowed to abort.

To understand Scheidler's reply, it is necessary to rearrange the information already established into the context – just as in the jokes discussed in Part I. Scheidler has effectively frame-shifted Reagan's initial scenario. Reagan's initial inquiry presents Scheidler with a state of affairs concerning his daughter and invites him to respond to how he would feel about his daughter getting an abortion in these circumstances. The sort of response that Reagan has invited depends on evoking the cultural models of rape and unjust punishment discussed in the last chapter.

Instead, Scheidler blends the Kinship frame with the established scenario. His shift maintains identity mappings, but leaves the Rape frame behind.

Perhaps because we are usually successful in interpreting the relevance of each other's statements, we tend to underestimate the extent to which meaning is underdetermined. For instance, Reagan's hypothetical scenario can be expanded in an indefinite number of ways. Scheidler might reply that his daughter loves being pregnant (and so would never consider abortion); that she might be driven to a suicidal frenzy (making abortion a moot point); or even that she is sterile and thus pregnancy is not possible.

Scheidler's actual expansion accentuates some dimensions of Reagan's initial scenario while downplaying others. For example, Reagan presumably mentions the case of Scheidler's daughter to appeal to the sense of compassion engendered by the father and daughter's familial relationship. Scheidler employs a similar strategy by evoking the relationship he sees as existing between himself and the fetus. Scheidler's response highlights the existence of the fetus, implicates the immorality of abortion, and makes it much harder to interpret the response in the context of the acceptability scale. While Nathanson explicitly rejects the Rape frame as irrelevant, Reagan merely shifts to the Kinship frame that renders the moral judgment unambiguous.

Reagan's strategy of personalizing the scenario of the rape victim backfires when he extends it to the incest victim. Whereas the strategy in the first case was intended to make Scheidler empathize with the rape victim, Reagan's suggestion that Scheidler's daughter might be the victim of incest serves only to offend Scheidler because (if only hypothetically) he fills the role of the incestuous rapist. Scheidler's response attempts to dismiss the question on the grounds that it is far-fetched. However, Reagan substantiates those grounds by arguing that the possibility of pregnancy as the result of incest is not far-fetched at all. Moreover, he depersonalizes the question by moving to general terms. Scheidler's second response is that he is reluctant to answer the question because to do so would involve considering a scenario he finds repulsive. Reagan, to get an answer to his question, says, "I agree, but, the, should the result of that unnatural act be expunged in any way?"

Reagan has alluded to the double hierarchy that links the morality of the cause to the morality of the effect. Reagan's choice of words here may be motivated by the desire to provoke Scheidler or, alternatively,

may reflect his view on the issue. Note, for example, that Reagan refers to the fetus in a dehumanizing manner ("the result of that unnatural act") bound to offend Scheidler. Moreover, Reagan is not construing the abortion as a medical procedure that results in the loss of life, but rather as the ablation of an earlier moral wrong.

Reagan's and Scheidler's statements reflect differing conceptions of the ontology of the fetus. Although they share the proposition that abortion eliminates the fetus, their differing cultural models have different inferential properties and correspondingly different moral implications. Reagan's model presupposes that a pregnancy that results from incest never should have happened in the first place. Abortion has the effect of making the situation equivalent to one in which incest had never happened, and thus of expunging an unnatural act. On the other hand, Scheidler's model presupposes that all pregnancies (including those that result from incest) represent the beginning of life. This makes the idea of pregnancy as an event that never should have happened much less plausible. Instead, the pregnancy that results from rape or incest is the good outcome of a bad event, and abortion to expunge an unnatural act is misdirected.

Because the audience is quite aware of Pro-Life models, Scheidler's statement, "I would never kill a child," readily evokes a model of abortion with slots for Agent and Child. Interestingly, the Agent in this example maps onto Scheidler as a potential advocate of abortion. Although not explicitly evoked, the woman of the hypothetical scenario is presumably available as a referent. As Scheidler continues with "because of the sins of the father," he evokes the frames for Incest and Unjust Punishment. Playing off of the culturally available construal of the incest victim's pregnancy as unjust punishment of the woman, Scheidler proposes another way of integrating the Unjust Punishment model onto the scenario.

The Agent is mapped onto the father who commits incest; the Action is mapped onto the incest ("sins"); the Victim is mapped onto the unborn child; the Punishment is mapped onto the abortion ("kill a child"). Conceptual integration results in the inference that the abortion, construed as unjust punishment, is immoral. The listener who evokes the pregnancy-as-unjust-punishment framing cued by Reagan's question will have to reorganize her conceptual representation dramatically upon hearing Scheidler's actual response.

9.3 CONCLUSION

Our examples underscore the fact that people's discourse about the morality of abortion does not consist of demonstrating how certain conclusions follow from agreed upon premises. Rather, it involves presenting alternative framings for morally ambiguous scenarios. The participants in the arguments examined here are engaged in either building or avoiding pragmatic scales, and appealing simultaneously to a whole host of inferences supported by various cultural models.

Moreover, analysis reveals yet another example of how semantic and pragmatic levels of meaning construction operate in reasoning, narrative structure, and other high-level aspects of communication. First, the linguistic accounts of scalar operators reviewed in section 1 demonstrate the import of pragmatic factors in the routine interpretation of individual words and phrases. Moreover, in demonstrating real-life examples of scalar implicature in abortion rhetoric, examples in section 2 show how the same sorts of structures that speakers use to interpret words in a sentence are also employed to interpret sentences in a text and speakers in a dialogue.

To understand P's argument, it is necessary to construe trauma in a scalar fashion, to link the degree of trauma to the acceptability of abortion, and to derive the implicature in accordance with the dictates of the scalar model. Moreover, Nathanson's argument proceeds on the assumption that other speakers are likely to make such links of their own accord. Similarly, the sequencing of Reagan's questioning makes sense only if it is understood as occurring in the context of a scale of acceptable reasons for abortion and a group of speakers who will compute the appropriate scalar implicatures.

Although the abortion debate has been portrayed as a case where different people hold mutually incompatible models, the models themselves are available to all speakers. It is not the case that people on different sides of the debate do not understand one another's models. In the examples examined in this chapter, speakers understood their opponents arguments well enough to frame these ambiguous scenarios strategically so as to support unambiguous moral judgments. Though speakers were indeed constrained by cultural framing of events such as rape, pregnancy, and abortion, they were frequently able to capitalize on the dynamic character of meaning construction to reconstrue any given scenario to suit their own ideological position.

Perhaps the most banal observation is that speakers with different ideological perspectives sometimes appeal to different cultural models.

For example, several models of abortion are used in discussions of the rape exception. Proponents of the rape exception tend to use a model of abortion in which it is a means by which the rape victim can alleviate the trauma caused by impregnation and pregnancy. In contrast, opponents of the rape exception use a model of abortion as the intentional killing of a child.

More interesting is how speakers on all sides of the debate make use of the *same* models. For example, all parties in the disputes examined in this chapter admit the existence of the unjust punishment model and find it to be a morally compelling frame. The only discrepancy is in the way different speakers integrate the punishment model with the scenario involving the pregnant rape victim. Integrating the rape scenario with the concept of responsibility from the normal reproduction model frames the pregnancy as unjust punishment of the victim. Integrating the punishment model with the abortion scenario implies that the abortion is unjust punishment of the fetus.

Arguments made by opponents of the rape exception are not aimed at explicitly denying claims made by its proponents, but at evoking a different set of cultural models to frame the disputed scenario in a new way. Opponents of the rape exception, for example, do not argue with the idea that the rape victim's pregnancy is traumatic, or that the victim is somehow responsible for getting pregnant. As noted previously, Nathanson does not argue with the pregnant rape victim's placement on the trauma scale. Rather, he evokes an alternative model of abortion that makes it more difficult to forge the link between trauma and acceptability.

However, as we noted in Chapter 8, merely posing a framing for an ambiguous scenario is not sufficient for moral disambiguation. Listeners need not fully integrate framings suggested by their argumentative opponents, and frequently reframe the scenarios for themselves. Moreover, explicitly dismissing the relevance of a scale such as the acceptability scale does not make it go away. Indeed, the examples in this chapter suggest that the best way to avoid the scales that surround the rape exception just may be to change the subject. Scheidler, in the excerpt from the Ron Reagan show, demonstrates great facility in his efforts to avoid the inferences generated by the acceptability scale. He simply refuses to acknowledge the existence of his opponents' models by making statements that force the listener to invoke the models in a way that supports his Pro-Life position. Rather than be drawn into an argument about the logic of the rape exception, Scheidler employs frame-shifting to complicate interpretation in the scalar context.

Formally, the frame-shifting that occurs in moral arguments is identical to that which occurs in jokes, but the discourse events examined in this chapter reveal subtle differences between the two. In the joke, pragmatic ambiguity stems from information deliberately withheld by the joke teller, and frame-shifting is triggered by seemingly accidental misanalysis. However, in moral argumentation, pragmatic ambiguity is an inherent facet of the situation, and the frame-shifting is of a more deliberate character. So, when Scheidler frame-shifts structure evoked by Reagan's statement, the implication is that the initial (preshifted) interpretation was not merely misguided, but downright wrong.

Both jokes and moral arguments turn on the way the significance of an event can vary from context to context. Consequently, slight differences in background assumptions can result in dramatic changes in the event's evaluative significance. In fact, the point of frame-shifting in an argument is to demonstrate the tenuous nature of the assumptions of one background frame or another. Because of the taken-for-granted nature of cultural models, speakers may not always be aware that more than one frame is available to provide background assumptions. Moreover, speakers may not always be aware of the extent to which the conclusions depend on assuming one frame rather than another. By switching from one frame to another, a speaker induces frame-shifting in listeners, thereby accentuating the existence of alternative background assumptions and their differing implications.

The ability to frame morally ambiguous scenarios differentially constitutes an important cognitive capacity that warrants further investigation. By employing diverse rhetorical strategies, speakers adapt cultural models to suit a variety of ideological outlooks and argumentative needs. In this way, we negotiate the salience of various parts of the models and supplement our understanding of social reality. The dynamic negotiation occurring in conversational debates can thus be seen as a microcosm of the cultural negotiation of reality. Cultural models, pragmatic scales, and rhetorical strategies are all tools we use to construct and reconstruct a cultural understanding of the world we both inhabit and create.

