Edge

Summer Reading from the Archive

PHILOSOPHY IN THE FLESH

A Conversation with George Lakoff

[March 1999]

> *We are neural beings. Our brains take their input from the rest of our bodies. What our bodies are like and how they function in the world thus structures the very concepts we can use to think. We cannot think just anything—only what our embodied brains permit.*

Introduction

"We are neural beings," states Berkeley cognitive scientist George Lakoff. "Our brains take their input from the rest of our bodies. What our bodies are like and how they function in the world thus structures the very concepts we can use to think. We cannot think just anything—only what our embodied brains permit."

His new book *Philosophy In The Flesh*, co-authored by Mark Johnson, makes the following points: "The mind is inherently embodied. Thought is mostly unconscious. Abstract concepts are largely metaphorical."

Lakoff believes that new empirical evidence concerning these finding of cognitive science have taken us over the epistemological divide: we are in a new place and our philosophical assumptions are all up for grabs.

He and Johnson write: "When taken together and considered in detail, these three findings from the science of the mind are inconsistent with central parts of Western philosophy, and require a thorough rethinking of the most popular current approaches, namely, Anglo-American analytic philosophy and postmodernist philosophy."

According to Lakoff, metaphor appears to be a neural mechanism that allows us to adapt the neural systems used in sensory-motor activity to create forms of abstract reason. "If this is correct, as it seems to be," he says, "our sensory-motor systems thus limit the abstract reasoning that we can perform. Anything we can think or understand is shaped by, made possible by, and limited by our bodies, brains, and our embodied interactions in the world. This is what we have to theorize with."

He then raises the interesting question: "Is it adequate to understand the world scientifically?"

—JB

—

GEORGE LAKOFF is emeritus professor of cognitive science and linguistics at UC Berkeley and director of the Center for the Neural Mind & Society. Among his books are *Don't Think of an Elephant*, *The Political Mind*, and *Metaphors We Live By*.

PHILOSOPHY IN THE FLESH

JB: What is a body?

LAKOFF: That's an interesting question. Pierre Bourdieu has pointed out that our bodies and what we do with them differ significantly from culture to culture. Frenchmen do not walk like Americans do. Women's bodies are different than men's bodies. The Chinese body is not like the Polish body. And our understanding of what the body is has changed drastically over time, as postmodernists have often observed.

But nonetheless, our bodies do share a lot. We have two eyes, two ears, two arms, two legs, blood that circulates, lungs used to breathe, skin, internal organs, and on and on. The common conventionalized aspects of our conceptual systems tend to be structured by what our bodies have in common, which is a lot.

JB: But we go from being a machine to an information system, and eventually those orifices may not be part of the conversation.

LAKOFF: When you start to study the brain and body scientifically, you inevitably wind up using metaphors. Metaphors for the mind, as you say, have evolved over time—from machines to switchboards to computers. There's no avoiding metaphor in science. In our lab, we use the Neural Circuitry metaphor ubiquitous throughout neuroscience. If you're studying neural computation, that metaphor is necessary. In the day to day research on the details of neural computation, the biological brain moves into the background while the Neural Circuitry introduced by the metaphor is what one works with. But no matter how ubiquitous a metaphor may be, it is important to keep track of what it hides and what it introduces. If you don't, the body does disappear. We're careful about our metaphors, as most scientists should be..

JB: There were no information processing metaphors 35-40 years ago—and so is the body real, or is it invented?

LAKOFF: There's a difference between the body and our conceptualization of it. The body is the same as it was 35 years ago; the conception of the body is very different. We have metaphors for the body we didn't have then, with relatively advanced science built on those metaphors. In this respect, the contemporary body and brain, conceptualized in terms of neural circuitry and other information processing metaphors, were "invented." Such inventions are crucial to science. Our emerging understanding of the embodiment of mind would not be possible without them.

JB: How does this approach depart from your early work?

