

IRENE HEIM: I'm Irene Heim, and I will be moderating the language and thought panel. So before I introduce my illustrious panelists, let me just say some general words about the topic.

The capacity for language is obviously one of the very most striking manifestations of human intelligence. Language is not only our most powerful tool for sharing our thoughts and feelings with others, it is also a most essential medium for thought itself. Not all thinking and reasoning is verbal thinking and reasoning, but a very large amount of it is. Part of what we're going to talk a little bit about today is just how much is language, and how it interfaces with other cognitive systems.

Contemporary psychologists and linguists, like other cognitive scientists, aim to discover the mental data structures and algorithms that are involved in cognitive capacities. In this case, the structures and algorithms that are involved in thinking and reasoning, in understanding speech, producing speech, as well as those involved in acquiring fluency in a natural language in the first place. So this is a vast and very complex subject matter.

Over the last few decades, there has been a lot of progress which also has led to specialized fields that deal with different aspects of this. If you consider the process by which a stream of speech sounds that hits your ears eventually gives rise to a mental image or sets off a chain of deduction, this is completely spontaneous and instantaneous, and yet this process breaks down into a complex and long sequence of very different sub-routines, from auditory perception and categorization to inferring various kinds of hierarchical symbolic structures.

So accordingly, in these fields that we are representing here today-- phoneticians, phonologists, morphologists, syntacticians, semanticists-- all of whom study specific properties of these heterogeneous processes, yet study them as a system that hangs together.

From a developmental perspective, also language acquisition by young children unfolds gradually, over a number of years, and specialists study, among other things, the core concepts that are available in pre-linguistics infants, the ways these get mapped to words in the acquisition of vocabulary, and then also the time course of relatively late developing aspects of language use and reasoning.

So in the brief segment of this symposium that we have devoted to this vast complex of the study of language and thought, of course we will only be able to spotlight very small parcels. Particularly, they will be on the semantics side of the field-- the interface of language and thought.

I would like to introduce my speakers in the order in which they will be speaking. The first person to speak will be Professor Patrick Winston, who is Ford Professor of Artificial Intelligence and Computer Science here at MIT. He's been at MIT for a very long time. He joined the faculty in 1970, and he was director of the MIT Artificial Intelligence Laboratory from 1972 to 1997.

His groups studies how vision, language, and motor faculties account for intelligence, integrating work from a number of allied fields, including computer science, neuroscience, and cognitive science, including linguistics. His publications include numerous books, including major textbooks on artificial intelligence and programming languages. He's also chairman and co-founder of Ascent Technology, a company that develops products which solve complex resource planning and situation assessment problems, and he served as a member and also chairman of the Naval Research Advisory Committee. He's a past president of the American Association for Artificial Intelligence, and a leader of the Human Intelligence Enterprise.

Our next speaker after Patrick will be Steven Pinker, who is Harvard College Professor and Johnston Family Professor of Psychology at Harvard. He's also taught at Stanford and for very many years here at MIT. His research on visual cognition and psychology of language has won prizes from the National Academy of Science, Royal Institution of Great Britain, the American Psychological Association, the Cognitive Neuroscience Society. He has several honorary doctorates, teaching awards, numerous prizes for his books, including *The Language Instinct*, and his most recent book, *The Stuff Of Thought-- How Language is a Window Into Human Nature*. He's been named Humanist of the Year, and many other honors.

And then, we have Professor Susan Carey, who is Henry and Elizabeth Morse Professor of Psychology, and also the chair of the Psychology Department at Harvard University, and she previously taught at NYU, but also for a very long time here at MIT. She has played a leading role in transforming our understanding of cognitive development. Her research has been on child development, the history and evolution of human cognition, and she's written a number of books, including in particular a recent book called *The Origin of Concepts*.

She's a recipient of numerous prizes, including the David Rumelhart prize, which she was the first person to receive for theoretical contributions. It was previously given only for computational modeling work. Among the many other honors she's received is the Distinguished Scientific Contribution Award from the American Psychological Association and the Nicod Prize from Paris. She's a member of many other prestigious societies.

Finally, our last speaker will be Professor Gennaro Chierchia, who is Haas Foundations Professor of Linguistics and Chair of the Department of Linguistics at Harvard. Before he came to Harvard relatively recently, he had taught at Brown University, Cornell University, and for a long time at the University of Milan. His interests are in semantics and syntax/semantics interface. He's worked on a very wide range of foundational topics in these fields-- anaphora, the structure of noun phrases, the polarity system, the theory of [INAUDIBLE]-- all from a cross-linguistic perspective.

His work is mostly comparative linguistic analysis, but he also takes part in psycholinguistics experimental research. The main focus of his research is how the spontaneous capacity to attach interpretations to sound patterns, or to a range of signs in sign language, depends on innate natural logic to sort that out from how it depends on specifically non-logical abilities.

So these are my four panelists, and I would like to have Patrick Winston start and then proceed.

**PATRICK
WINSTON:**

Many of you expressed disappointment that Marvin wasn't able to answer the last question put to him by Steve Pinker last night. Only Noam got a chance to answer. So I'd like to fix that up before I start. Marvin agrees with Noam. And as far as I know, this is the only issue on which they've ever agreed, so it was really a milestone moment and a brilliant question.

For my own presentation, I want to finish the talk I started last night. I'm going to ask myself two questions today. I'm going to address the question of what is the goal of science-side artificial intelligence, and then I'm going to address the question of how should it be done.

So take note, I said science-side AI. Those of us who do science-side AI are in a minority. We've sometimes joked about starting a special interest group in the AI Society on the subject of AI. We regard ourselves fundamentally as lunatic fringe psychologists, but we are also lunatic fringe computer scientists, too, so we are at the intersection of two kinds of lunacy, I suppose.

But I want to answer my questions before I really start on the defense of my answers. My first answer is that those of us who are interested in science-side artificial intelligence are interested in developing a computational account of intelligence that includes human intelligence. So my answer is embedded in my title. The answer to my other question is also embedded in my title, or rather the other question is embedded in my title. My answer to that is, we have to sit at a round table with people in allied fields, who we need to work with.

And now, the rest of my talk will be buttressing my position on those two issues. So first of all, to set the problem up, I'd like to go back 40 years to the early 1970s. That was a time when, at MIT, many of us were graduate students who felt that we were working on intelligence, independent of whether it was human or machine.