CHAPTER TEN

The Space Structuring Model

> With false hope of a firm foundation gone, with the world displaced by words that are but versions, with substance dissolved into function, and with the given acknowledged as taken, we face the questions of how worlds are made, tested, and known.
>
> – Nelson Goodman, *Ways of Worldmaking*

> This approach requires . . . a fundamental change in perspective, such that the contingence of action on a complex world of objects, artifacts, and other actors, located in space and time, is no longer treated as an extraneous problem with which the individual actor must contend, but rather is seen as the essential resource that makes knowledge possible and gives action its sense.
>
> – Lucy Suchman, *Plans and Situated Actions*

We have seen that on-line meaning construction is a nontrivial process in which a speaker assembles utterance meaning in response to linguistic clues. Because meaning is considerably underdetermined by the overt structure of language, it often requires the creative application of background knowledge. The importance of background knowledge is especially obvious when the objects of analysis are stretches of connected discourse. Schank & Abelson (1977), for example, point to the somewhat surprising difficulty of constructing a computational model capable of understanding simple stories like this one:

> Seana went to a restaurant.
> She ordered chicken.
> She left a large tip.

Although one might conceivably build a model that could construct meanings for each individual sentence, Schank & Abelson argued that

such a model would fail to compute a number of things human readers would naturally assume to have transpired.

Seana went to a restaurant.
(She sat down.)
She ordered chicken.
(She ate the chicken.)
(She paid the check.)
She left a large tip.
(She left the restaurant.)

The problem by now is a familiar one: The interpretation of the combined utterances is much richer than one that might derive from formal composition of its components. Moreover, it arises because the literal content of the sentences underspecifies Seana's adventures in the restaurant. Because understanding what happened to Seana in this story seems to require a more general understanding of what goes on in restaurants, Schank & Abelson reasoned that representations of stereotyped situations might be useful for natural language processing systems. Consequently, in the course of building a system that could understand stories such as Seana's, Schank & Abelson (1977) utilized *scripts*, data structures they proposed as being analogous to Minsky's (1975) frames. Scripts represent stereotyped sequences of common events and provide the information the system needs to fill in gaps in the story.

But how often do we encounter stories like Seana's, where she went to a restaurant, ordered chicken, and left a large tip? That's about as interesting as a story about a guy watching paint dry. In fact, the latter would probably be *more* interesting since one could at least hope for the story to provide hints about the protagonist's motives and explain his unusually subdued behavior. Consider instead the following story about Seana and Rodney, a young couple who meet for dinner shortly after the Christmas holidays.

10.1 MY DINNER WITH RODNEY

As Seana enters, she sees Rodney seated at the sushi bar and heads across the restaurant to join him. He rises to greet her with a kiss, and helps her out of her coat.

"It's so good to see you!" she says, sitting down next to him. They order drinks and chat about their respective Christmases. Seana gives an enthusiastic review of her brief stay in the wintery Midwest. "Of course it's nice

to be back in California, but ice skating on the lake, with snow everywhere – it was amazing."

"Sounds fantastic! Actually, I had a pretty good time, too," says Rodney. "Me and my brother dusted off our surfboards and headed down to the beach."

"Excellent," says Seana. Turning her attention to the sushi bar, she says, "Hey, Hiro, what's that you're making?"

The story so far is pretty simple, but already its complexity exceeds the capacity of most comprehension models. As in the previous example, the story here is about going to a restaurant, and consequently recruits information from general knowledge of restaurants, as well as more specific information about Japanese restaurants. But besides the dinner, this is a story about two people on a date, a story about two people discussing their Christmas vacations, and, to a lesser extent, a story about what happened on those Christmas vacations. Though the different storylines aren't all of equal importance, there's a sense in which they really are all part of the story.

In the space structuring model, meaning construction involves partitioning the information in the message-level representation into different mental spaces. One of the things that creates the need for partitioning is the fact that different subsets of knowledge are necessary for understanding different storylines. For instance, while knowledge about surfing is important for understanding what Rodney and his brother did when they went to the beach, it's not important for understanding why Rodney kisses Seana as she enters the restaurant. Moreover, while some of the same characters crop up in the different stories, they do not have equal significance in all of them. Hiro, for example, is a central character in the dinner story, but a peripheral one in the story about the couple's date.

To understand this story, we need one space for the dinner, which recruits structure from the restaurant script, and one space for the date, which recruits structure from knowledge about dating practices. Because it is primarily a story about a date, the dinner space is embedded in the date space. As information comes in, it is added to spaces in which it applies. Comprehension is aimed at establishing links between counterpart elements in active spaces, and mapping structure from space to space. For example, the kiss is part of the greeting in the dinner space, and is a show of intimacy in the date space. The frozen lake in Seana's vacation story is represented by an element in her vacation space (a space that is itself embedded as a conversation topic in the

dinner space), but has no counterpart in the outer spaces. Seana's exclamation, "Excellent!" is a reaction to structure set up in the vacation space, but gets its relevance from the date space, where it signifies her appreciation of Rodney's story.

However, in turning her attention from Rodney to Hiro, Seana shifts the focus from the date space to the dinner space by asking Hiro what he is making.

"Spicy tuna roll," answers Hiro.
Seana looks at Rodney and says, "Spicy tuna roll?"
Rodney looks at Hiro and says, "Spicy tuna roll."

As we have seen, language motivates the construction of utterance meaning but does not definitively determine it. In the preceding excerpt, the definition of "spicy tuna roll" is necessary for a complete understanding of the meaning of these utterances, but it is far from sufficient. Here "spicy tuna roll" means, "I'm making a spicy tuna roll," "Would you like a spicy tuna roll?" and, "Please make me a spicy tuna roll." Moreover, this is only what the utterances convey in the *dinner* space. In the date space, "spicy tuna roll" means, "I would like a spicy tuna roll and I want to know if you share my desire," as well as, "I share your desire (for a spicy tuna roll) and I want to make you happy by fulfilling it."

Because understandings evoked by particular utterances are far richer and often quite different from the understanding of expressions in a null context, conventional word meanings have limited utility for the online construction of utterance meaning. On the space structuring model, meaning construction involves the integration of linguistic, contextual, and background knowledge to create the cognitive models speakers use to participate in the ongoing activity. Given the contextuality of utterances, perhaps a context-sensitive retrieval mechanism would be useful for explaining why the same word can mean different things in different contexts, as well as why the same utterance can mean different things in different social interactions. These sorts of considerations collapse locutionary and illocutionary force, and suggest that the same sorts of cognitive processes subserve them both.

Going back to their conversation about the holidays, Seana asks Rodney how his older brother's doing.

"Oh, he's doing great as usual. This fall he started a fellowship in cardiothoracic surgery, and my parents think he can do no wrong. Me, on the other hand..."

"Come on, Rodney," she reassures him, "I'm sure your parents love you both!"

"No," he says, "My parents always hated me. When I was a kid, my bath toys were a toaster and a radio."

Seana bursts into laughter.

In the preceding excerpt, Rodney shifts subtly from talking about his parents' feelings toward him and his brother to a joke about childhood bath toys. Because it is a joke, Rodney's utterance is not understood as conveying information about events in his own childhood. The "I" in this utterance refers primarily to Rodney's counterpart in the joke space, and there's no implication that Rodney's real parents tried to kill him as a child by letting him play with electrical appliances in the tub. However, Rodney really *is* telling a joke, and the fact that he's telling a joke has ramifications in the date space as an attempt to amuse Seana. Since her laughter is an inappropriate response to the events that occur within the imaginary world of the joke, it triggers a shift in focus to the date space, where it signals Seana's assent to Rodney's attempt to amuse her.

The demands of meaning construction at this level involve taking information evoked by each of the characters' actions and utterances and assimilating it to relevant structure in the dinner and date spaces. Though the content of those utterances is clearly related to their significance in the various spaces, the precise relationship between the two is quite complex. Further, just as the speakers' utterances are related to one another via integration, with structure being built at higher levels, people attempt to integrate sentences within a single speaker's utterance. Consequently, when Rodney follows, "My parents always hated me," with, "When I was a kid, my bath toys were a toaster and a radio," the listener attempts to construct a model consistent with the content of both sentences.

Moreover, the same sorts of coherence-establishing processes that seem to be needed to establish the relationship between sentences in a text are also needed to relate the elements in a single sentence. For instance, understanding the punchline in Rodney's joke requires the integration of information about baths and bath toys with information about electrical appliances. While cognitive scientists have long realized the importance of background knowledge for understanding such leaps, most models have not adequately addressed the *way* in which knowledge is brought to bear. In most models, knowledge is instantiated by

binding contextual elements to preexisting data structures, in a way that grossly underestimates the extent to which filling a slot – in this case, binding "toaster" to "bath toy" – involves imaginative work.

For one thing, it's unlikely that we have a preexisting representation either for electronic bath toys or for ways you know your parents hate you. So, rather than instantiating a high-level schema, the speaker must construct one by integrating information from several different domains. Most importantly, models of how children generally play with bath toys are blended with the particular elements proposed in the joke: a childhood Rodney counterpart, one or more of Rodney's parents, a bathtub, a toaster, and a radio. Elaboration via the cultural model in which human actions are caused by the intention to produce their consequences can yield an integrated model of the cause of the bathtub scenario, as well as its intended consequences for Rodney. On any account, the comprehension of Rodney's joke requires activation of the knowledge that electrical appliances submerged in water can cause electrocution. Conceptual integration theory suggests that this knowledge operates by animating an imaginary scenario in which one of Rodney's parents hands him an electrical appliance to play with in the tub.

> *"Actually," says Rodney, "The best time I had this weekend was going for sushi with my brother and his girlfriend."*
>
> *"Oh?" says Seana.*
>
> *"Sushi virgin," says Rodney with a knowing smile.*
>
> *"So, did you get her to try to sea urchin?" asks Seana, laughing.*
>
> *"Oh, please, we couldn't even get her to try the cucumber roll! We started out slow, ordering cucumber roll and shrimp – you know, thinking she might ease into it," he chuckles. "But then when she wouldn't try the tame stuff, we took turns ordering everything from octopus, to monk fish liver, to Ikura and quail egg yolk."*

With the utterance, "we couldn't even get her to try the cucumber roll," Rodney employs scalar implicature to accentuate his brother's girlfriend's culinary temerity. In doing so, he makes use of a scale – shared by himself, his brother, and Seana – of exotic sushi. The scale begins with cucumber roll and shrimp and proceeds out to octopus, monk fish liver, and sea urchin. So, by denying the sushi virgin's willingness to try an item close to the origin (cucumber roll), Rodney implicates, and indeed emphasizes, her fear of all things sushi. In contrast to the semanticist's isolated expression meanings, utterances are always interpreted

against a rich background of shared frames and scales that speakers can strategically exploit.