LAKOFF: My really early work was done between 1963 and 1975, when I was pursuing the theory of Generative Semantics. During that period, I was attempting to unify Chomsky's transformational grammar with formal logic. I had helped work out a lot of the early details of Chomsky's theory of grammar. Noam claimed then—and still does, so far as I can tell—that syntax is independent of meaning, context, background knowledge, memory, cognitive processing, communicative intent, and every aspect of the body.

In working through the details of his early theory, I found quite a few cases where semantics, context, and other such factors entered into rules governing the syntactic occurrences of phrases and morphemes. I came up with the beginnings of an alternative theory in 1963 and, along with wonderful collaborators like Haj Ross and Jim McCawley, developed it through the sixties. Back in 1963, semantics meant logic—deductive logic and model theory—and our group developed a theory of Generative Semantics that united formal logic and transformational grammar. In that theory, semantics (in the form of logic) was taken as prior to syntax on the basis of evidence that semantic and pragmatic considerations entered into generalizations governing syntactic structure. Chomsky has since adopted many of our innovations, though he fought them viciously in the '60s and '70s.

In 1975, I became acquainted with certain basic results from the various cognitive sciences pointing toward an embodied theory of mind—the neurophysiology of color vision, prototypes and basic-level categories, Talmy's work on spatial relations concepts, and Fillmore's frame semantics. These results convinced me that the entire thrust of research in generative linguistics and formal logic was hopeless. I set about, along with Len Talmy, Ron Langacker, and Gilles Fauconnier, to form a new linguistics—one compatible with research in cognitive science and neuroscience. It is called Cognitive Linguistics, and it's a thriving scientific enterprise. In 1978, I discovered that metaphor was not a minor kind of trope used in poetry, but rather a fundamental mechanism of mind. In 1979, Mark Johnson visited in the Berkeley Philosophy Department and we began working out the details and their implications for philosophy. We've been collaborating for 20 years. Mark is now Chair of Philosophy at Oregon.

JB: Distinguish cognitive science from philosophy?

LAKOFF: That is a deep and important question, and central to the enterprise of *Philosophy In The Flesh*. The reason that the question doesn't have a simple answer is that there are two forms of cognitive science, one fashioned on the assumptions of Anglo-American philosophy and one (so far as we can tell) independent of specific philosophical assumptions that determine the results of the inquiry.

Early cognitive science, what we call "first-generation" cognitive science (or "disembodied cognitive science"), was designed to fit a formalist version of Anglo-

American philosophy. That is, it had philosophical assumptions that the determined important parts of the content of the scientific "results." Back in the late 1950s, Hilary Putnam (a noted and very gifted philosopher) formulated a philosophical position called "functionalism." (Incidentally, he has since renounced that position.) It was an apriori philosophical position, not based on any evidence whatever. The proposal was this:

The mind can be studied in terms of its cognitive functions—that is, in terms of the operations it performs—independently of the brain and body.

The operations performed by the mind can be adequately modeled by the manipulation of meaningless formal symbols, as in a computer program.

This philosophical program fit paradigms that existed at the time in a number of disciplines.

In formal philosophy:

The idea that reason could be adequately characterized using symbolic logic, which utilizes the manipulation of meaningless formal symbols.

In generative linguistics:

The idea that the grammar of a language can be adequately characterized in terms of rules that manipulate meaningless formal symbols.

In artificial intelligence:

The idea that intelligence in general consists in computer programs that manipulate meaningless formal symbols.

In information processing psychology:

The idea that the mind is an information-processing device, where information-processing is taken as the manipulation of meaningless formal symbols, as in a computer program.

All of these fields had developed out of formal philosophy. These four fields converged in the 1970s to form first-generation cognitive science. It had a view of mind as the disembodied manipulation of meaningless formal symbols.

JB: How does this fit into empirical science?