There were four theses that are often regarded as remarkable. Terry Winograd implemented the program that engaged in a model-back dialogue about things moving around in the blocks world. I implemented the program that learned from a few samples what an arch is in the blocks world. Dave Waltz implemented the program that understood line drawings, a model-back program in the blocks world. And Jerry Sussman implemented the program that learned how to manipulate blocks. It was a model-back program in the blocks world.

And you'll notice a couple of common themes. These were in the blocks world, and they were model-backed. These things proved to be remarkably easy, and within the scope of single PhD theses. It was also a time when Roger Schank was developing story understanding programs, and that turned out to be a lot harder, much beyond the reach of a single PhD thesis. And in fact, after Shank's pioneering research on story understanding proceeded for about a decade, it faded. It didn't disappear, but it was no longer the sort of thing that you could consider to be at the center of artificial intelligence research.

So why was that? It might be because in the blocks world, the modeling problem is very easy. We know how to model a block. What we don't know how to model is all of the sorts of stuff that occur in the kinds of stories that we humans tell each other.

So now, I have to fast forward to today. Last night, I alluded to research that I do that's focused on story understanding, and here's half of the text of a simple rendering of the plot in Shakespeare's *Macbeth* story. There are a couple of points of interest here. First of all, if you only read this story and don't know anything, you don't know that a few people end up dead. You don't know why McDuff killed Macbeth. You don't know that this is a story about revenge, and if you're a Google search engine you won't find that if you search on Shakespeare and revenge.

But we humans read this story and we say, of course it's a story about revenge. And how do we do that? Well, that's the subject of the one hour version of this talk, but I'd like to not leave the subject before I show you the answers provided by the kinds of systems that I built, along with my students.

So suppose you ask the question, why did McDuff kill Macbeth? Here is the answer on two levels. So on a kind of commonsense level, there are two answers, one by Doctor Jekyll and the other by Mr. Hyde-- pardon the cutesy names-- and on the reflective level, there are also two answers. So Doctor Jekyll is taking a kind of dispositional point of view. It must have happened because of some evil intent inside of McDuff. Whereas Mr. Hyde is taking a more situational point of view, which says that McDuff must have been embedded in some kind of situation that made him do it. So believe it or not, this is a blocks world first of a program that reads the same story from the perspective of two different and separate persona, one reflecting a more Western one point of view and one reflecting a more Eastern point of view.

So we're beginning to have the sorts of things we wished we had in 1970-- story understanding systems that can read stories from multiple points of view, from the point of view you might think of as coming from multiple kinds of experiences in the background of the program. Well, I think it's pretty cool, and as is usual in scientific presentations I would like to thank my collaborators. Here are my collaborators.

Now there are two interesting things about this list of collaborators. First of all, Liz Spelke is a developmental psychologist, Ian Tattersall is a paleoanthropologist, and Ray Jackendoff is a linguist. That's the first point of interest. The second point of interest is that none of these people know they've been collaborating with me. And in fact, it might be worse than that. Some of them may be either embarrassed or angry to be on this list. But the fact is, we couldn't have done what we've done without the inspiration of the work of these people. I'd like to explain that in the next three slides.

First of all, what Liz has done, along with her students and other collaborators outside of our group-- it's the wonderful blue wall experiment. There's a rectangular room, a blue wall, a turntable of some sort in the middle, and initially, let's put a rat on the turntable. We disorient it-- it's already watched us put food in one corner-- then we disorient it, spin it around, and to our amazement, the rat is semi-intelligent. It goes to the diagonally opposite corners, which would be correct if there weren't a blue wall. It knows about the geometry of the room.

Amazingly, that's also true of a small child. You have to get to the level of a human adult before you get the correct answer. That's me. This can be done by people who are not university professors. In fact, the interesting question is, when does a small child become an adult? And the answer is-- the answer I found to be astonishing is-- that it doesn't happen until the child is about five or six, and it doesn't happen until the child has begun to use the words left and right in their productive vocabulary. They understand what left and right mean before that, but they don't use those words until approximately the same time that they can break the asymmetry.

Now the longer version of this talk talks about how you can reduce the human adult to a level of a rat by asking him to-- without drugs-- but I'll leave that aside for another time. So what did I learn from Spelke? I learned what Spelke tells me, that we human adults have a kind of combinator capability that allows us to tie together information that's coming from the blueness of the wall along with information that's coming from the geometry of the room. So it's the combinator capacity that I find extraordinarily influential.

Now what about Tattersall? I droned on about this a bit last night, so I won't continue, I won't drone on about it now. I'll merely say that it might be that that magic quality that made us different from the Neanderthals and all other primates today is that combinator capability, because that combinator capability allows description, and once you've got description, there's a trail that can lead you all the way up to what I call the strong story hypothesis. This is one of the fundamental things that makes us humans different from the chimpanzees and Neanderthals.

Now if we believe that, being a lunatic fringe computer scientist, I have to implement one. As Tomaso said, you don't really understand it until you can build it. You can build it before you understand it, but being able to build it is a necessary-- if not a sufficient-- condition. So we have to try to build it, and how do we build it?

Well now, we borrow information from our friends in linguistics. This is page 168 of Ray Jackendoff's *Semantics and Cognition*. And you don't have to read it, because what it's about, it's about trajectory, and he discusses how much trajectory is embedded in our everyday language. We're constantly talking about things moving along paths, sometimes physical things moving along physical paths, sometimes abstract things moving along abstract paths. If I say the president wanted Iraq to move toward democracy, that's movement still, it's just in an abstract space.

So that's the story. Our story understanding capability, our story understanding construction, has been heavily influenced and guided by help from our friends in psychology, paleoanthropology, and linguistics. So what I argue is that we need to sit at a round table, because we borrowed stuff from them. But not just that, it has to go the other way, and here are some of the questions that I would like to have answered that my kind of technology doesn't answer.

For example, if you tell a story you're talking about a sequence. Where does that sequence get processed? Is it up here, in this gigantic piece of human brain that doesn't exist in such a gigantic form in the Neanderthals? I wish Matt Wilson could stick 100 probes into a human hippocampus to see if our hippocampus is engaged in understanding a story, or alternatively if I could only tell a story to a rat, I guess I could do it with that.

But the fact is, we don't know anything about what goes on in here, and we don't know all sorts of answers to all sorts of questions having to do with sequence, with story, what happens when we try to understand a hard story, what happens when I study a story so that I can retrieve it and recall it later on, what happens when I do retrieve it. Why is it the case that I can think that a country can be a Cinderella, but I might have difficulty thinking about an old man being a Cinderella?