In fact, the sushi scale structures much of this excerpt. Because sea urchin occurs at an extreme point on the scale, it makes Seana's question a potentially informative one. Though a straightforward negative answer would not be particularly informative, a positive answer would suggest Rodney's brother's girlfriend tried everything they offered her. Scalar implicatures are set up by single sentences, as well as by combinations of sentences in a discourse. Moreover, the sushi scale can even be seen to structure the way Rodney and his brother order – starting with items close to the origin and working their way out.

> *Giggling, Seana says, "Rodney, you're making me hungry!"*
> *Rodney turns to the chef and says, "Hirosan, how about one of your famous Shogun TNT rolls?"*

The semi-opaque character of phrases such as "Shogun TNT roll" (and, to a lesser extent, "sushi virgin") presents a serious challenge for traditional approaches to meaning construction. On such approaches, the lexicon supplies atomic components of meaning that are combined via syntactic instructions. However, in this case, the meanings in the lexicon do not combine in any straightforward way. A *Shogun* was the title given to military governors in Japan prior to the revolution of 1867, *TNT* is a flammable toxic compound, and a *roll* can be any number of things – though in this context it is a cylindrical food item composed of vinegared rice, fish, and vegetables. But even if the speaker specifies the sushi-related context, the listener unfamiliar with the meaning of "Shogun TNT roll" is highly unlikely to deduce that it is a piece of raw tuna, surrounded by rice, wrapped in Nori (seaweed), and topped with a very hot, red pepper paste.

Since we could not produce the description of the referent on the basis of the linguistic information alone, it is clear that the relationship between "Shogun TNT roll" and the individual meanings of "Shogun," "TNT," and "roll" is not compositional. However, the relationship isn't random either. Presumably, a competent speaker could pick a Shogun TNT roll out of a sushi lineup without too much difficulty. In Langacker's (1987) terminology, this noun phrase is *analyzable* in that the meaning of the whole can be related to the meanings of its component parts. In particular, its analyzability lies in its appeal to a number of mappings that organize the integration, especially between information evoked by "TNT" and the contextually relevant sense of "roll."

Besides referring to the chemical trinitrotoluene, "TNT" is often used to refer to sticks of dynamite. Though perceptually the roll differs substantially in size and color from a stick of dynamite, their roughly cylindrical shape is apparently enough to establish a mapping between the two. Dynamite is used for its explosive properties, and though the Shogun TNT roll does not literally explode, when the red peppers activate the eater's pain receptors, she will experience a sudden sensation that can be mapped onto the explosion of dynamite. This yields a further mapping between the dynamite's wick and the pepper paste, as each can be construed as potentially triggering the explosion.

"Shogun" refers metonymically to Japan, and provides a context – albeit a historically implausible one – for interpreting "TNT." While it's unlikely that any Shogun ever used TNT (or dynamite), it's conceivable that a warlord might want to make use of an explosive device. The two-word combination "Shogun TNT" triggers the sort of blending often used in creative category extensions, like the French Watergate or the poor man's opera. In Turner's (1990) example, "Sex is the poor man's opera," the listener is invited to map between two spaces, one with a poor man and sex, the other with a rich man and opera, and to infer analogically shared structure. Similarly, here the listener is invited to map between a space with a Shogun and a sushi roll, and a modern warrior with TNT.

In fact, in context, a speaker might easily understand this phrase without knowing the exact definitions of either *Shogun* or *TNT*. In the space structuring model, conceptual structure recruited for meaning construction consists of partial cognitive models that, though built from frames in long-term memory, are not identical to them (see Barsalou, 1993; Glenberg et al., 1994, for similar suggestions). Because cognitive models are constructed for particular contexts, and for particular goal-directed activities, the role of background knowledge is to constrain the elaboration of the idealized models that speakers construct. The nature of meaning construction, then, is such that it allows for the formal composition of meanings, as well as for the more creative concept combination often prompted by modified noun phrases.

In their study of children's novel word acquisition, Smith & Jones (1993) argue that observed productivity does not primarily reflect structures in long-term memory, but rather "real-time, real-task processes that flexibly adjust attention to find the most likely referent of a specific utterance of some unknown word" (Smith & Jones, 1993: 184). We have seen that the establishment of various sorts of pragmatic

mappings provides a key component of the meaning construction in these cases. Moreover, while background and contextual knowledge are clearly recruited by these real-time processes, their role appears to be far more extensive than supposed by compositional approaches to meaning. Background knowledge is not called upon primarily to fill an informational gap, and nor is contextual knowledge being recruited to disambiguate between preexistent interpretations. Rather, speakers selectively recruit and integrate contextual information with selective aspects of background knowledge from three different domains. It is this integration that enables our protagonists to comprehend the connection between the noun phrase "Shogun TNT roll" and Hiro's spicy concoction.

"So Hiro," says Seana, "How's life treating you."

As he skillfully rolls the rice into the Nori, Hiro shakes his head and says, "Oh, not so good, guys. Last week, I came home early from work and saw a guy jogging naked. So I go, 'Why are you doing that?' And he says, 'Because you came home early.'"

Almost in unison, Seana and Rodney groan in appreciation of Hiro's joke.

In this joke, Hiro begins by setting the scene, locating it temporally with respect to the present, and contextualizes the scene as something that happened when he came home early from work. In this minimal context, a guy jogging naked is something out of the ordinary and presents itself as something in need of an explanation. The line, "Because you came home early," prompts the listeners to construct an explanation, as well as triggering frame-shifting. As we have seen, frame-shifting is the dramatic reanalysis that occurs when the frame being used to structure interpretation of a given strip of activity proves to be inadequate.

Usually prompted by violation of slot-filling constraints, frame-shifting proceeds by mapping information from the extant message-level representation into a new space structured by a different frame. Because the shifted interpretation requires blending models from the initial interpretation with the newly evoked frame, it can have dramatically different implications and assume new evaluative significance. In Hiro's joke, for instance, the narrator's early return from work is initially interpreted as part of the background for the story. However, in the shifted interpretation, it assumes a causal role. Moreover, while Hiro's promotion to the cause of the naked jogger's behavior is explicitly cued by the language of the joke, the reader must construct for herself the

precise understanding of *how* his early return from work could inspire naked jogging.

Although we don't have very many scripts with slots for naked joggers, at least two scenarios involve naked people running, and might conceivably be recruited to make sense of this utterance. One involves insanity, a common recourse for explaining events that don't fit into general patterns. Another involves an adulterer fleeing from his lover's husband. Integrating the latter model with the extant model of Hiro coming home from work early suggests that Hiro has a wife, that his wife has a lover, and that he is a cuckold. Indeed part of the humor of this joke is that the shifted interpretation evokes a counterfactual space in which Hiro comes home from work on time and the naked jogger gets his exercise in another way.

Rodney says, "Yeah, I know what you mean. Last week I went to the doctor because I'd swallowed a bottle of sleeping pills. My doctor told me to have a few drinks and get some rest."

In Rodney's reply to Hiro's joke, we see again how background knowledge helps both to structure expectations and to revise them in light of incoming information. Moreover, the relationship that comprehension bears to contextual and background knowledge goes beyond disambiguation and the provision of missing information. The contextual knowledge that Hiro has just told a self-deprecating joke, for instance, helps us to understand the relationship between Rodney's narrative and the overall discourse scenario, but does not *disambiguate* between two preestablished interpretations. In fact, construal of the story as reflecting badly on the joke's protagonist requires the creative integration of information from a number of input domains.

For example, we see the influence of background knowledge about a general medical problem scenario, the particular medical problem of an overdose, and a model about the combined effect of sleeping pills and alcohol on the human. The general medical problem scenario appeals to a scriptlike model in which the narrator sees a doctor about the medical problem, and the doctor advises him to take medicine and get some rest. But while the general model holds for many ordinary doctor visits, it doesn't work very well with the particular medical problem in question: an overdose. A normative model for treatment of an overdose is to pump the patient's stomach.

The event sequence specified in the joke, then, can be seen as the composition of two sequences that, relative to the joke, are counterfactual:

the beginning of the overdose sequence, and the end of the medical problem sequence. But the joke does not concern the incompatibility of the general medical problem model with the particular demands of an overdose victim. Rather, the humor here turns on the way in which "drinks" fills the medication slot projected from the general medical problem model. This unorthodox role assignment is apparently enough to elicit a cognitive model of the combined effects of sleeping pills and alcohol.

Neither the combined effects model nor the mappings between this general model and the particulars of Rodney's scenario is cued grammatically. Nonetheless, integration and elaboration proceed naturally to yield the inference that the doctor's advice could kill the protagonist. Because of the action model, we assume the doctor wasn't just incompetent, but acted from the intention to kill. Moreover, the expectation that this might be self-deprecating humor suggests the nature of the motive for the doctor's intention to kill. Once again, the emergent structure produced by conceptual blending does not primarily result from the instantiation of a higher-level schema to fill gaps in the narrative sequence. As in the previous bath toys joke, background knowledge animates the scenario set up in the blend.

> *As they all laugh, Seana reaches under the bar and touches Rodney's knee. "How do you feel about unagi, Rod?" Unagi, a freshwater eel that Rodney jokingly claims is an aphrodesiac, is usually the last thing the couple orders. Rodney takes the question as a private signal, saying, "How about an order of unagi, Hirosan? And the check."*

The ease with which we construct the couple's private meaning for "unagi" provides us with an emphatic demonstration that meaning is not specified determinately by the grammar. Moreover, it highlights an important limitation of the lexicon, a limitation that is more profound than the lack of a metaphoric or metonymic meaning. Because the utterance meaning appeals crucially to the couple's shared experiences and understandings, information associated with the lexical entry for "unagi" is almost useless.

In contrast, the space structuring approach attempts to account for the variability of utterance meaning by appealing to flexible initial activation of structure supplemented with the creative processes of cross-space mapping and conceptual integration. By invoking a constructivist paradigm for *all* cases of meaning construction, the model is aimed at providing a unified account of (i) straightforward, compositional instances of meaning construction, (ii) the contextual variability

of utterance meaning, (iii) the way that speakers adjust their expectations based on incoming information, and even (iv) nonconventional meanings such as the one in the preceding example.