LAKOFF: This view was not empirically based, having arisen from an apriori philosophy. Nonetheless it got the field started. What was good about it was that it was precise. What was disastrous about it was that it had a hidden philosophical worldview that mascaraded as a scientific result. And if you accepted that philosophical position, all results inconsistent with that philosophy could only be seen as nonsense. To researchers trained in that tradition, cognitive science was the study of mind within that apriori philosophical position. The first generation of cognitive scientists was trained to think that way, and many textbooks still portray cognitive science in that way. Thus, first generation cognitive science is not distinct from philosophy; it comes with an apriori philosophical worldview that places substantive constraints on what a "mind" can be. Here are some of those constraints:

Concepts must be literal. If reasoning is to be characterized in terms of traditional formal logic, there can be no such thing as a metaphorical concept and no such thing as metaphorical thought.

Concepts and reasoning with concepts must be distinct from mental imagery, since imagery uses the mechanisms of vision and cannot be characterized as being the manipulation of meaningless formal symbols.

Concepts and reasoning must be independent of the sensory-motor system, since the sensory motor system, being embodied, cannot be a form of disembodied abstract symbol-manipulation.

Language too—if it was to fit the symbol-manipulation paradigm—had to be literal, independent of imagery, and independent of the sensory-motor system.

From this perspective, the brain could only be a means to implement abstract "mind"—wetware on which the "programs of the mind" happened to be implementable. Mind on this view does not arise from and is not shaped by the brain. Mind is a disembodied abstraction that our brains happen to be able to implement. These were not empirical results, but rather followed from philosophical assumptions.

In the mid-1970s, cognitive science was finally given a name and outfitted with a society and a journal. The people who formed the field accepted the symbol-manipulation paradigm. I was originally one of them (on the basis of my early work on generative semantics) and gave one of the invited inaugural lectures at the first meeting of the Cognitive Science Society. But just around the time that the field officially was recognized and organized around the symbol-manipulation paradigm, empirical results started coming in calling the paradigm itself into question.

This startling collection of results pointed toward the idea that mind was not disembodied—not characterizable in terms of the manipulation of meaningless symbols independent of the brain and body, that is, independent of the sensory-motor system and our functioning in the world. Mind instead is embodied, not in the trivial sense of being implementable in a brain, but in the crucial sense that conceptual structure and the mechanisms of reason arise ultimately and are shaped by from the sensory-motor system of the brain and body.

JB: Can you prove it?

LAKOFF: There is a huge body of work supporting this view. Here are some of the basic results that have interested me the most: The structure of the system of color categories is shaped by the neurophysiology of color vision, by our color cones and neural circuitry for color. Colors and color categories are not "out there" in the world but are interactional, a nontrivial product of wave length reflactances of objects and lighting conditions on the one hand, and our color cones and neural circuitry on the other. Color concepts and color-based inferences are thus structured by our bodies and brains.

Basic-level categories are structured in terms of gestalt perception, mental imagery, and motor schemas. In this way the body and the sensory-motor system of the brain enters centrally into our conceptual systems.

Spatial relations concepts in languages around the world (e.g, in, through, around in English, sini in Mixtec, mux in Cora, and so on) are composed of the same primitive "image-schemas," that is, schematic mental images. These, in turn, appear to arise from the structure of visual and motor systems. This forms the basis of an explanation of how we can fit language and reasoning to vision and movement.

Aspectual concepts (which characterize the structure of events) appear to arise from neural structures for motor control.

Categories make use of prototypes of many sorts to reason about the categories as a whole. Those prototypes are characterized partly in terms of sensory-motor information.

The conceptual and inferential system for reasoning about bodily movements can be performed by neural models that can model both motor control and inference. Abstract concepts are largely metaphorical, based on metaphors that make use of our sensory-motor capacities to perform abstract inferences. Thus, abstract reason, on a large scale, appears to arise from the body.