These kinds of things are puzzles that I can pose, and which come up in the course of my attempts to build computational models and computational systems that deal with and understand stories, but they require a kind of dialogue in order to come to grips with them.

So that concludes what I wanted to say, and these are the things that I hope I contributed not just to this particular seminar, but to the symposium as a whole. There is a kind of AI that focuses on the science-side, on the questions of what it is of building computational models of human intelligence, not on the engineering side of building artifacts like Watson and Deep Blue, that do their jobs by whatever means and do them so extraordinarily well.

The second point that I want to contribute is the idea that story understanding is very fundamental to any real understanding of human intelligence. Watson doesn't understand stories. Watson doesn't have a model of itself. Watson can't blend stories together. These are the kinds of questions that we have to answer if we are going to have a model of what goes on inside our own heads. And finally, the round table story-- it's not just about exchanging technology, it's also, I think, about exchanging questions.

IRENE HEIM: Thank you, Patrick. Steve Pinker will be our next speaker.

**STEVEN
PINKER:** Is this the right clicker? What should I be aiming at? Let's try this. No luck? Oh, very good. Okay. Okay.

I'm going to share some thoughts on thought and language by explaining an idea that I have worked intensely on for the last couple of decades, which I'll call the conceptual structure hypothesis. This has been inspired by the work of many linguists, the most prominent of which are Ray Jackendoff and Beth Levine, and that I've worked out in my books *Learnability and Cognition* and *The Stuff of Thought*.

According to the conceptual structure hypothesis, the human mind contains a level of representation with the following properties. It's categorical, it's structured. This level of representation underlies the meaning of words and sentences as the speaker understands them. It affects certain syntactic phenomena, in particular argument structure alternations. It also underlies many high-level cognitive processes, including many forms of inference, our intuitive physics, the argumentation that takes place in political and moral arenas, and our shared social and moral norms.

I'll give you a taste of this idea by walking through three examples of conceptual structures: space, time, and causation. The representation of space in language has a number of distinctive, and perhaps even peculiar, properties. One is that even though space of course is three dimensional and analogue, language treats it as digital. Language tends to make categorical distinctions like near versus far, on versus off, in versus out.

Also, language treats shape schematically. Even though, in reality, all objects are complex three dimensional arrangements of matter, language schematizes them in terms of parts and boundaries that are extended in zero, one, two, or three dimensions. Similar in a way to David Marr's theory of generalized cones, and Irv Biedermann's geons.

For example, we have the word "line" that refers to a one-dimensional entity, but also the word "road" that refers to an entity that we think of as fundamentally one-dimensional, though we also recognize that it has a finite width, and "beam", another word that we think of as fundamentally one-dimensional, though we grant it a finite thickness. We have the word "surface" for a two-dimensional object, but also the word "slab" for an object that we think of as fundamentally two-dimensional while recognizing that it has a non-zero thickness.

This idealized geometry governs the way we choose our prepositions. For example, the preposition "along" requires an object that we think of as fundamentally one-dimensional. You can say the ant walked along the line, or along the road, or along the beam, but it sounds a bit odd to say the ant walked along the plate, or walked along the ball.

This geometry is what goes into our choice of nouns to pertain to shapes. Even though a geometer would say that a wire is just a long skinny cylinder, it sounds odd to use the word cylinder referring to a wire because we treat its thickness as if it didn't exist. Likewise, we wouldn't use the word cylinder to apply to a CD, even though technically it's just a short, fat cylinder, because at this level we conceptualize its thickness as not existing.

I think it goes into our general sense of shape. What objects remind us of what other objects, as in the little girl who said to her father, "I don't want the little crayon box, I want the box that looks like an audience." That is, not the flat, eight-crayon Crayola box, but the 64-crayon box where the crayons are arranged in stepped rows like the tiers of an auditorium.

Well, why do we represent space this way? I don't think it's because of any demand of language per se, but rather the geometry of spatial language divides up the world into regions that tend to have different causal consequences, and I'll give one anecdote to illustrate what I mean by that.

This is a story that I clipped out of the *Boston Globe* a couple of years ago, which reported that a woman who fell through thin ice Sunday and was underwater for 90 minutes died yesterday. The Lincoln Fire Department said a miscommunication between the caller who reported the accident and the dispatcher significantly delayed her rescue. The rescue workers believe that a woman had fallen on the ice, not through it, and that left the rescuers combing the woods to find the scene of the accident. Now in terms of analog coordinates, the difference between on and through is a matter of a few feet, but it's obviously a few feet that had important-- indeed literally-- life and death consequences.

Let me switch to causality in language. The model of causality in language can be summarized more or less in this diagram. That is, an actor directly impinges on an entity, making it move or change. And this is a model of causality that enters into causative constructions in a wide variety of languages.

My favorite illustration of the phenomenon comes from a set of experiments by Philip Wolfe, who showed subjects animations of Sarah causing the door to open, either by pushing the door open directly or by opening a window, whereupon a breeze came in and blew the door open. If you describe this periphrastically and asked people, "Did Sarah cause the door to open?", in both cases they say yes. But if you now pack the causality inside the verb, and ask people did Sarah open the door, then when she physically manhandled the door they say yes but when she initiated the first link in a causal chain that results in the door opening, they say no.

Why do we represent causality in this way? I think it's because it's the directly caused events that are most likely to be foreseeable and intended, hence those for which we feel we can hold people responsible. And when a causal chain has multiple links, our sense of moral and legal responsibility breaks down. I'll illustrate this with an anecdote from American history.

In 1881, President James Garfield was felled by two bullets fired by an assassin while waiting to board a train in a Baltimore station. The bullets missed his major organs and arteries, and needn't have been fatal even in Garfield's time, but unfortunately he was handed over to the best physicians of the day, who subjected him to the harebrained medical practices of the day. For example, they probed his wounds without washing their hands. They decided that it would be a good idea to feed him through his rectum instead of through his mouth. Well, Garfield lingered on his death bed for three months before finally succumbing to the effects of infection and starvation. At the murder trial the assassin, Charles Guiteau, said, "The doctors killed him. I just shot him." The jury was unpersuaded, and he went to the gallows, another illustration of the consequential implications of the mental models underlying verbs.

