OF MINDS AND LANGUAGE

A Dialogue with Noam Chomsky in the Basque Country

Edited by
Massimo Piattelli-Palmarini,
Juan Uriagereka, and
Pello Salaburu

OXFORD
UNIVERSITY PRESS

CHAPTER 15

Innate Learning and Beyond*

Rochel Gelman

15.1 Relevance, similarity, and attention

I usually start my presentations on this topic by asking the members o audience to participate in an experiment. I show them slides with a pair of i and ask them to rate their similarity using a scale of 1 to 10, where 1 is, *Cou be less similar*, and 10 is, *Very, very similar*. Their task is simply to call c number that reflects how similar they perceive the pair of stimuli in the sli be. A sample stimulus pair is presented in Fig. 15.1.

As expected, they normally rate the pairs as very similar, presumably bec they look very much alike on the surface. Then I inform them that the iten the slide were taken in two different places. One of the pair was taken at a and one was taken on the shelf of a store that specializes in fine ceramic co Now, with this as background information and a mindset that distingui these environments, I ask them to rate the pair of items again. This time adult audience also does as expected: they now rate the exact same pai stimuli as very dissimilar, switching from the top end of the similarity s to the bottom end of it.

Let us turn now to what 3- and 4-year-olds do when they are shown the and store pictures. When a child comes into the room, he finds the experime on her knees, surrounded by forty-two pictures, taken of twenty-one pai real and fabricated animals. She tells the child that she just dropped her pict and asks if they will help to put the zoo pictures in the zoo book, and the s pictures in the store book. The child is then given the items, one at a time. I

* Partial support for this chapter was provided by NSF ROLE Grant REC-0529579 research funds from Rutgers University.

A B

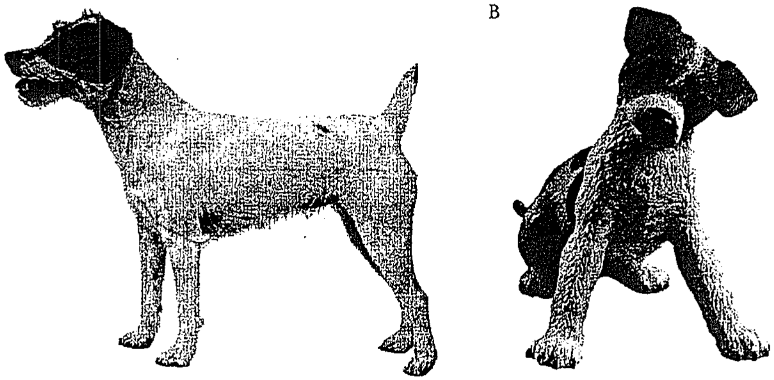

Fig. 15.1. Photographs of dogs that are similar on the surface, although one is of an animate and the other of a fabricated dog. An example of displays used in Gelman and Brenneman (2004).

age groups do this extremely well. They do not fall for the overall surface similarity as might be expected given any Piagetian, stage, or association theory about preschool competence. According to such theories, preschoolers are perception-bound. If so, our young subjects should treat pairs that are perceptually very similar on the surface as the same. Therefore their placements should be at chance. But they are not. In fact, in one such study (Gelman and Brenneman 2004), 67% and 100% of the 3- and 4-year-olds, respectively, turned in performance that met a criterion of $p < .026$. For the children to succeed on this task, they had to be able to look for details in the photographs of the live and fabricated version of the same kind that provided clues regarding their different ontological categories. But to do this, they had to have available a framework providing hints as to what constitutes relevant information for animate as opposed to inanimate objects.