These are the results most striking to me. They require us to recognize the role of the body and brain in human reason and language. They thus run contrary to any notion of a disembodied mind. It was for such reasons that I abandoned my earlier work on generative semantics and started studying how mind and language are embodied. They are among the results that have led to a second-generation of cognitive science, the cognitive science of the embodied mind.

JB: Let's get back to my question about the difference between cognitive science and philosophy.

LAKOFF: OK. Cognitive science is the empirical study of the mind, unfettered by apriori philosophical assumptions. First-generation cognitive science, which posed a disembodied mind, was carrying out a philosophical program. Second-generation cognitive science, which is working out the nature of the mind as it really is—embodied!—had to overcome the built-in philosophy of earlier cognitive science.

JB: Does "second-generation cognitive science" presuppose a philosophy?

LAKOFF: We have argued that it does not, that it simply presupposes commitments to take empirical research seriously, seek the widest generalizations, and look for convergent evidence from many sources. That is just what science is committed to. The results about the embodied mind did not begin from, and does not presuppose, any particular philosophical theory of mind. Indeed, it has required separating out the old philosophy from the science.

JB: Where does this leave philosophy?

LAKOFF: In a position to start over from an empirically responsible position. Young philosophers should be thrilled. Philosophy is anything but dead. It has to be rethought taking the empirical results about the embodied mind into account. Philosophy considers the deepest questions of human existence. It is time to rethink them and that is an exciting prospect.

JB: What about the academic wars between postmodern and analytic philosophy?

LAKOFF: The results suggest that both sides were insightful in some respects and mistaken in others. The postmodernists were right that some concepts can change over time and vary across cultures. But they were wrong in suggesting that they all concepts are like that. Thousands are not. They arise around the world in culture after culture from our common embodiment.

Postmodernists were right in observing that there are many places where the folk theory of essences fails. But they were wrong in suggesting that such a failure undercuts our conceptual systems and makes them arbitrary. The analytic tradition insightfully characterized the theory of speech acts. Although formal logic does not work for all, or even most, of reason, there are places where something akin to formal logic (much revised) does characterize certain limited aspects of reason. But the analytic tradition was wrong in certain of its central theses: the correspondence theory of truth, the theory of literal meaning, and the disembodied nature of reason.

The academic world is now in a position to transcend both positions, each having contributed something important and each needing revision.

JB: Is there an East Coast and West Coast divide?

LAKOFF: Dan Dennett referred to the "East Pole" and "West Pole" back in the early-to-mid 1980s, as if the proponents of the disembodied mind were all on the East Coast and the proponents of the embodied mind were all on the West Coast. Research on the embodied mind did tend to start on the West Coast, but even then the geographical characterization was oversimplified. By now, both positions are represented on both coasts and throughout the country. Cambridge and Princeton in the past have tended largely toward the old disembodied mind position, at least in certain fields. But there are so many interesting thinkers on both coasts and spread across the country that I think that any geographical divisions that still exist won't last long.

When Dennett first made that distinction, the great revolutions in neuroscience and neural modeling were just starting. Cognitive linguistics was just coming into existence. *Metaphors We Live By* had barely come out and *Women, Fire, And Dangerous Things* had not yet been written. Nor had Edelman's *Bright Air, Brilliant Fire* nor Damasio's *Descartes Error*, nor Regier's *The Human Semantic Potential*, nor the various books by Pat and Paul Churchland. Over the past decade and a half, neuroscience and neural computation have changed the landscape of cognitive science and they will change it even more in next decade or two. Those changes will inevitably move us further toward an appreciation of the embodiment of mind. You cannot think anything without using the neural system of your brain. The fine structure of neural connections in the brain, their connections to the rest of the body, and the nature of neural computation will keep being developed. The more we discover about the details, the more we will come to understand the detailed nature of how reason and the conceptual systems in which we reason are embodied.

The idea of disembodied reason was an apriori philosophical idea. It lasted 2500 years. I can't imagine it lasting another 30 years in serious scientific circles.