Finally, let me just say a couple of words about the treatment of time in language. I'm relying on an analysis by Bernard Comrie. Like the treatment of space in language, time, which of course is essentially analog, is digitized. That is, we make categorical distinctions in the grammar of language such as past, present, and future.

What is the coordinate system by which we reckon these intervals in time? Well, they are defined relative to the moment of speaking. Languages typically-- if they have tense systems-- divide the timeline into two to five regions, defined relative to the moment of speaking. One of the origins of this coordinate system is an entity that William James called the specious present, a sliding interval of about three seconds that corresponds to our general sense of now-ness. The neurophysiologist Ernst [INAUDIBLE] has proposed a law that psychological events as they're conceived by the experiencer fall into a three second window.

Turns out that three seconds is the duration of a deliberate action, like a handshake or a hug. It's the duration of a quick decision, how long you stay on a channel while channel surfing, the decay of unrehearsed working memory or short term memory, the duration of a line of poetry in all the world's languages, and the duration of a musical motif, such as the opening notes of Beethoven's Fifth Symphony. Which people don't experience as, well, there goes a note, oh there's another note, there's yet another note, but rather, we take it in as a gestalt, and there's a sense in which we experience it all at the same time. And in fact, philosophers who study consciousness, when asked to define it, often say well, it's first-person present tense experience. The fact that they use the vocabulary of grammar to talk about this psychological entity, I think, is informative.

Of course, we also have other intervals that are picked out by the tense system. The past tense, for example, refers to everything that happened from about three seconds ago stretching back to the Big Bang, which is why Groucho Marx, in one of his movies, said to a dinner hostess, "I've had a wonderful evening, but this wasn't it."

Then we have this other interval picked out by tense systems. Basically anything from about three seconds of now until the heat death of the universe is given the same timestamp, as far as tense concerned. I won't be able to provide a full answer to why the human mind represents time this way for the purpose of expressing it in tense. The bottom line is that stretches of time that are defined relative to the moment of speaking have different consequences for knowledge and action. They correspond to what is knowable, factual, and willable.

But I'll just sum up some of the implications for human affairs by asking one question. Is the representation of time in language really relevant to human concerns? And the answer is, it depends on what the meaning of "is" is. So to sum up, conceptual structure is a categorical structured representation at the interface of language and thought, and I've given some examples. Space is represented as a set of objects extended in zero to three dimensions, located relative to a place. Causation is thought of as the direct impingement of an actor upon an entity, and time is conceived as activities and events extended in one dimension, located upon a time line. Thanks very much.

IRENE HEIM: Thank you, Steve. So we'll have Susan Carey next.

SUSAN CAREY: The project I'm interested in is, how is it we're the only animals who can think thoughts formulated over concepts like cancer, or infinity, or wisdom, or carburetor? What makes this possible? What makes the acquisition of concepts possible? This question is complicated by the fact that it's also often very hard.

For example, one half of American high school students do not have a concept of a fraction. That's what divides over 500 in the SATs and under 500 on the SATs. 80% of students who take a year course in physics in high school or college-- not at MIT-- don't have the Newtonian concept of force at the end of that year studying physics. So there's really two questions that we have to consider together. What makes concept acquisition-- conceptual development which is really the human conceptual repertoire-- possible, and what makes it so hard?

Now, as a matter of logic, a theory of conceptual development has to have three parts. We have to be able to specify what the innate representational resources are, we need to be able to specify what the adult conceptual repertoire is, and we have to describe the difference between the two so we know what kind of problem we're facing. And what it is to specify a representation resource, this we turn to work in AI or linguistics. We need to understand the nature of the symbols / what are their content? What do they represent? What's their format, how are they written? And what's their computational role? That's what a theory of conceptual representations, at any level, will look like.

And with respect to comparing the innate representational system and the adult one, just at a descriptive level, we want to know whether there are discontinuities, important qualitative differences between the infant innate repertoire and the adult one, because that will be relevant to understanding why conceptual development is hard. But even more important, it tells us whether we have an easy learning problem or a hard one.

Okay. Now, what I'd like to do is illustrate the answers to this sort of abstract question with a single case example-- a case study. And I'm going to take a case study of a concept of an integer. So both the concept of an integer abstractly, but also the representation of specific integers like seven, or 423, or $10^{29} + 3$. The famous constructivist mathematician Kronecker famously said, "The integers were created by God. All else is man-made."

Now, he was making a metaphysical claim, but you can turn this into an epistemological one and say that Kronecker's hypothesis was that natural selection gave us the integers, and the rest of mathematics is a construction. Now, if natural selection were nice enough to do that, mathematicians have told us how we can construct much of mathematics out of the integers.

But I'm here to tell you that Kronecker was wrong. Natural selection did not give us the integers. To give you a little feel for the evidence that that's true, let me just mention several facts. One is that young children learn lexical words for integers before they're two years old. They learn to count. One, two, three, four, five. However, they also learn what one means very early on, so if you ask them to give you a penny they give you a penny, but it takes them nine months to learn what two means. So if you asked such a child to give you two pennies, they pick up a handful of pennies and they hand it to you and say, here's two.

Not only that, adults in societies that don't have an external representational system for integers, like a counting routine, cannot construct sets that match each other in exact number. There have been recent modern studies of Amazonian Indians for which this is true. And it's not just a matter of culture that they don't need these concepts, because deaf people who don't have access to a sign language, and so they construct a system of representation called home sign, use their fingers to represent number, but they use it to represent a proximate number. They are exactly like the Piraha and Mundurucu. They cannot-- if I tap you seven times and ask you to do the same to me, they cannot do that. They cannot construct a representation of a set of exactly seven.

So that's the kind of evidence that we don't get representation of exact cardinal values for free. And I also want to tell you that this is very important. I mean, this is a deceptively simple case study to illustrate my scientific points, but the single most important predictor for success in elementary school is whether children have constructed a counting algorithm and know how it represents number in preschool. That is what separates success and failure in elementary school, across the board.

So it's hard, and it's important. So what makes it possible, and why is it hard? Well, what makes it possible is, in fact, there's a tremendous amount of innate support for number representations. We have rich, innate representations for number, they're just not integers. So one system of representation takes sets as an input, and creates an analogue model of the approximate cardinal value of a set.

That's what the Piraha and Mundurucu and homesigners can do. If I show you 10 objects and ask them how many, they put up roughly 10 fingers, so they're approximating. It's clearly a representation of cardinal value, but it isn't integers because it's not based on the successor function.