10.2 TURNING THE CAT ON ITS HEAD

Traditional approaches to meaning in cognitive science have imported the assumption from linguistics that natural language is a formal system, or at least an approximation of one. Perhaps due to the interdisciplinary nature of the field, this limited view of language has spread to other areas, and even become institutionalized in the practice of cognitive science. For example, the idea that the meaning of a sentence is built from the meanings of component parts suggests that the language research program can and should be divided accordingly. An unfortunate consequence of this assumption is that much of modern psycholinguistics research is directed at processes of word recognition and sentence processing in a way that assumes it can be addressed independently of the formulation of the message-level representation.

Moreover, the assumption that expression meaning is the fundamental component of utterance meaning underlies cognitive scientists' tendency to build and test models that explain simple language use in minimal contexts. Unfortunately, by restricting our attention to cats on mats in typical settings (complete with their gravitational fields and their standard orientational frames), we are prevented from seeing how contextual knowledge and background assumptions interact, and their constitutive role in on-line meaning construction. For the speaker, it's quite a nice thing that we can rely on knowledge of typical situations for meaning construction. But for the researcher, it's often more valuable to turn the cat upside down and look at the atypical, and even the exotic. It is these exotic examples that allow us to assess more adequately the nature of the information and the processes speakers recruit in on-line meaning construction.

Our investigation of the semantic leaps in jokes, arguments, counterfactuals, and analogies suggests on-line meaning construction requires a multitiered integration system that involves hierarchical slot-filler structures all the way down. At bottom, sentential integration is a process in which speakers integrate abstract grammatical constructions with more specific frames evoked by lexical items. Grammatical information such as space builders and verbal morphology can cue the construction of new spaces. Similarly, other sorts of grammatical information

such as clause and sentence boundaries can cue the spaces' delimitation. However, both the opening and closing of spaces can go unmarked by grammatical features, relying on language users to detect the changing background assumptions necessary for continued interpretation.

Data that suggest that meaning is underspecified by the grammar indicate that grammatical information is not always necessary nor sufficient for the computation of on-line meanings. Such data reveal the constructive nature of comprehension and point to the crucial role played by noncompositional processes. Rather than a deterministic process of composition in which the satisfaction of truth conditions for the resultant sentence can be seen to be truth-functionally related to the meaning of the parts, the space structuring model appeals to the constructional productivity of the imaginative processes of conceptual integration.

The space structuring model does not eliminate the need for combinatorial mechanisms of meaning construction, nor does it preclude the possibility of dedicated structures or specialized processes for parsing. However, by casting the noncompositional processes of blending in a leading role, the model does undermine the significance of the parser. Structural regularities in the language still help us to figure out who did what to whom, but our understanding of how these regularities do so is quite different from the original vision of, for example, Chomsky (1965) or Montague (1974), in which an autonomous syntax composes semantic representations that connect with the language of thought. Instead, structural regularities are construed as having semantic content of their own, albeit with very abstract content.

In many ways, structure in the blended space can be seen as an overly literal interpretation of the linguistic input. However, while the assumption in many traditional accounts of meaning is that such interpretations are rejected as irrelevant to the resultant interpretation, in the space structuring model, the implausible representations constructed in the blend are related to resultant message-level representations in principled ways. For instance, speakers often recruit metaphoric, metonymic, and other sorts of cross-space mappings that rely on the induction of shared relational structure.

A recurring theme has been that certain processes appear to operate at a number of different levels: partitioning, mapping, contextual variation of meaning, conceptual integration, scalar implicature, frame-shifting, and the relevance of salient counterfactuals. In Chapters 3 and 4, we reviewed empirical data from psychology and cognitive neuroscience that suggest higher-level discourse factors operate in the processing of

single words. Moreover, we have seen analytically how the demands of sentential integration are not qualitatively different from those of text processing.

Of course, the finding that discourse-level considerations affect the processing of individual words is counterintuitive only from the traditional building blocks approach to meaning construction. That is, once we abandon the notion that contextual and background knowledge are brought to bear *after* the assembly of a context-invariant meaning, the finding that the same factors operate at the lexical, sentential, and intersentential levels should come as no surprise. If language is designed to prompt the construction of cognitive models, the cuing of projections, and the elaboration of blends, we should actually predict that words, sentences, and groups of sentences can prompt the same sorts of operations.

In fact, as cognitive scientists, we need to make our own leap, akin to the shift Suchman recommends in the quote at the chapter's outset. This leap involves abandoning the old assumption that the systematicity and productivity of human cognition are the necessary result of a system that formally composes static symbols. Moreover, it involves embracing the situated character of on-line meaning, the constructive nature of comprehension, and the constitutive role of context. Because language use, in particular, is firmly rooted in human experience and social interaction, we need to construe meaning construction as a set of routines for assembling cognitive models that enable interpretation, action, and interaction. Besides acknowledging the crucial role of the physical and social world within which we function, the leap toward situativity is congruent with the rising consciousness in cognitive neuroscience of the importance of the motor system (see, e.g., Rizzolatti & Craighero, 1998), and the growing realization that attention, perception, and memory are all intimately connected with action (see, e.g., Arbib & Rizzolatti, 1996; Ballard, Harhoe, Pook, & Rao, 1997; Milner & Goodale, 1995).

Neuroscience gives us a picture of information processing as involving partitioning of sensory information into parallel streams, each computing different sorts of information, and each with its own hierarchical structure (Van Essen, Anderson, & Felleman, 1992; Ungerleider & Mishkin, 1982). The massively interconnected systems allow for information to be continuously mapped and remapped between intertwined processing streams. Similarly, the space structuring model, though motivated by very different issues and sorts of data, portrays meaning construction in an analogous way: the partitioning of information into

parallel streams, extensive mapping, and the integration of disparate information needed for an adequate message-level comprehension. While the establishment of abstract mappings in mental space theory is not directly comparable to mapping in the visual system, perhaps computationally similar mechanisms of information regulation underlie flexibility evident in both meaning construction and visual processing.

In fact, recent research in cognitive psychology (Barsalou, 1999; Glenberg et al., 1994; Mandler, 1993) points to the import of what Barsalou (1999) calls *perceptual symbols*. Perceptual symbols are mental representations that are neither perceptual – that is, strictly dependent on sensory input systems – nor symbolic – that is, completely amodal. As outlined in Barsalou (1999), schematic representations of perceptual experience are stored around a common frame that promotes schematized simulations. Importantly, such simulations need not be accompanied by the experience of visual imagery, and are not to be construed as mental "pictures." Indeed, perceptual symbols recruit neural machinery activated in perceptual experience from all modalities – auditory, olfactory, somatosensory, and kinesthetic, as well as visual. As abstracted perceptual experience, perceptual symbols develop to support categorization, inference, and interaction with the world around us. Frames built from perceptual symbols present themselves as representations that can sustain the creative blending mechanisms of composition, completion, and elaboration while maintaining the representational advantages of hierarchically organized slot-filler structures.

10.2.1 Connections

We began by posing two problems with a frame-based approach to meaning construction. The first problem contrasted the simplicity of frames with the complexity of human behavior, while the second contrasted the static nature of these data structures with the flexibility of the behavior they have been hypothesized to support. In answer to the first, the issue of how simple representations can sustain the complex cognitive operations that people seem to perform, part of the explanation lies in the fact the people use language to set up construals rather than to evoke objectively valid models of the world. Although the models used in meaning construction are often schematic and occasionally even implausible, they are supported by our ability to integrate models dynamically with contextual information, to exploit culturally defined scales, to shift frames, and to blend different models together.

In answer to the second question, of how static frames can sustain the flexible language use evidenced by speakers, we effectively state that they cannot. Though data support the contention that human knowledge representation is hierarchically structured, that slot-filler structure is fundamental, and that people do assume default values when none are provided, it seems that knowledge representation is more dynamic than originally proposed. Rather than just retrieving and instantiating frames, speakers are continuously and creatively building and blending cognitive models to yield new concepts, construals, and activities.

In underlining the importance of imagination in information integration, and understanding the dynamic, flexible character of utterance meaning as arising from the demands of situated cognition, the space structuring model affords a promising approach to the fascinating domain of on-line meaning construction. In its present form, the main contribution of the approach is a way of thinking about some old problems. However, it can and should be translated into empirical predictions and computational models, which in turn can be tested with the methods of cognitive anthropology, linguistics, psychology, and neuroscience. Given the tools at our disposal, we should soon be well on our way to the formulation of a precise understanding of the emergence of on-line meanings. Sometimes a journey of a thousand steps begins with a single leap.

Bibliography

[1] Allen, J. (1987). *Natural Language Processing*. Philadelphia, PA: Benjamin/Cummings Publication.

[2] Anderson, A., Garrod, S., & Sanford, A. (1983). The accessibility of pronominal antecedents as a function of episode shifts in narrative text. *Quarterly Journal of Experimental Psychology* 35A: 427–440.

[3] Anderson, R. C. & Ortony, A. (1975). On putting apples into bottles – a problem of polysemy. *Cognitive Psychology* 7: 167–180.

[4] Arbib, M. & Rizzolatti, G. (1996). Neural expectations: A possible evolutionary path from manual skills to language. *Communication & Cognition* 29: 393–424.

[5] Aristotle. (1952). *Aristotle: Great Books of the Western World*. Vol. 8. Chicago, IL: Encyclopedia Britannica.

[6] Armstrong, S., Gleitman, L., & Gleitman, H. (1983). What some concepts might not be. *Cognition* 13(3): 263–308.

[7] Attardo, S., Attardo, D. H., Baltes, P., & Petray, M. J. (1994). The linear organization of jokes: Analysis of two thousand texts. *Humor: International Journal of Humor Research* 7(1): 27–54.

[8] Ballard, D., Hayhoe, M., Pook, P., & Rao, R. (1997). Deictic codes for the embodiment of cognition. *Behavioral & Brain Sciences* 20: 723–767.

[9] Barclay, J., Bransford, J. D., Franks, J. J., McCarrell, N. S., & Nitsch, K. E. (1974). Comprehension and semantic flexibility. *Journal of Verbal Learning and Verbal Behavior* 13(4): 471–481.