JB: And what do we have to look forward to?

LAKOFF: Cognitive science and neuroscience are triggering a philosophical revolution. Philosophy In The Flesh is just part of the first wave. Over the next decade or two, the neural theory of language should develop sufficiently to replace the old view of language as meaningless disembodied symbol manipulation that one finds in the old Chomskyan tradition. But the biggest, and one of the most important, changes will come in our understanding of mathematics.

The precursor of that change was Stanislas Dehaene's *The Number Sense*, which reviewed the evidence from neuroscience, child development, and animal research indicating that we (and certain other animals) have evolved with a part of our brains dedicated to enumeration and simple arithmetic up to a small number of objects (around four). Rafael Núñez and I begin with those findings and ask how sophisticated arithmetic (with the laws of arithmetic) developed, that is, how could ordinary conceptual mechanisms for human thought have given rise to mathematics?

Our answer is that the ordinary embodied mind, with its image schemas, conceptual metaphors, and mental spaces, has the capacity to create the most sophisticated of mathematics via using everyday conceptual mechanisms. Dehaene stopped with simple arithmetic. We go on to show that set theory, symbolic logic, algebra, analytic geometry, trigonometry, calculus, and complex numbers can all be accounted for using those everyday conceptual mechanisms. Moreover, we show that conceptual metaphor is at the heart of the development of complex mathematics. This is not hard to see. Think of the number line. It is the result of a metaphor that Numbers Are Points on a Line. Numbers don't have to be thought of as points on a line. Arithmetic works perfectly well without being thought of in terms of geometry. But if you use that metaphor, much more interesting mathematics results. Or take the idea, in set-theoretical foundations for arithmetic, that Numbers Are Sets, with zero as the empty set, one as the set containing the empty set, and so on. That's a metaphor too. Numbers don't have to be thought of as being sets. Arithmetic went on perfectly well for 2000 years without numbers being conceptualized as sets. But if you use that metaphor, then interesting mathematics results. There is a third less well-known metaphor for numbers, that Numbers Are Values of Strategies in combinatorial game theory. So which is it? Are numbers points? Are they sets? Are numbers fundamentally just values of strategies in combinatorial games?

These metaphors for numbers are part of the mathematics, and you make a choice each time depending on the kind of mathematics you want to be doing. The moral is simple: Conceptual metaphor is central to conceptualization of number in mathematics of any complexity at all. It's a perfectly sensible idea. Conceptual metaphors are cross-domain mappings that preserve inferential structure. Mathematical metaphors are what provide the links across different branches of mathematics. One of our most interesting results concerns the conceptualization of infinity. There are many concepts that involve infinity: points at infinity in projective and inversive geometry, infinite sets, infinite unions, mathematical induction, transfinite numbers, infinite sequences, infinite decimals, infinite sums, limits, least upper bounds, and infinitesimals. Núñez and I have found that all of these concepts can be conceptualized as special cases of one simple Basic Metaphor of Infinity. The idea of "actual infinity"—of infinity not just as going on and on, but as a thing—is metaphorical, but the metaphor, as we show turns out to quite simple and exists outside of mathematics. What mathematicians have done is to provide elaborate carefully devised special cases of this basic metaphorical idea.

What we conclude is that mathematics as we know it is a product of the human body and brain; it is not part of the objective structure of the universe—this or any other. What our results appear to disprove is what we call the Romance of Mathematics, the idea that mathematics exists independently of beings with bodies and brains and that mathematics structures the universe independently of any embodied beings to create the mathematics. This does not, of course, result in the idea that mathematics is an arbitrary product of culture as some postmodern theorists would have it. It simply says that it is a stable product of our brains, our bodies, our experience in the world, and aspects of culture. The explanation of why mathematics "works so well" is simple: it is the result of tens of thousands of very smart people observing the world carefully and adapting or creating mathematics to fit their observations. It is also the result of a mathematical evolution: a lot of mathematics invented to fit the world turned out not to. The forms of mathematics that work in the world are the result of such an evolutionary process.