Another system of representation with numerical content is we create working memory models consisting of one symbol in our memory for each individual in a set, and we can compare those sets on the basis of one-to-one correspondence. There's a lot of implicit numerical content in that, but we can only represent sets of one, two, and three with that system of representation.

And then the conceptual structures that Steve was talking about, which takes set representations as their input and express aspects of number in language, underlie the quantifiers of natural language, singular or plural distinction, determiners-- we have a lot of numerical information at the interface between language and thought. It's just that none of those systems, at their outset, have any symbols for any integers over three or four.

So that's the discontinuity, right? We, as adults, once we've mastered a numeral list representation of number, or scientific notation-- there's lots of different notation systems that we have for number-- we have a system of representation with more expressive power than any of its input. So the question is, how does learning achieve this at two timescales, both in cultural history and in ontogenesis?

And I don't have time to give you the full answer to that question, but there are many learning mechanisms that underlie discontinuities such as that. But it's the essence of human intelligence that we culturally create representational systems with more expressive power than, or that are incommensurable with, the initial innate representational resources. And then each child in ontogenesis has to master how those representational systems work.

And one mechanism that underlies this is a kind of bootstrapping process that's been described in the philosophy of science, in which we use language to formulate hypotheses in which their content is given by the interrelations of the concepts between themselves directly. And then we use these to model the world that we represent, in terms of our innate representational resources, and that's a mechanism that allows us to create representational resources that greatly increase our representational power.

So I've exceeded my 10 minutes, and what I've tried to show you here is that cognitive science has the theoretical and methodological tools to answer both questions that I've repeated several times, which is what makes conceptual development possible and what makes it so hard. I just want to end with the observation that probably the most important seminal work on this classical set of issues has happened here at MIT, both in the Linguistics Department, starting with Chomsky, and AI Department and the Brain and Cognitive Science Department. Thank you.

IRENE HEIM: Thank you Susan. So we'll have our last speaker, Gennaro Chierchia.

**GENNARO
CHIERCHIA:** I'm waiting for something to happen. I never used technology, for good reasons. Hello? All right, there we go.

At any rate, what I have to tell you won't be the thrill of your life, but it's one example that I would like to discuss and see how it relates to the issue of language, understanding, and intelligence.

So, "There are any cookies left." A bad sentence. Each word is fine, but that order doesn't work. "I believe that there are any cookies left." Still bad, no? And surprising. You take the sentence, you embedded it, still bad. "I doubt that there are any cookies left." Now, it becomes perfect. The same sequence becomes perfect. Conditionals are a lot of fun.

If you put the same tense, the single words in the consequent, "If we go home, there are any cookies left," it doesn't work. But, if you put it in the antecedent, "If there are any cookies left, we'll be fine," it becomes perfect. Right?

So the first thing that you want to ask is, what is it that doubt and the antecedent of conditionals have in common, as opposed to belief and the consequent of conditionals? And as it turns out, there's a rather abstract logical property having to do with inference. And in particular, in the consequent you get superset inference. That is to say, going from small sets to large ones. And the antecedent is the other way around.

So for example, if it's true that if I am nervous-- like say now-- I eat pizza with anchovies, it has got to be true that if I'm nervous, I eat pizza. So the inference goes from anchovy pizza eaters to pizza eaters. But in this spot, it's the other way around. So if it's true that if I eat pizza I get sick, it's got to be true that if I eat pizza with anchovies, I get sick.

And doubt has the same characteristics. So it reverses entailment pattern, and the distribution of words like "any" or "ever" is sensitive to this, is sensitive to this abstract, inference-related property. And every language has words like this.

And so, on to the next question. These things are called negative polarity items. And why, what are they? Why does every language bother? Let me sketch a quick story. Let's assume that any is roughly the same thing as some. Now, this makes entomological sense. Any derives from the proto-Indo-European word for one, plus an adjectival ending, so it means one-y, one-like. And the cognate in German is "einig," which is a plain existential quantifier.

It also makes sense from an interpretive point of view, so something like "I doubt that John saw anyone" roughly means the same as "I doubt that John saw someone." I do not believe that there is something, in a certain domain, that John saw.

And then, let me give you a quick story on how the distribution of this element might work. This is the part-- this is the sort of most technical part of [INAUDIBLE]. I might lose those of you that don't know elementary logic. Just bear with me for one second.

So suppose that this is a story-- whoops. One second. It's not time for Marsha yet. So imagine that a sentence containing "any" is associated with an operator that is defined just in case this sentence entails this sentence, where the only thing that you change is a quantification of domain, and you make it smaller. So in some sense, "any" is associated with a certain domain, and its alternative is smaller domains. Whenever that is entailed, the sentence is okay. And whenever it is not, it is not okay.

So for example, "I doubt that there is someone in Boston that John saw" entails "I doubt that there is someone at MIT that John saw." And so the entailment goes from the large set to the small set, and "any" is fine. When instead the entailment goes from the small set to the large set, "any" is not fine. This is the idea. I'm simply re-couching what I just said in formal terms.

Now, this is the next thing to this interesting step, and this data comes from an MIT thesis from 30 years ago. It's a nice round number, and it's one fifth of the age that we're celebrating, and it's also roughly my generation. So these are the facts. These are cases where what I just said looks more elusive.

The first sentence seems to be OK. "Dogs do not have ears because they have anything we don't. They have ears because they have ears." Good sentence, fine. Notice that its form is of this kind. It is not the case that P because Q, and the "any" is inside here. You see that?

This sentence is isomorphic to that, but it's bad. "Dogs don't have ears because they have any eyes. They have ears because they have ears." Don't you want to get rid of the second "any" there? "Dogs don't have ears because they have eyes. They have ears because they have ears." So "any" is good here, and bad here, in the same spot. Why? Hopeless. Let's change jobs. Let's look into something else.

And this is pretty regular. "I didn't help him because I have any sympathy for urban guerrillas." Fine. "I doubt that grass is green because it has any chlorophyll." Not so good. What is going on?

Here is what is going on. Take this sentence. "Dogs don't have ears because they have anything we don't." Now here, we are presuming that this sentence is false, right? We are presuming that it's false that dogs have something we don't. So this sentence suggests, implicates, presupposes, that the cause is false. This sentence, on the other hand-- "Dogs don't have ears because they have any eyes"-- here, the cause is true. Dogs do have eyes. You get it?