[10] Barsalou, L. W. (1987). The instability of graded structure: Implications for the nature of concepts. In U. Neisser (ed.), *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*. Cambridge, UK: Cambridge University Press, 101–140.

[11] Barsalou, L. W. (1989). Intra-concept similarity and its implications for interconcept similarity. In Stella Vosniadou and Andrew Ortony (eds.), *Similarity and Analogical Reasoning*. New York: Cambridge University Press, 76–121.

[12] Barsalou, L. W. (1992). Frames, concepts, and conceptual fields. In Adrienne Lehrer and Eva Feder Kittay (eds.), *Frames, Fields, and Contrasts:*

New Essays in Lexical and Semantic Organization. Hillsdale, NJ: Lawrence Erlbaum Associates.

[13] Barsalou, L. W. (1993a). Challenging assumptions about concepts. *Cognitive Development* 8(2): 169–180.

[14] Barsalou, L. W. (1993b). Flexibility, structure, and linguistic vaguary in concepts: Manifestations of a compositional system of perceptual symbols. In A. F. Collins, S. E. Gathercole, M. A. Conway, P. E. Morris (eds.), *Theories of Memory*. Hilldale, NJ: Lawrence Erlbaum Associates, 28–102.

[15] Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22: 577–609.

[16] Barwise, J. & Etchemendy, J. (1989). Model-theoretic semantics. In M. I. Posner (ed.), *Foundations of Cognitive Science*. Cambridge, MA: MIT Press, 207–244.

[17] Bateson, G. (1976;1972;1954). *Steps to an Ecology of Mind*. New York: Ballantine Books.

[18] Beeman, M. (1993). Semantic processing in the right hemisphere may contribute to drawing inferences from discourse. *Brain & Language* 44: 80–120.

[19] Beeman, M., Friedman, R., Grafman, J., Perez, E., Diamond, S., & Lindsay, M. (1994). Summation priming and coarse semantic coding in the right hemisphere. *Journal of Cognitive Neuroscience* 6: 26–45.

[20] Bihrle, A., Brownell, H., & Gardner, H. (1986). Comprehension of humorous and nonhumorous materials by left- and right-brain damaged patients. *Brain and Cognition* 5: 399–411.

[21] Black, M. (1962). *Models and Metaphors*. Ithaca, NY: Cornell University Press.

[22] Black, M. (1993). More about metaphor. In Andrew Ortony (ed.), *Metaphor and Thought*, Second Edition. Cambridge, UK: Cambridge University Press, 19–41.

[23] Bloom, P. A. & Fischler, I. (1980). Completion norms for 329 sentence contexts. *Memory & Cognition* 8: 631–642.

[24] Boden, M. A. (1988). *Computer Models of Mind*. Cambridge, UK: Cambridge University Press.

[25] Bower, G. H., Black, J. B., & Turner, T. J. (1979). Scripts in memory for text. *Cognitive Psychology* 11: 177–220.

[26] Brachman, R. J. (1985). "I lied about the trees" – or, defaults and definitions in knowledge representation. *The AI Magazine* 6: 80–93.

[27] Brownell, H., Michel, D., Powelson, J., & Gardner, H. (1983). Surprise but not coherence: Sensitivity to verbal human in right-hemisphere patients. *Brain and Language* 18: 20–27.

[28] Brownell, H., Potter, H., Bihrle, A., & Gardner, H. (1986). Inference deficits in right brain-damaged patients. *Brain and Language* 27: 310–321.

[29] Buckner, R. (1996). Beyond HERA: Contributions of specific pre-frontal brain areas to long-term memory retrieval. *Psychonomic Bulletin and Review* 3: 149–158.

[30] Caplan, D. (1972). Clause boundaries and recognition latencies for words in sentences. *Perception and Psychophysics* 12(1-B): 73–76.

[31] Carlon, K., Davidse, K., & Rudzka-Ostyn, B. (eds.). (1994). *Perspectives on English: Essays in Honour of Emma Vorlat*. Leuven, Belgium: Peeters.
[32] Charniak, E. (1983). Passing markers: A theory of contextual influence in language comprehension. *Cognitive Science* 7: 171–190.
[33] Chomsky, N. (1965). *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
[34] Churchland, P. (1989). *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*. Cambridge, MA: MIT Press.
[35] Cicourel, A. V. (1988). Elicitation as a problem of discourse. In T. v. Dijk (ed.), *Handbook of Discourse Analysis*. Vol. 4: *Discourse Analysis in Society*. Amsterdam: North Holland Publications, 193–202.
[36] Clark, H. H. (1979). Responding to indirect speech acts. *Cognitive Psychology* 11(4): 430–477.
[37] Clark, H. H. (1983). Making sense of nonce sense. In Giovanni B. Flores D'Arcais, R. J. Jarvella (eds.), *The Process of Language Understanding*. Chichester, UK: John Wiley & Sons, 297–332.
[38] Coulson, S. (1992). Is incest best? The role of pragmatic scales and cultural models in abortion rhetoric. *CRL Newsletter* 7(2). San Diego, CA: Center for Research in Language.
[39] Coulson, S. (1997). Semantic leaps: The role of frame-shifting and conceptual blending in meaning construction. University of California, San Diego, Ph.D. Dissertation.
[40] Coulson, S. & Fauconnier, G. (1999). Fake guns and stone lions: Conceptual blending and privative adjectives. In B. Fox, D. Jurafsky, L. Michaelis, (eds.), *Cognition and Function in Language*. Stanford, CA: CSLI.
[41] Coulson, S. & Flor, N. V. (1994). Rational choice and framing devices: Argumentation and computer programmers. *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.
[42] Coulson, S., King, J., & Kutas, M. (1998). Expect the unexpected: Grammaticality and probability effects on the ERP. *Language and Cognitive Processes* 13: 21–58.
[43] Coulson, S. & Kutas, M. (1998). Frame-shifting and sentential integration. University of California San Diego Cognitive Science Department Technical Report.
[44] Dale, A. (1994). Source localization and discriminant analysis of event-related potentials: Linear approaches. University of California San Diego, Ph. D. Dissertation.
[45] D'Andrade, R. (1987). A folk model of the mind. In D. Holland, N. Quinn (eds.), *Cultural Models in Language and Thought*. Chicago, IL: University of Chicago Press, 112–150.
[46] Dinsmore, J. (1991). *Partitioned Representations: A Study in Mental Representation, Language Understanding, and Linguistic Structure*. Dordrecht, Boston: Kluwer Academic Publishers.
[47] Donchin, E. (1981). Presidential Address, 1980. Surprise!... Surprise? *Psychophysiology* 18(5): 493–513.

[48] Downing, P. A. (1977). On the creation and use of English compounds. *Language* 53: 810–842.

[49] Dowty, D. R., Wall, R. E., & Peters, S. (1981). *Introduction to Montague Semantics*. Dordrecht, Netherlands: D. Reidel Publishing Company.

[50] Dyer, M. G. (1990). Symbolic neuroengineering for language processing: A multi-level research approach. In J. Barnden, J. Pollock (eds.), *Advances in Connectionist and Neural Computation Theory*, 32–86. Norwood, NJ: Ablex.

[51] Fauconnier, G. (1988). Quantification, roles, and domains. In U. Eco (ed.), *Meaning and Mental Representations*. Bloomington, IN: Indiana University Press, 61–80.

[52] Fauconnier, G. (1994;1985). *Mental Spaces: Aspects of Meaning Construction in Natural Language*. Cambridge, UK: Cambridge University Press.

[53] Fauconnier, G. (1997). *Mappings in Thought and Language*. Cambridge, UK: Cambridge University Press.

[54] Fauconnier, G. & Sweetser, E. (1996). *Spaces, Worlds, and Grammar*. Chicago, IL: The University of Chicago Press.

[55] Fauconnier, G. & Turner, M. (1994). Conceptual projection and middle spaces. Technical Report no. 9401, Department of Cognitive Science, University of California, San Diego.

[56] Fauconnier, G. & Turner, M. (1996). Blending as a central process of grammar. In A. E. Goldberg (ed.), *Conceptual Structure, Discourse, and Language*. Stanford, CA: CSLI Publications, 113–131.

[57] Fauconnier, G. & Turner, M. (1998). Conceptual integration networks. *Cognitive Science* 22: 133–187.

[58] Feldman, J. A. (1989). Neural representation of conceptual knowledge. In L. Nadel, L. A. Cooper, P. Cullicover, R. M. Harnish (eds.), *Neural Connections, Mental Computation*. Cambridge, MA: MIT Press, 68–103.

[59] Fillmore, C. J. (1968). The case for case. In E. Bach, R. T. Harm (eds.), *Universals of Linguistic Theory*, New York: Holt, Rinehart and Winston, 1–90.

[60] Fillmore, C. J. (1977). Topics in lexical semantics. In R. Cole (ed.), *Current Issues in Linguistic Theory*. Bloomington: Indiana University Press, 76–138.

[61] Fillmore, C. J. (1982). Frame semantics. In Linguistic Society of Korea (ed.), *Linguistics in the Morning Calm*. Seoul: Hanshin, 111–137.

[62] Fillmore, C. J., Kay, P., & O'Connor, M. C. (1988). Let alone: Regularity and idiomaticity in grammatical constructions. *Language* 64: 501–539.

[63] Firebaugh, M. (1988). *Artificial Intelligence*. Boston, MA: PWD-Kent Publishing Company.

[64] Fisher, C., Hall, G., Rakowitz, S., & Gleitman, L. (1991). When it is better to receive than to give: Syntactic and conceptual constraints on vocabulary growth. IRCS Report 91-41. Philadelphia, PA: University of Pennsylvania.

[65] Flavell, J. H. (1999). Cognitive development: Children's knowledge about the mind. *Annual Review of Psychology* 50: 21–45.

[66] Flores D'Arcais, G. B. & Schreuder, R. (1983). The process of language understanding: A few issues in contemporary psycholinguistics. In

G. B. Flores D'Arcais, R. J. Jarvella (eds.), *The Process of Language Understanding*. Chichester, UK: John Wiley & Sons, 1–41.

[67] Fodor, J. (1983). *The Modularity of Mind: An Essay on Faculty Psychology*. Cambridge, MA: MIT Press.

[68] Fodor, J. & Pylyshyn, Z. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition* 28: 3–71.

[69] Forbus, K. & Gentner, D. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science* 19(2): 141–205.