It is important to know that we create mathematics and to understand just what mechanisms of the embodied mind make mathematics possible. It gives us a more realistic appreciation of our role in the universe. We, with our physical bodies and brains, are the source of reason, the source of mathematics, the source of ideas. We are not mere vehicles for disembodied concepts, disembodied reason, and disembodied mathematics floating out there in the universe. That makes each embodied human being (the only kind) infinitely valuable—a source not a vessel. It makes bodies infinitely valuable—the source of all concepts, reason, and mathematics.

For two millenia, we have been progressively devaluing human life by underestimating the value of human bodies. We can hope that the next millenium, in which the embodiment of mind will come to be fully appreciated, will be more humanistic.

JB: Where are you headed next?

LAKOFF: I've plunged myself as fully as possible into the research that Jerry Feldman and I have been doing for the past decade at the International Computer Science Institute on the Neural Theory of Language. That's where most of my technical research effort is going to go for quite a while.

Jerry developed the theory of structured connectionism (not PDP connectionism) beginning in the 1970s. Structured connectionism allows us to construct detailed computational neural models of conceptual and linguistic structures and of the learning of such structures.

Since 1988, we've been running a project takes up a question that has absorbed both of us: From the perspective of neural computation, a human brain consists of a very large number of neurons connected up in specific ways with certain computational properties. How is it possible to get the details of human concepts, the forms of human reason, and the range of human languages out of a lot of neurons connected up as they are in our brains? How do you get thought and language out of neurons? That is the question we are trying to answer in our lab through the computational neural modeling of thought and language.

JB: How do you connect structures in the brain to ideas of space?

LAKOFF: Terry Regier has taken the first step to figuring that out in his book *The Human Semantic Potential*. He has hypothesized that certain types of brain structures—topographic maps of the visual field, orientation-sensitive cells, and so on—can compute the primitive spatial relations (called "image-schemas") that linguists have discovered. The amazing thing to me is that not only do we actually have a reasonable idea of how certain types of neural structures can give rise to spatial relations concepts. Recent neural modeling research by Narayanan has similarly given us an idea of how brain structures can compute aspectual concepts (which structure events), conceptual metaphors, mental spaces, blended spaces, and other basics of human conceptual systems. The next breakthrough, I think, will be a neural theory of grammar.

These are remarkable technical results. When you put them together with other results about the embodiment of mind coming from neuroscience, psychology, and cognitive linguistics, they tell us a great deal about things that are important in the everyday lives of ordinary people—things that philosophers have speculated about for over 2500 years. Cognitive science has important things to tell us about our understanding of time, events, causation, and so on.

JB: Like what?

When Mark Johnson and I looked over these results from the cognitive sciences in detail, we realized that there were three major results that were inconsistent with almost all of Western philosophy (except for Merleau-Ponty and Dewey), namely:

The mind is inherently embodied.

Most thought is unconscious.

Abstract concepts are largely metaphorical.

This realization led us to ask the following question in *Philosophy In The Flesh*: What would happen if we started with the new results about the mind and reconstructed philosophy from there? What would philosophy look like?

It turns out that it looks entirely different from virtually all the philosophy that went before. And the differences are differences that matter in your life. Starting with results from cognitive semantics, we discovered a lot that is new about the nature of moral systems, about the ways that we conceptualize the internal structure of the Self, even about the nature of truth.

JB: This seems like a distinctively new kind of enterprise.