And so, I think that this is what is going on. "Any" is not just sensitive to the literal meaning, but is sensitive to this other implicated, presupposed meaning that is flimsy. So in particular, if there is a positive implicature of sorts, "any" is disrupted. "Any" wants subset licensing inferences, and if there is a positive suggestion around, that destroys the logical environment that "any" needs.

Well, this is working it out. I will maybe spare you that, because I'm already past my time, and jump to the main morals. So the first and foremost main moral is that grammar is a logical system through and through. In fact, even at the interface with pragmatics. We have been hearing that the main sort of building block of language is a capacity for structure formation, sort of a computational facility for building structures. Well, that is a necessary condition, but is not sufficient. What you really, really need is a logic, which isn't the same thing as a recursive device. And that is necessary, in my opinion, also to make the conceptual systems that Stephen, for instance, was talking about, not blind. To use a conceptual system, you need to run through a logic.

Now, what we've been painstakingly doing is figuring out the blueprint of this logic. And you know, this means, what operators are there? How do they interact? Some of them will be familiar from classical logics, and some of them will not. The operator that I've posited, that O that I had there, is in fact "only." That's what it means. And this is something that, again, every language has. Just like every language has "even." So these little words play an enormous role and develop in extremely complex structures. Natural simply means that it grows spontaneously in us, so there is a rich system of operators that occurs across languages, with certain points of variations that constitutes this natural logic, and we just need to keep going.

IRENE HEIM: Thank you. I would like to start with asking a few questions to the panelists, and then we'll engage the audience. I'd actually like to start with a question for Susan, but also perhaps it's for Gennaro also.

So in light of what Gennaro has sketched a little bit about the many levels of meaning, of the literal meaning and this enriched meaning-- and the literal meaning, of course, is sort of computed from a syntactic structure that is not completely what meets the eye. So I'm wondering how this affects your research, for example, when you work with young children and test their understanding of number boards. In these tasks, you have to give them full sentences. You have to ask them things like, "Can you give me two cookies," or something. So I just would like to know how you think it's possible-- if you want to give us a detailed linguistic analysis of what goes into "Give me two cookies," that perhaps breaks it down into an existential quantification and a predicate of sets. And possibly-- we don't know for sure, as Gennaro has shown us, what part is literal meaning and what part might be some enriched meaning. So how do you negotiate this issues in your research?

SUSAN CAREY: So there's really two issues here. I completely agree with Gennaro and Steve about the importance of structured representations, and also in the role of logic. When you try to characterize what the meaning is for some other creature, like a two-year-old, or in the history of science, what did force mean for Galileo, you have no choice but to look at the use of that term in a wide range of contexts. And it's always hard to know what's underlying that should be drawn from the meaning of that term, as opposed to broader systems pragmatic, et cetera.

And for the purpose of what I'm saying about the concept of number, it doesn't necessarily matter. Because two might mean at least two for a one-knower, right? And so, you wouldn't say that they don't in fact know what two means. But, they don't know what three means, because they don't have "at least three."

You still need the concept that's embedded in that "at least two" representation, so although I think the problem is extremely interesting and important to separate the level of representation that you might want to call meaning and the level of representation that you want to put in larger computational role or pragmatics, for some projects-- which is, where do you get representations with that content at all-- that's neutral with respect to that question.

IRENE HEIM: Any other panelists want to say something to that? Do you want to say something to that, Gennaro?

**GENNARO
CHIERCHIA:** No, I don't think-- when it comes to how numbers are acquired, Sue is magistra suma. You know, what can you say? When it comes to the adult system, we can develop arguments.

And so in particular, something that is now very clear-- and in fact, even maybe experimentally-- is that the interpretation of numbers is sensitive to the very same structures that I discussed in connection with "any." So in particular, take the following simple contrast. How much will you get for shoveling that driveway? Answer A. If we're in luck, we'll get \$20. There you interpret 20 [INAUDIBLE] exactly. 20, and not significantly more.

Same scenario, slightly different answer. If we get \$20, we are in luck. Now here, you are not saying if you get 20 and no more you are in luck. You're interpreting their 20 as at least 20. And notice that this scenario is exactly the same. The [INAUDIBLE] knowledge of the [INAUDIBLE] is exactly the same. The words are exactly the same. So this shift in preferred interpretation-- which can be experimentally assessed-- simply seems to stem from whether you put it in the consequent or in the antecedent of a conditional. Yet another illustration of how we work and think in terms of this spontaneous logic, without any kind of formal training.

IRENE HEIM: I'd like to ask a question for Patrick. The linguistic semantics that you have used in the things that you were talking about-- it's lexical semantics, and also the things that Steve was talking about, so how you represent verbs, and the concepts of locational changes both in physical space and metaphorically, and things like causation-- I'm wondering if you see any role for the use of more logical part of the vocabulary. Whether you think that really needs to be modelled in a precise way to make sense of people understanding stories, and what prospects you perhaps see there for more collaboration with linguistics.

**PATRICK
WINSTON:** Yes. Would you like a more-- well, let me actually reinterpret the question. That's a technique I learned in humanities classes when I was an undergraduate. I like to reinterpret that question as, what can we learn from the first two sessions-- this morning's sessions-- about how we humans can make ourselves smarter? And the answer is, we can learn a lot because what we've seen in both sessions is that we have tremendous problem solving power in our visual and linguistic systems. So if we want to make ourselves smarter, we have to engage them.

And one way to engage them-- one way to engage our visual system-- is to draw pictures, and one way we can engage our language system is to talk. So if we want to make ourselves smarter-- if you want to make yourself smarter-- draw pictures and talk. And you can't do that when you're answering your email, you can only do that when you're taking notes with a pencil, and drawing, and engaging your language system.

So much of the technology put to use in education has been put to use in ways that merely pass information past us faster, not in ways that better engage our linguistic system and our visual system. So in the future, I think that education will have to benefit from the kind of stuff we've been talking about this morning. The miracle of our human visual and language processing systems, which have problem-solving power and reasoning power and thinking power far beyond anything we've been able to even conceive in applications so far.

IRENE HEIM: So one more question. This is a little bit broader. Implicitly, Gennaro already spoke to that, but-- and really, it's a little bit also to what Noam Chomsky was talking about last night, about how the study of language and linguistics has led to an appreciation that a lot of the external things like linearizing things and so on, this is part of what we need to do for communication. But language is a vehicle of thought. It might, for example, consist of hierarchical structures that might be much more common between languages. This is an old idea, that there is a language of thought and that language learning is partly like mapping the natural languages that we learn on to the primitives of the language of thought. But I think there's perhaps another idea that's suggested by work in linguistics, that natural language is the language of thought. I think this is a question, maybe, for several of you. I saw Steve raising his eyebrows, so perhaps you might tell us what you think about this.