[70] Forster, K. I. (1979). Levels of processing and the structure of the language processor. In W. E. Cooper, E. C. T. Walker (eds.), *Sentence Processing: Psycholinguistic Studies Presented to Merrill Garrett*. Hillsdale, NJ: Lawrence Erlbaum Associates, 257–287.

[71] Foss, D. J. & Speer, S. R. (1991). Great expectations: Context effects during sentence processing. In G. Flores D'Arcais, R. J. Jarvella (eds.), *The Process of Language Understanding*. Chichester, UK: John Wiley & Sons, 169–191.

[72] Franks, B. (1995). Sense generation: A "quasi-classical" approach to concepts and concept combination. *Cognitive Science* 19: 441–505.

[73] Frazier, L. (1995). Issues in representation in psycholinguistics. In Joanne Miller, Peter Eimas (eds.), *Speech, Language, and Communication: Handbook of Perception and Cognition*, Second Edition. San Diego, CA: Academic Press, 1–27.

[74] Frege, G. (1970;1892). On sense and reference. In P. Geach, M. Black (eds.), *Translations from the Philosophical Writings of Gottlob Frege*. Oxford, UK: Basil Blackwell.

[75] Freyd, J. (1983). Shareability: The social psychology of epistemology. *Cognitive Science* 7(3): 191–210.

[76] Fruzzetti, A. E., Toland, K., Teller, S. A., & Loftus, E. F. (1992). Memory and eyewitness testimony. In M. M. Gruneberg, P. E. Morris (eds.), *Aspects of Memory*. Vol. 1: *The Practical Aspects*. Second Edition. London: Routledge, 18–50.

[77] Gainotti, G. (1972). Emotional behavior and hemispheric side of the lesion. *Cortex* 8: 41–55.

[78] Gardner, H., Ling, P., Flamm, L., & Silverman, J. (1975). Comprehension and appreciation of humour in brain-damaged patients. *Brain* 98: 399–412.

[79] Garnham, A. (1979). Instantiation of verbs. *Quarterly Journal of Experimental Psychology* 31: 207–214.

[80] Garrod, S. & Sanford, A. J. (1980). Bridging inferences and the extended domain of reference. Paper presented at 9th International Attention and Performance Conference, Cambridge, UK.

[81] Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science* 7(2): 155–170.

[82] Gentner, D. (1989). The mechanisms of analogical learning. In Stella Vosniadou, Andrew Ortony (eds.), *Similarity and Analogical Reasoning*. New York: Cambridge University Press, 199–241.

[83] Gentner, D. & Collins, A. (1987). How people construct mental models. In D. Holland, N. Quinn (eds.), *Cultural Models in Language and Thought*. Chicago, IL: University of Chicago Press, 243–265.

[84] Gentner, D. & Gentner, D. R. (1983). Flowing waters or teeming crowds: Mental models of electricity. In Dedre Gentner, A. R. Stevens, *Mental Models*. Hillsdale, NJ: Lawrence Erlbaum Associates, 99–129.

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

[86] Gibbs, R. J. (1994). *The Poetics of Mind: Figurative Thought, Language, and Understanding*. Cambridge, UK: Cambridge University Press.

[87] Gibbs, R. J. & Steen, G. (eds). (1999). *Metaphor in Cognitive Linguistics*. Philadelphia, PA: John Benjamins Publishing Company.

[88] Gick, M. L. & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology* 12(3): 306–355.

[89] Gick, M. L. & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology* 15(1): 1–38.

[90] Glenberg, A. M., Kruley, P., & Langston, W. E. (1994). Analogical processes in comprehension: Simulation of a mental model. In Morton Ann Gernsbacher (ed.), *Handbook of Psycholinguistics*. San Diego: Academic Press, 609–640.

[91] Goffman, E. (1974). *Frame Analysis: An Essay on the Organization of Experience*. New York: Harper and Row.

[92] Goldberg, A. E. (1995). *Constructions: A Construction Grammar Approach to Argument Structure*. Chicago, IL: University of Chicago Press.

[93] Goldstein, K. (1948). *Language and Language Disturbances*. New York: Grune and Stratton.

[94] Goodman, N. (1973). *Fact, Fiction, and Forecast*. Indianapolis, IN: Bobbs-Merrill.

[95] Goodman, N. (1978). *Ways of Worldmaking*. Indianapolis, IN: Hackett Publishing Company.

[96] Grady, J., Oakley, T., & Coulson, S. (1999). Blending and metaphor. In R. Gibbs, G. Steen (eds.), *Metaphor in Cognitive Linguistics*. Philadelphia, PA: John Benjamins Publishing Company, 101–124.

[97] Graesser, A. C., Woll, S. B., Kowalski, D. J., & Smith, D. A. (1980). Memory for typical and atypical actions in scripted activities. *Journal of Experimental Psychology: Human Learning and Memory* 6: 503–515.

[98] Grice, H. P. (1975). Logic and conversation. In P. Cole, J. L. Morgan (eds.), *Syntax and Semantics*. Vol. 3: *Speech Acts*. New York: Academic Press.

[99] Haberlandt, K., Berian, C., & Sandson, J. (1980). The episode schema in story processing. *Journal of Verbal Learning and Verbal Behavior* 19: 635–650.

[100] Haenggi, D., Gernsbacher, M., & Kintsch, W. (1993). Spatial situation models and text comprehension. *Discourse Processes* 19: 173–199.

[101] Hampton, J. A. (1988). Overextension of conjunctive concepts: Evidence for a unitary model of concept typicality and class inclusion. *Journal of Experimental Psychology: Learning, Memory and Cognition* 14: 12–32.

[102] Hampton, J. A. (1991). The combination of prototype concepts. In Paula J. Schwanenflugel (ed.), *The Psychology of Word Meanings*. Hillsdale, NJ: Lawrence Erlbaum Associates, 91–116.

[103] Hatcher, A. G. (1960). An introduction to the analysis of English noun compounds. *Word* 16: 356–373.

[104] Hess, D. J., Foss, D. J., & Carroll, P. (1995). Effects of global and local context on lexical processing during language comprehension. *Journal of Experimental Psychology: General* 124: 62–82.
[105] Hillyard, S. A. & Picton, T. W. (1987). Electrophysiology of cognitive processing. *Annual Review of Psychology* 34: 33–61.
[106] Hobbes, T. (1965). *Hobbe's Leviathan*. Reprinted from the edition of 1651 with an essay by the late W. G. Pogson Smith. Oxford, UK: Clarendon Press.
[107] Hofstadter, D. & Gabora, L. (1989). Frame blends. *Humor* 2: 417–440.
[108] Holland, D. & Skinner, D. (1987). Prestige and intimacy: The cultural models behind Americans' talk about gender types. In D. Holland, N. Quinn (eds.), *Cultural Models in Language and Thought*. Chicago, IL: University of Chicago Press, 78–111.
[109] Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1986). *Induction*. Cambridge, MA: MIT Press.
[110] Holyoak, K. J. & Thagard, P. (1995). *Mental Leaps: Analogy in Creative Thought*. Cambridge, MA: MIT Press.
[111] Hutchins, E. (1980). *Culture and Inference: A Trobriand Case Study*. Cambridge, MA: Harvard University Press.
[112] Indurkjya, B. (1991). *Metaphor and Cognition: An Interactionist Approach*. Dordrecht, Boston: Kluwer Academic.
[113] Israel, M. (1996). Polarity sensitivity as lexical semantics. *Linguistics and Philosophy* 19: 619–666.
[114] Israel, M. (1998). Cognitive grammar and scalarity. University of California San Diego Ph.D. Dissertation.
[115] Jackendoff, R. (1975). On belief contexts. *Linguistic Inquiry* 6: 53–93.
[116] Jarvella, R. (1973). Coreference and short-term memory for discourse. *Journal of Experimental Psychology* 98: 426–428.
[117] Jesperson, O. (1942). *A Modern English Grammar on Historical Principles*. Copenhagen, Denmark: Ejnar Munksgaard.
[118] Joanette, Y. & Goulet, P. (1987). Inferencing deficit in right brain damaged: Absence of evidence. Paper presented at the 10th European Conference of the International Neuropsychological Society, Barcelona. Cited in Joanette, Y., Goulet, P., & Hannequin, D. (1990).
[119] Joanette, Y., Goulet, P., & Hannequin, D. (1990). *Right Hemisphere and Verbal Communication*. New York: Springer-Verlag.
[120] Johnson-Laird, P. N. (1981). Mental models of meaning. In A. K. Joshi, B. L. Webber, I. A. Sag, (eds.), *Elements of Discourse Understanding*. Cambridge, UK: Cambridge University Press.
[121] Johnson-Laird, P. N. (1983). *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*. Cambridge, UK: Cambridge University Press/Cambridge, MA: Harvard University Press.
[122] Johnson-Laird, P. N. (1989). Analogy and the exercise of creativity. In S. Vosniadou, A. Ortony (eds.), *Similarity and Analogical Reasoning*. New York: Cambridge University Press, 313–331.
[123] Johnson-Laird, P. N. (1993). The mental representation of the meaning of words. In A. I. Goldman (ed.), *Readings in Philosophy and Cognitive Science*. Cambridge, MA, and London: MIT Press, 561–584.

[124] Jurafsky, D. (1996). A probabilistic model of lexical access and syntactic disambiguation. *Cognitive Science* 20: 137–194.

[125] Just, M. A., Carpenter, P., & Wooley, J. D. (1982). Paradigms and processes in reading comprehension. *Journal of Experimental Psychology: General* 111(2): 228–238.

[126] Kamp, H. (1975). Two theories about adjectives. In E. L. Keenan (ed.), *Formal Semantics of Natural Language*. Cambridge, UK: Cambridge University Press, 123–155.

[127] Kay, P. (1990). Even. *Linguistics and Philosophy* 13(1), 59–111.

[128] Kellermann, K., Broatzmann, S., Lim, T.-S., & Kitao, K. (1989). The conversation MOP: Scenes in the stream of discourse. *Discourse Processes* 12: 27–61.

[129] King, J. W. & Kutas, M. (1995). Who did what and when? Using word- and clause-related ERPs to monitor working memory usage in reading. *Journal of Cognitive Neuroscience* 7: 378–397.

[130] Kintsch, W. (1974). *The Representation of Meaning in Memory*. Hillsdale, NJ: Lawrence Erlbaum Associates.

[131] Kintsch, W. (1988). The use of knowledge in discourse processing: A construction-integration model. *Psychological Review* 95: 163–182.