LAKOFF: It's an interesting enterprise to take philosophy as a subject matter for empirical study in cognitive science. Most philosophers take philosophy as an apriori discipline, where no empirical study of the mind, reason, and language is necessary. In the Anglo-American tradition, you are taught to think like a philosopher and then it is assumed that you can, on the basis of your philosophical training, make pronouncements about any other discipline. Thus, there are branches of philosophy like the Philosophy of Language, the Philosophy of Mind, the Philosophy of Mathematics, and so on. Johnson and I realized that philosophy itself, which consists of systems of thought, needed to be studied from the perspective of the cognitive sciences, especially cognitive semantics, which studies systems of thought empirically. Our goal has been to bring a scientific perspective to philosophy, especially a perspective from the science of mind.

JB: How does this connect with traditional philosophy?

Lakoff: It is a startling thing to realize that most of Western philosophy is inconsistent with fundamental results from the science of the mind. But that is negative. We respect and value philosophy. Our work comes out of a deep love for philosophy and a disappointment over what it has been over the past couple of decades. We wanted to look at great moments in the history of philosophy—the Presocratics, Plato, Aristotle, Descartes, Kant—even the analytic philosophers—and show what shining the light of cognitive science oh philosophy could reveal about the nature of philosophy.

What we discovered was fascinating: Each major philosopher seems to take a small number of metaphors as eternal and self-evident truths and then, with rigorous logic and total systematicity, follows out the entailments of those metaphors to their conclusions wherever they lead. They lead to some pretty strange places. Plato's metaphors entail that philosophers should govern the state. Aristotle's metaphors entail that there are four causes and that there cannot be a vacuum. Descartes' metaphors entail that the mind is completely disembodied and that all thought is conscious. Kant's metaphors lead to the conclusions that there is a universal reason and that it dictates universal moral laws. These and other positions taken by those philosophers are not random opinions. They are consequences of taking commonplace metaphors as truths and systematically working out the consequences.

JB: What's the import of recognizing that metaphors are central to the work of earlier philosophers?

LAKOFF: It is not just earlier philosophers, but contemporary philosophers as well. Our moral is not that their work should be disregarded because it is metaphorical. Quite the opposite. Because most abstract thought is, and has to be, metaphorical, all rigorous abstract systems of thought will be like those of the great philosophers whose systems of thought we analyze. Moreover, everyone's everyday reasoning is often of the same character, though hardly as consistent overall. A cognitive perspective on philosophy not only teaches us how the great philosophers thought, but it gives us deep insights on how all of us think—at least when we're being consistent and systematic. It also tells us that, in most cases, the answers to the deepest questions of human existence will most likely be metaphorical answers. There is nothing wrong with this. We just need to be aware of just what our metaphors are and what they entail.

Another positive thing we sought to do was to look at the most fundamental of philosophical concepts from the perspective of cognitive semantics. Mark made a list of the basics. In addition to Truth, we looked in detail at Time, Causation, Events, The Mind, The Self, Morality and Being. Luckily, a fair amount of work had already been done on these within cognitive semantics. We pulled the results together, unified them, and worked out further details. Not surprisingly all of these abstract concepts turned out to be mostly metaphorical, using multiple metaphors, each with a different logic. Thus, there is not one concept of causation, but around 20, each metaphorical and each with different inference patterns. Thus, causes can be links, paths, sources, forces, correlations, essences, and so on. Pick a metaphor for causation and different inferences come with the metaphor.

The science and the social sciences all use causal theories, but the metaphors for causation can vary widely and thus so can the kinds of causal inferences you can draw. Again, there is nothing wrong with this. You just have to realize that causation is not just one thing. There are many kinds of modes of causation, each with different logical inferences, that physical, social, and cognitive scientists attribute to reality using different metaphors for causation. Again, it is important to know which metaphor for causation you are using. Science cannot be done without metaphors of all sorts, starting with a choice of metaphors for causation. Most interestingly, if you look at the history of philosophy, you will find a considerable number of "theories of causation." When we looked closely at the philosophical theories of causation over the centuries, they all turned out to be one or another of our commonplace metaphors for causation. What philosophers have done is to pick their favorite metaphor for causation and put it forth as an eternal truth.

JB: Where does morality come into all this?