STEVEN PINKER: Yeah, I think it's very unlikely that any language that we speak, such as English or German or Japanese, is a language of thought. For one thing, there are a lot of experiments that show that very soon after we process a stretch of language, much of our memory for the surface form evaporates, but we do represent the gist of the passage which must be in a more abstract format than actual sentences.

And also, for a number of reasons, the actual sentences that we speak don't make the kinds of distinctions that are relevant for inference. And vice versa, there are a lot of distinctions that are relevant for inference that aren't made explicit in the actual sentences we speak. And I think that what we heard from Chomsky last night, arguing that language was evolved for thought, depends on a very idiosyncratic definition of language that basically stipulates that anything that is specific to communication, Chomsky just did not call language. He called that the mapping of language to the interfaces, such as linear order.

Noam has his reasons for defining language that way, but it's not what most people think of as language. If you met someone who had all of the hierarchical structure of English but none of the linear order, so that every pair of words in a phrase could be flipped, you'd say, does that person speak English? For one thing, you wouldn't understand most of what he or she was saying, and you'd say, I don't know what that guy is talking about, but he's certainly not speaking English.

Now, Chomsky's definition of language that we heard last night would say that yes, he is speaking English. What he just doesn't know is how to map his knowledge of English onto the articulators. Personally, I think that's an eccentric definition of language, but it's only by that definition that one could maintain that language evolved for thought. It also leaves out of the picture a lot of devices, like agreement, like case, where it's completely unclear why any reasoning system would need subject verb agreement, would have irregular past tense forms-- a whole slew of devices which are clearly there in order to help a listener recover information about who did what to whom from the surface string, but that an internal computational system could very well do without.

So for those and other reasons, I'm very skeptical of the idea that language as it's ordinarily understood, namely English sentences, is what we think in. Although I do think, as I suggested, that there is a representation that in some regards is language-like-- it makes categorical distinctions, it is hierarchically structured-- that's very closely related, though not identical, to what we think of as the semantic representations underlying sentences, and that also enters into forms of pure thought, namely inference, norms, shared understanding, and so on.

SUSAN CAREY: Another problem with identifying language with thought is that if you do so, you stipulate that other animals can't think. And maybe you want that stipulation-- I mean, certainly there are philosophers who say that, but it partly depends on what you mean by thought, as well as what you mean by language, because other animals have very abstract, powerful, abstract conceptual representations that they can draw inferences over. They have many of the same core cognition systems that human babies do. You'd have to define that as not thought if you want to make that stipulation.

IRENE HEIM: Yeah. I think it would be good to not argue about terminology in all of these instances. I think from what I hear, both Steve and you say and think that in the end-- I don't disagree with it, and I don't think, actually, Chomsky would accept for the terminological logical issues. Because indeed, let's say, never mind what we want to call language. Steve ended his contribution saying there is something that we process, each time we hear a sentence, that is a semantic representation, that is a hierarchical symbolic structure. Probably, as Gennaro says, it has a logic that allows us to draw inferences. It has elements that are like logical words that, in the way, are defined purely by their role in inferencing. It is computed online, when you speak and understand.

And then what Susan says. In addition these systems of cognition-- and Patrick also emphasized this yesterday and today-- that there's some kind of reasoning-- or if you want to call it that-- thinking, that really evokes an image, and you see something in that image. These other ways of thinking-- and sure, we want to call that thinking also. I mean, one thing that is exciting about doing linguistics and studying different languages, and precisely like getting from the differences in all these things like the word order, the agreement, patterns, and so on, the irregularities, to the structure that's computed, that underlies thinking, that that is, perhaps, really what we should think of as the language of thought, if there is such a thing. I don't know. Gennaro, Patrick, want to say something?

GENNARO
CHIERCHIA: Well, the only thing that I would like to add is that there is this traditional distinction between open-class concepts and close-class concepts. A sort of naive view of what language is about puts all of the weight on the open-class. You know, knowledge of the lexicon, in that sense. And in a way, that is not uninteresting, but it's the thing that we know the least about, and the thing that we spend the least of our time on.

We really work on the functional architecture of language, and that involves such interesting words as and, or, if, any--

SUSAN CAREY: Not.

GENNARO
CHIERCHIA: Even, only, -ing, and so on. And that is where you find what is most common throughout the languages of the world, and also where some key dimensions and variations take place. So that is what I think are the ingredients of the language of thought, ultimately.

STEVEN
PINKER: In that regard verbs are an interesting intermediate case, as Dedre Gentner argued many decades ago, because there are distinctions that are in the open-class vocabulary of verbs in some languages but are in the closed-class category in other languages. Like causative morphemes, like whether actions are distributed over--

GENNARO
CHIERCHIA: Applicatives and, yep.

STEVEN
PINKER: Or directed at a single object, and so on. So it's likely it shades off, at least in the case of verbs.

GENNARO
CHIERCHIA: That's right. So when I say that, I certainly do not mean that we know-- we go to the next language and we know what it's functional structure is. It's just a program for the search. We know where to look, and also the fact that we cannot tell doesn't mean that the distinction isn't unfounded.

SUSAN CAREY: So getting back to Irene's question, I think it's an open and a very important project, whether there are completely nonlinguistic representations with the same content and computational role as the functional vocabulary of languages. So, can animals represent a conditional? Do they have anything like the concept "not" or "or"? The evidence is probably yes, but not clear.

PATRICK
WINSTON: Well let me stick my oar in, if you don't mind. I talk to myself quite a lot, and I try not to do it out loud because it's fairly dense with expletives. In any event, I think there is an inner language, and I think that it came before our communication, partly because there's no point in saying anything if you don't have anything to say. And I think our external language emerged in part because we're social animals, and our external language is of great benefit in thinking because I think there's a great deal of creativity in creative misunderstanding of what other people say.