[132] Kleiman, G. E. (1980). Sentence frame contexts and lexical decisions: Sentence-acceptability and word-relatedness effects. *Memory and Cognition* 8(4): 336–344.

[133] Kutas, M. (1993). In the company of other words: Electrophysiological evidence for single-word and sentence context effects. *Language and Cognitive Processes* 8: 533–572.

[134] Kutas, M. & Hillyard, S. A. (1980). Reading senseless sentences: Brain potentials reflect semantic incongruity. *Science* 207: 203–205.

[135] Kutas, M. & Hillyard, S. A. (1984). Brain potentials during reading reflect word expectancy and semantic association. *Nature* 307: 161–163.

[136] Kutas, M. & King, J. W. (1996). The potentials for basic sentence processing: Differentiating integrative processes. In T. Innue, J. McClelland (eds.), *Attention and Performance XVI: Information Integration in Perception and Communication*. Cambridge, MA: MIT Press.

[137] Kutas, M. & Van Petten, C. (1994). Psycholinguistics electrified. In M. A. Gernsbacher (ed.), *Handbook of Psycholinguistics*. San Diego, CA: Academic Press, 83–143.

[138] Kutas, M., Van Petten, C., & Besson, M. (1988). Event-related potential asymmetries during the reading of sentences. *Electroencephalography and Clinical Neurophysiology* 69: 218–233.

[139] Lakoff, G. (1987). *Women, Fire, and Dangerous Things: What Categories Reveal About the Mind*. Chicago, IL: University of Chicago Press.

[140] Lakoff, G. (1988). Cognitive semantics. In U. Eco (ed.), *Meaning and Mental Representations*. Bloomington, IN: Indiana University Press, 119–154.

[141] Lakoff, G. (1993). The contemporary theory of metaphor. In A. Ortony (ed.), *Metaphor and Thought*, Second Edition. Cambridge, UK: Cambridge University Press, 202–251.

[142] Lakoff, G. & Johnson, M. (1980). *Metaphors We Live By*. Chicago, IL: University of Chicago Press.

[143] Lakoff, G. & Kovecsces, Z. (1987). The cognitive model of anger inherent in American English. In D. Holland, N. Quinn (eds.), *Cultural Models in Language and Thought*. Chicago, IL: University of Chicago Press. 195–221.

[144] Landau, B. & Gleitman, L. (1985). *Language and Experience: Evidence from the Blind Child*. Cambridge, MA: Harvard University Press.

[145] Landman, F. (1986). *Towards a Theory of Information*. Dordrecht, Netherlands: Foris.

[146] Langacker, R. W. (1987). *Foundations of Cognitive Grammar: Theoretical Prerequisites*. Stanford, CA: Stanford University Press.

[147] Langacker, R. (1996). A dynamic account of grammatical function. Unpublished manuscript.

[148] Lange, T. E. & Dyer, M. G. (1989). High-level inferencing in a connectionist network. *Connection Science* 1: 181–217.

[149] Lebowitz, M. (1983). Generalization from natural language text. *Cognitive Science* 7: 1–40.

[150] Lehnert, W. G. (1988). The analysis of nominal compounds. In U. Eco, M. Santambrogio, P. Violi (eds.), *Meaning and Mental Representation*. Bloomington and Indianapolis, IN: Indiana University Press, 155–180.

[151] Lehrer, A. (1996). Identifying and interpreting blends: An experimental approach. *Cognitive Linguistics* 7(4): 359–391.

[152] Lemmens, M. (1994). The transitive-ergative interplay in a polysemic complex: The issue of abortion. In K. Carlon, K. Davidse, B. Rudzka-Ostyn (eds.), *Perspectives on English: Essays in Honour of Emma Vorlat*. Leuven: Peeters, 345–385.

[153] Leonard, R. (1984). *The Interpretation of English Noun Sequences on the Computer*. Amsterdam: North-Holland.

[154] Leslie, A. (1987). Pretense and representation: The origin of "theory of mind." *Psychological Review* 94: 412–426.

[155] Levi, J. (1978). *The Syntax and Semantics of Complex Nominals*. New York: Academic Press.

[156] Levinson, S. C. (1983). *Pragmatics*. New York: Cambridge University Press.

[157] Lewis, D. K. *Counterfactuals*. Cambridge, MA: Harvard University Press.

[158] Loftus, E. F. (1992). When a lie becomes memory's truth: Memory distortions after exposure to misinformation. *Current Directions in Psychological Science* 1(4): 121–123.

[159] MacDonald, S. & Wales, R. (1986). An investigation of the ability to process inferences in language following right hemisphere brain damage. *Brain and Language* 29: 68–80.

[160] Mandelblit, N. (1997). Grammatical blending: Creative and schematic aspects in sentence processing and translation. University of California San Diego, Ph.D. Dissertation.

[161] Mandler, J. (1993). The foundations of conceptual thought in infancy. *Cognitive Development* 7: 273–285.

[162] Mangun, G. R., Hillyard, S. A., & Luck, S. J. (1993). Electrocortical substrates of visual selective attention. In D. E. Meyer, S. Kornblum (eds.), *Attention and Performance 14: Synergies in Experimental Psychology, Artificial*

Intelligence, and Cognitive Neuroscience. Cambridge, MA: MIT Press, 219–243.

[163] McClelland, J. L. & Rumelhart, D. E. (eds.). (1986). *Explorations in Parallel Distributed Processing: A Handbook of Models, Programs, and Exercises.* Cambridge, MA: MIT Press.

[164] McClelland, J. L., Rumelhart, D. E., & Hinton, G. E. (1986). The appeal of parallel distributed processing. In D. E. Rumelhart, J. L. McClelland (eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition.* Vol. 1: *Foundations.* Cambridge, MA: MIT Press. 3–44.

[165] McKoon, G. & Ratcliff, R. (1986). Automatic activation of episodic information in a semantic memory task. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 12: 108–115.

[166] Medin, D. L. & Shoben, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology* 20: 158–190.

[167] Meyer, D. E. & Schvaneveldt, R. W. Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology* 90: 227–234.

[168] Michaelis, L. (1996). Cross-world continuity and the polysemy of adverbial *still.* In G. Fauconnier, E. Sweetser (eds.), *Spaces, Worlds, and Grammar.* Chicago, IL: University of Chicago Press, 179–226.

[169] Milner, A. & Goodale, M. (1995). *The Visual Brain in Action.* Oxford, UK: Oxford University Press.

[170] Minsky, M. (1975). Frame system theory. In P. N. Johnson-Laird, P. C. Wason (eds.), *Thinking: Readings in Cognitive Science.* Cambridge, UK: Cambridge University Press, 355–376.

[171] Minsky, M. (1980). Jokes and the logic of the cognitive unconscious. AI Memo No. 603. Massachusetts Institute of Technology.

[172] Montague, R. (1974). *Formal Philosophy: Selected Papers of Richard Montague.* R. H. Thomason (ed.). New Haven, CT: Yale University Press.

[173] Moxey, L. M. & Sanford, A. J. (1993). *Communicating Quantities: A Psychological Perspective.* Hillsdale, NJ: Lawrence Erlbaum Associates.

[174] Murphy, G. L. (1988). Comprehending complex concepts. *Cognitive Science* 12: 529–562.

[175] Nunberg, G. (1978). *The Pragmatics of Reference.* Bloomington, IN: Indiana University Linguistics Club.

[176] Oakley, T. van. (1995). Presence: The conceptual basis of rhetorical effect. University of Maryland, Ph.D. Dissertation.

[177] Osherson, D. N. & Smith, E. E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition* 9: 35–58.

[178] Osherson, D. N. & Smith, E. E. (1983). Gradedness and conceptual combination. *Cognition* 12: 299–318.

[179] Perelman, C. & Olbrechts-Tyteca, L. (1969). *The New Rhetoric: A Treatise on Argumentation.* Translated by John Wilkinson and Purcell Weaver. Notre Dame, IN: University of Notre Dame Press.

[180] Perfetti, C. A. & Hogaboam, T. (1975). Relationship between single word decoding and reading comprehension skill. *Journal of Educational Psychology* 67: 461–469.

[181] Pickering, M. & Traxler, M. (1995). Plausibility and the processing of unbounded dependencies: An eye-tracking study. *Journal of Memory and Language* 35(3): 454–475.

[182] Potter, M. C. & Faulconer, B. A. (1979). Understanding noun phrases. *Journal of verbal learning and verbal behavior* 18: 509–522.

[183] Quinn, N. & Holland, D. (1987). Culture and cognition. In D. Holland, N. Quinn (eds.), *Cultural Models in Language and Thought*. Chicago, IL: University of Chicago Press, 3–43.

[184] Quintillian (1921–1933). *The Institutio Oratoria of Quintillian*. Translated by H. E. Butler, Loeb Classical Library. New York: G. P. Putnam's Sons.

[185] Reiter, R. & Criscuolo, G. (1981). On interacting defaults. *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, 270–276.

[186] Rips, L. J. (1995). The current status of research on concept combination. *Mind & Language* 10: 72–104.

[187] Rizzolatti, G. & Craighero, L. (1998). Spatial attention: Mechanisms and theories. In M. Sabourin, F. Craik (eds.), *Advances in Psychological Science*. Vol. 2: *Biological and Cognitive Aspects*. Hove, UK: Taylor & Francis, 171–198.

[188] Robin, N. & Holyoak, K. J. (1995). Relational complexity and the functions of pre-frontal cortex. In M. I. Gazzaniga (ed.), *The Cognitive Neurosciences*. Cambridge, MA: MIT Press.

[189] Ross, E. (1981). The aprosodias. Functional-anatomic organization of the affective components of language in the right hemisphere. *Archives of Neurology* 38: 561–569.

[190] Roth, E. & Shoben, E. J. (1983). The effect of context on the structure of categories. *Cognitive Psychology* 15: 346–378.

[191] Rumelhart, D. E. (1980;1987). Understanding understanding. Paper presented at international Workshop on "Psycholinguistic Models on Production," held in Kassel, Federal Republic of Germany, from July 13–17. In Hans W. Dechert, Manfred Raupach (eds.), *Psycholinguistic Models of Production*. Norwood, NJ: Ablex, 209–228.

[192] Rumelhart, D. E. (1980). On evaluating story grammars. *Cognitive Science* 4(3): 313–316.