LAKOFF: One of the most satisfying set of results is the collection of metaphors governing moral thought. We found that they all seem to arise naturally in an embodied way from forms of well being—health, wealth, uprightness, light, wholeness, cleanliness, and so on. A particularly interesting result is that moral systems as a whole seem to organized metaphorically around alternative models of the family. Again, this should not be surprising, since it is in our families that we learn what we take as moral behavior.

We are now in a position to study the metaphorical structure of various moral systems. We think that cognitive science allows one to give much more detailed and insightful analyses of metaphorical systems than has ever been available before. For example, in our study of Kant's moral theory, we argue that this great intellectual edifice arose from just four basic metaphors, and that this allows us to see just how the various aspects of Kant's moral theory fits together.

Cognitive Science not only sheds light on the conceptual structure of moral systems, but also on politics and social issues. Some colleagues and I are now in the process of forming a political think tank to apply these methods of cognitive analysis to everyday political and social issues.

Perhaps the most sobering result is the most fundamental. We are neural beings. Our brains take their input from the rest of our bodies. What our bodies are like and how they function in the world thus structures the very concepts we can use to think. We cannot think just anything—only what our embodied brains permit.

Metaphor appears to be a neural mechanism that allows us to adapt the neural systems used in sensory-motor activity to create forms of abstract reason. If this is correct, as it seems to be, our sensory-motor systems thus limit the abstract reasoning that we can perform. Anything we can think or understand is shaped by, made possible by, and limited by our bodies, brains, and our embodied interactions in the world. This is what we have to theorize with. Is it adequate to understand the world scientifically?

There is reason to think that our embodied conceptual resources may not be adequate to all the tasks of science. We take case studies from physics and discuss them in our sections on Time and Causation. General relativity is a good example.

JB: So, what's the big change here?

LAKOFF: In characterizing space-time, Einstein, like Newton before him, used the common metaphor that time is a spatial dimension. My present time and location is metaphorically conceptualized as a point in a four-dimensional space, with the present as a point on the time axis. In order for there to be curvature in space time, the time axis must be extended—it cannot be just one point, the present. In addition to the present, the time axis must include portions of the time axis understood as future and past if there is to be enough of the time axis to form a curved space time. This seems to imply, as philosophers have repeatedly observed, that at least portions of the future and past coexist with present. And if the future exists at present, then the universe is deterministic. Frankly, it seems nutty to say that the past, present, and future are coexistent—and yet the curvature of space-time seems to imply it.

JB : Does the problem lie with the physical theory or the mathematics used to express it?

LAKOFF: It lies with the common metaphor "Time Is A Spatial Dimension," which is used to understand Einstein's mathematical theory of the physical universe. The philosophical entailment of determinism is coming not out of the mathematical physics, but out of that metaphor applied to the mathematical physics. Does that mean that we should—or can—try to jettison the metaphor?

For better or worse, we cannot get rid of it—even if it does have a nutty entailment. Physics is about something. We need to link the mathematics of relativity to an understanding of space and time. "The Time Is A Spatial Dimension metaphor does that job." We have no better metaphor and no literal concept arising from our embodied minds to replace it with. The commonplace metaphor may be imperfect in having a nutty entailment, but it's the best that embodied human conceptual systems are likely to come up with. What this means is that it is important to separate the mathematical physics from the commonplace metaphors used to comprehend it. And it is vitally important not to take those metaphors literally, even if that leaves us with no literal understanding at all. We should not take time literally to be a spatial dimension; we should recognize that we are using a common metaphor, and that the metaphor has the unwanted baggage of determinism—the entailment that present, past, and future coexist.

The moral is that you cannot take conceptual systems for granted. They are neither transparent nor simple nor fully literal. From the perspective of the science of mind, science itself looks very different from what we are commonly taught it is. Scientific understanding, like all human understanding, must make use of a conceptual system shaped by our brains and bodies.

##