But I think there are computational reasons why the inner language ought to be simpler, perhaps somewhat more canonical, than our outer language, because we have to compare things and we have to match them, and it would be extraordinarily difficult from a computational point of view to try to do that with a raw, uncanonicalized, language stream. So we can't match a story in the newspaper to Cinderella and say that's a Cinderella story unless, I think, we have reduced the story and Cinderella to a kind of inner language in which there is movement through space, there is a capacity to express emotion, there's a capacity to express social relationships, and I think those kinds of things constitute the inner language. And when they're rendered in simple form, they make possible this kind of analogizing and precedent-based reasoning that is fundamental to human thought. And I'll bet chimpanzees don't do it, because I think they'd be a lot smarter if they did.

IRENE HEIM: Shall we open up for questions from the floor? I think somebody needs to bring the microphones around. Is that how it works?

AUDIENCE: Is this on? I heard the first use of the word analogy just now, in a discussion of what is effectively Jerry Fodor's question. I'd like to try to channel Douglas Hofstadter for just a moment to ask, in all of the stories, and in discretizing the continuous, and in finding an analogy between a million and one versus the same difference between zero and one, is analogy one of the fundamental logical operators of language, Gennaro, that you are trying to find? How is analogy interwoven throughout all of these?

GENNARO
CHIERCHIA: Well, I can only say that for me, it's real, but it's not the solution. It's a problem.

SUSAN CAREY: In the bootstrapping processes that I've studied in the case studies of both the history of science and in the developmental case studies, analogy is an essential part of the process, where you're using an external structure that you've symbolized in some external set of symbols to model some other domain. So, we need analogy. There's no doubt that we do, that it's an incredibly important part of creating new representational resources, which doesn't mean we understand how it works.

STEVEN PINKER: Patrick mentioned a widespread phenomenon in lexical semantics, which is that certain distinctions, like an object moving along a path to a goal, seem to have counterparts in many domains that are not literally spatial, such as in possession. Giving something to someone else, we use the proposition to, which ordinarily implies direction towards a goal. Or, the light changed from green to red.

In fact, if you look at almost any stretch of language, it's very difficult to find language that is not, in some sense, metaphorical, in that historically, and even to some extent in the consciousness of contemporary speakers. You can see a far more concrete image or schema underlying the abstract verbiage supporting a position, or Obama's poll numbers rose, and so on. Try it some time, and just about every sentence has some concrete image behind it.

That is what you find when you look at language itself, and the question is, how analogical is the human mind? That is, are some of these metaphors dead metaphors, in that they apply to some stroke of analogical insight that our linguistic ancestors may have had that then got frozen in the language, so that people don't even see through the literal verbiage to the underlying spatial model?

This is a problem that I've been quite interested in, and I've been struck by a phenomenon that's been noted independently by two computer scientists, Douglas Hofstadter and Roger Schank, that often, when one thing reminds someone of something else, often the common thread is not sensory, as in Proust's taste of the Madeleine that time-transported him back to childhood, but often highly, highly abstract. Schank gave the example of going to the barber, and he ask the barber to cut his hair short, and he was annoyed that the barber didn't cut it as short as he wanted. And he immediately thought of an episode a couple of weeks before, in which he asked his wife to cook a steak well done, and he was annoyed that she didn't cook it as well done as he wanted.

Now what's the common thread between them? Well it's not anything sensory, because well done steak and short hair don't have anything in common at the sensory level, but there is a common abstract structure in terms of a goal, telling someone else to achieve a goal, fall short of the goal, et cetera, suggesting that that is a kind of underlying matrix which the mind connects to otherwise disparate thoughts.

And I've been doing research to try to find out how commonly these analogical remindings come to people. Doug Hofstadter and Schank and I have our own diaries of analogical remindings with many dozens of examples, and we've tried to remind people of other things to see how often they are analogical, and we've been actually kind of disappointed that very few remindings that we get from other people in the lab are abstract and analogical like this. So this is a bit of a mystery why they occur so often to some cognitive scientists, they saturate our language to such an extent, but we are having trouble producing them on demand, which gives rise to a bit of a question as to how common this process really is.

AUDIENCE: In the discussion of the logic underlying language, not only in Gennaro's talk but across the board, one thing that was missing was any discussion of the interplay of proof theory versus model theory. There's a common computational conception that we have some symbolic representation system and do some proof theory on it, but also, whenever we talk about interface with vision, or interface with all these other things, the link between the linguistic-like representation and the non-language-like representations must also be very important. I'd like to hear something about that.

**GENNARO
CHIERCHIA:** A proof is essentially syntactic. A proof is a way of going from A to B. So in a way, there are-- let me answer your question in this way. There are two pieces to our semantic competence, two main pieces. One is what makes us say that "The bottle is on the floor" is true here, is true now. And there is something about the relation between my sentence and the state of the world. And this is partly of understanding, for sure, and it hinges on reference, and that is the notion that Chomsky finds very problematic. And I think he's right, it is very problematic. It is something that we know little about.

There is another type of semantic knowledge, which is that from the fact that the bottle is on the floor, I know that there is something on the floor. And that is inferential knowledge. That is knowledge about entailment, and that is different, because it's a relation between sentences and how we used them, and that can be mind internal, and sort of proof-theoretic.

**STEVEN
PINKER:** I'm not sure if the distinction between model-theoretic proof-theoretic corresponds to the distinction between an internal thought and connection to perception and action. That's probably not exactly what you had in mind, but to the extent that one might ask about how this level of conceptual structure ties us to the real world, I think it's an interesting research question in each case how these primitive symbols that seem to make so much of a difference to language, like an object that's basically extended in one versus two dimensions, or present versus non-present, map onto the other parts of the brain that interface ultimately with the senses, and that connect us to the actual world.

And there has been, recently, some work trying to do that. About 20 years ago, Ray Jackendoff and Barbara Landau speculated that the special mode of spatial thinking that seems to interface with a closed-class vocabulary such as prepositions might correspond to the dorsal visual system, particularly in the left hemisphere. That the kinds of distinctions like whether something is on, or over, or near, they speculated might be the kinds of categorical distinctions that the brain makes anyway, in trying to figure out the intuitive physics of a scene. And there has been recent work by David Kamerer and others, based on both neuroimaging and patient studies, suggesting that exactly this part of the spatial cognition system is tied to the understanding of that closed-class spatial vocabulary. But I think it asks the question of each one of these conceptual primitives, how is it connected to the senses, and hence to the world? And that's kind of an oblique, perhaps non-answer to the question of model-theoretic versus proof-theoretic semantics.

IRENE HEIM: Thank you. So I think we are out of time, so I'd like to thank my panelists.