[193] Rumelhart, D. E. (1993). Some problems with the notion of literal meanings. In Andrew Ortony (ed.), *Metaphor and Thought*, Second Edition. Cambridge, UK: Cambridge University Press, 71–82.

[194] Rumelhart, D. E. & Ortony, A. (1977). The representation of knowledge in memory. In R. C. Anderson, R. J. Spiro, W. E. Montague (eds.), *Schooling and the Acquisition of Knowledge*. Hillsdale, NJ: Lawrence Erlbaum Associates.

[195] Rumelhart, D. E., Smolensky, P., McClelland, J. L., & Hinton, G. (1986). Schemata and sequential thought processes in PDP models. In D. E. Rumelhart, J. L McClelland (eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. Vol. 2: *Psychological and Biological Models*. Cambridge, MA: MIT Press, 7–57.

[196] Ryder, M. E. (1994). *Ordered Chaos: The Interpretation of English Noun-Noun Compounds*. Berkeley, CA: University of California Press.

[197] Sanford, A. J. & Garrod, S. C. (1980). A demonstration of the situational basis of text-comprehension through implicit assignments of roles to

entities. Unpublished manuscript, Department of Psychology, University of Glasgow.

[198] Sanford, A. J. & Garrod, S. C. (1981). *Understanding Written Language: Explorations Beyond the Sentence*. Chichester, UK: John Wiley & Sons.

[199] Sawaguchi, T. & Goldman-Rakic, P. S. (1991). D1 dopamine receptors in prefrontal cortex: Involvement in working memory. *Science* 251: 947–950.

[200] Schank, R. C. (1982). Reminding and memory organization: An introduction to MOPS. In W. G. Lehnert, M. H. Ringle (eds.), *Strategies for Natural Language Processing*. Hillsdale, NJ: Lawrence Erlbaum Associates.

[201] Schank, R. C. (1983). *Dynamic Memory: A Theory of Learning in Computers and People*. Cambridge, UK: Cambridge University Press.

[202] Schank, R. C. & Abelson, R. P. (1977). *Scripts, Plans, Goals, and Understanding: An Inquiry into Human Knowledge Structures*. Hillsdale, NJ: Lawrence Erlbaum Associates.

[203] Schwanenflugel, P. J. & LaCount, K. L. (1988). Semantic relatedness and the scope of facilitation for upcoming words in sentences. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 14(2): 344–354.

[204] Schwanenflugel, P. J. & Shoben, E. J. (1985). The influence of sentence constraint on the scope of facilitation for upcoming words. *Journal of Memory and Language* 24: 232–252.

[205] Searle, J. R. (1979). *Expression and Meaning: Studies in the Theory of Speech Acts*. Cambridge, UK: Cambridge University Press.

[206] Sharkey, N. E. & Mitchell, D. C. (1988). Word recognition in a functional context: The use of scripts in reading. *Journal of Memory and Language* 24: 253–270.

[207] Shastri, L. & Ajjanagadde, V. (1993). From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. *Behavioral & Brain Sciences* 16(3): 417–494.

[208] Shoben, E. J. (1991). Predicating and nonpredicating combinations. In Paula J. Schwanenflugel (ed.), *The Psychology of Word Meanings*. Hillsdale, NJ: Lawrence Erlbaum Associates, 117–136.

[209] Simpson, G. B. (1984). Lexical ambiguity and its role in models of word recognition. *Psychological Bulletin* 96: 316–340.

[210] Simpson, G. B. (1994). Context and the processing of ambiguous words. In Morton Ann Gernsbacher (ed.), *Handbook of Psycholinguistics*. San Diego, CA: Academic Press, 359–374.

[211] Singer, J. L. (1995). Imaginative play in childhood: Precursor of subjunctive thoughts, daydreaming, and adult pretending games. In Anthony D. Pellegrini (ed.), *The Future of Play Theory: A Multidisciplinary Inquiry into the Contributions of Brian Sutton-Smith*. Albany, NY: State University of New York Press.

[212] Smith, L. B. & Jones, S. S. (1993). Cognition without concepts. *Cognitive Development* 8(2): 181–188.

[213] Smith, L. B., Ratterman, M. J., & Sera, M. D. (1988). "Higher" and "lower": Comparative categorical interpretations by children. *Cognitive Development* 3(4): 341–357.

[214] Smith, M. E. & Halgren, E. (1987). Event-related potentials during lexical decision: Effects of repetition, word frequency, pronounceability, and concreteness. *Electroencephalography & Clinical Neurosphysiology*, Suppl. 40: 417–421.

[215] Sperber, D. & Wilson, D. (1987). Precis of *Relevance: Communication and Cognition. Behavioral and Brain Sciences* 10: 697–754.

[216] Spivey-Knowlton, M. & Tannenhaus, M. (1994). Referential context and syntactic ambiguity resolution. In Charles Clifton, Jr., Lynn Frazier, Keith Rayner (eds.), *Perspectives on Sentence Processing*. Hillsdale, NJ: Lawrence Erlbaum Associates.

[217] Springer, K. & Murphy, G. (1992). Feature availability in conceptual combination. *Psychological Science* 2: 111–117.

[218] Stanovich, K. E. & West, R. F. (1979). Mechanisms of sentence context effects in reading: Automatic activation and conscious attention. *Memory and Cognition* 7(2): 77–85.

[219] Steele, J. (1992). TV's talking headaches. *Columbia Journalism Review* 31(2): 49–52.

[220] St. George, M., Mannes, S., & Hoffman, J. E. (1994). Global semantic expectancy and language comprehension. *Journal of Cognitive Neuroscience* 6: 70–83.

[221] St. John, M. F. (1992). The story gestalt: A model of knowledge intensive processes in text comprehension. *Cognitive Science* 16: 271–306.

[222] Suchman, L. (1987;1990). *Plans and Situated Actions*. Cambridge, UK: Cambridge University Press.

[223] Suls, J. M. (1983). Cognitive processes in humour appreciation. In P. E. McGhee, J. H. Goldstein (eds.), *Handbook of Humor Research*. New York: Springer-Verlag, 39–55.

[224] Sweetser, E. (1987). The definition of *lie*: An examination of the folk models underlying a semantic prototype. In D. Holland, N. Quinn (eds.), *Cultural Models in Language and Thought*. Chicago, IL: University of Chicago Press.

[225] Sweetser, E. (1990). *From Etymology to Pragmatics: Metaphorical and Cultural Aspects of Semantic Structure*. Cambridge, UK: Cambridge University Press.

[226] Sweetser, E. & Fauconnier, G. (1996). Cognitive links and domains. In Gilles Fauconnier, Eve Sweetser (eds.), *Spaces, Worlds, and Grammar*. Chicago, IL: University of Chicago Press.

[227] Swinney, D. (1979). Lexical access during sentence comprehension. (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior* 18: 645–659.

[228] Tabossi, P. (1983). Interpreting words in context. University of Sussex, D. Phil. Thesis.

[229] Tabossi, P. & Johnson-Laird, P. N. (1980). Linguistic context and the priming of semantic information. *Quarterly Journal of Experimental Psychology* 32: 595–603.

[230] Taylor, J. R. (1995). *Linguistic Categorization: Prototypes in Linguistic Theory*. Oxford, UK: Clarendon Press.

[231] Travis, C. (1981). *The True and the False: The Domain of the Pragmatic*. Amsterdam: Benjamins.

[232] Traxler, M., Sanford, A., Aked, J., & Moxey, L. (1997). Processing causal and diagnostic statements in discourse. *Journal of Experimental Psychology: Learning, Memory, & Cognition* 23: 88–101.

[233] Tucker, D. M. (1993). Spatial sampling of head electrical fields: The geodesic sensor net. *Electroencephalography and Clinical Neurophysiology* 87: 154–163.

[234] Turner, M. (1991). *Reading Minds: The Study of English in the Age of Cognitive Science*. Princeton, NJ: Princeton University Press.

[235] Turner, M. (1996). Conceptual blending and counterfactual argument in the social and behavioral sciences. In Philip Tetlock, Aaron Belkin (eds.), *Counterfactual Thought Experiments in World Politics*. Princeton, NJ: Princeton University Press, 291–295.

[236] Turner, M. & Fauconnier, G. (1995). Conceptual integration and formal expression. *Metaphor and Symbolic Activity* 10(3): 183–204.

[237] Turner, M. & Fauconnier, G. (1998). Conceptual integration in counterfactuals. In Koenig, J.-P. (ed.), *Discourse and Cognition*. Stanford, CA: Center for the Study of Language and Information (CSLI), 285–296.

[238] Ungerleider, L. & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, R. J. Mansfield (eds.), *Analysis of Visual Behavior*. Cambridge, MA: MIT Press, 549–556.

[239] Van Dijk, T. & Kintsch, W. (1983). *Strategies of Discourse Comprehension*. New York: Academic Press.

[240] Van Essen, D., Anderson, C., & Felleman, D. (1992). Information processing in the primate visual system: An integrated systems perspective. *Science* 255: 419–423.

[241] Van Petten, C. (1993). A comparison of lexical and sentence-level context effects in event-related potentials. *Language and Cognitive Processes* 8: 485–531.

[242] Van Petten, C. & Kutas, M. (1990). Interactions between sentence context and word frequency in event-related potentials. *Memory and Cognition* 18: 380–393.

[243] Van Petten, C. & Kutas, M. (1991). Influences of semantic and syntactic context on open and closed class words. *Memory and Cognition* 19: 95–112.

[244] von Eckhart, B. & Potter, M. (1985). Clauses and the semantic representation of words. *Memory and Cognition* 13: 371–376.

[245] Wapner, W., Hamby, S., & Gardner, H. (1981). The role of the right hemisphere in the apprehension of complex linguistic materials. *Brain and Language* 25: 144–159.

[246] Wilensky, R. (1986). Points: A theory of the structure of stories in memory. In B. J. Grosz, K. S. Jones, B. L. Webber (eds.), *Readings in Natural Language Processing*. Los Altos, CA: Morgan Kaufman Publishing, 459–473.

[247] Yeh, W. & Rips, L. (1992). Emergent properties for familiar and unfamiliar adjective-noun combinations. Unpublished raw data.

[248] Zadeh, L. A. (1965). Fuzzy sets. *Information and Control* 8: 338–353.

[249] Zwaan, R. A. & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin* 123: 162–185.

Index