Results like the above have led me to the view that there is a core domain which involves a high-level causal–conceptual distinction, one that makes principled distinctions between the nature of relevant energy sources for the movements and transformations of animate and inanimate separably moveable objects. For inanimate objects to move or be transformed, there has to be a transfer of external energy. Although animate objects obey the laws of physics, their particular motion paths and transformations are due to the generation of energy from within. I have dubbed these the Innards-Agent and External-Agent principles (Gelman et al. 1995). The idea is that the children benefited from an implicit, abstract causal framework, which informs the kind of perceptual information they take to be relevant and therefore salient for descriptions of similarity and actions. Thus, the framework provides input about what kind

of data are relevant to each sub-domain, in this case, cues for biological/living or inert things. The cues include ones that are relevant to the potential actions on the one hand, and potential functions, on the other hand. That is, the possible forms and details of each kind of object are part of implicit skeletal "blueprint" characterizations of the two ontological kinds.

Further evidence for this view was obtained in Massey and Gelman (1988). Children aged 3 and 4 were asked whether a series of objects could move themselves up and down a hill or whether they needed help. The objects all were novel. They included vertebrates and invertebrates, wheeled objects, statues that represented and shared parts of mythical human or animal creatures, and complex inanimate objects that resembled stick-like human figures. No graduate students could tell us what they were. Neither could the 3- and 4-year-olds, who successfully told us which objects could move by themselves both up and down a hill. What these young children said was most informative, as illustrated in the following sections from our transcripts.

Experimenter: *Could this (a statue)* [go up the hill by itself]?
Child: *No.*
Experimenter: *Why not?*
Child: *It doesn't have feet.*
Experimenter: *But look, it does have feet!*
Child: *Not really.*

In her own way, this child was telling us that the statue was not made up of the right kind of stuff. Another child told us that a statue was just a furniture statue, again an example from an inert category.

The results of this experiment also show that young children can use high-level, abstract causal principles, principles that outline the equivalence class of their entities, which differ for separably moveable animate and inanimate entities. Internal energy sources govern animates as well as the kinds of transformations, motions, and interactions that are permitted. External energy sources are taken as the source of the kinds of motions and transformations that inert objects exhibit. Of course animates honor the laws of physics, but they in turn have their own sources for generating goal-directed motions, responding in kind to other members of their species, and adjusting to unexpected features of the environment, such as holes, barriers, and so on (Gelman et al. 1995).

This brings me to the question of what counts as a *domain*. Randy Gallistel talked earlier about space and intentionality.[1] Simply put, a domain is a domain if

[1] See Chapter 4.

it has a set of coherent principles that form a structure and contains unique entities that are domain-specific. The domain of causality does not contain linguistic entities. It makes no sense to ask whether "movement" in a sentence – a linguistic variable – is due to biological energy or forces of nature. Similarly, it matters not how large an entity is when one engages counting principles (see below). When it comes to considerations of moving objects, the weight and size of an object is often paramount. To repeat: whenever we can state the principles that serve to capture the structure and the entities within it, either by themselves or ones generated according to the combination rules of the structure, it is appropriate to postulate a kind of domain-specific knowledge.

15.2 Core and non-core domains

I distinguish between *core* and *non-core domains* (Gelman and Williams 1998). The above account of a domain is neutral as to whether a given domain is innate or acquired. Like Spelke (2000), I reserve the phrase *core domain* for those that have an innate origin. I prefer to think of these as "skeletal." Of course the notion of "skeletal" is a metaphor meant to capture the idea that core domains do not start out being knowledge-rich. Nevertheless, no matter how nascent these mental structures, they are mental structures. And, like all mental structures, they direct attention and permit the uptake of relevant data in the environment. This leads me to favor structure-mapping as a fundamental learning mechanism. If we accept that young children have some core mental structures, we see that they have a leg-up when it comes to learning about the data that can put flesh on these.

Since non-core domains lack initial representational resources, it follows that learning about them will be hard. It is hard – in fact it is "hell on wheels" (HoW) – to master with understanding non-core domains. To do this, one has to both mount a structure and collect data that constitutes the knowledge in the domain. But we know that it is hard to acquire new conceptual structures. One has to work at the task for a considerable number of years and it helps to have formal tutoring. Often one's exposure to a new domain is incomprehensible. Imagine what beginning Chemistry students might think when they hear words like "bond," "attraction," and the like. They surely are not in a position to understand the technical meaning of these terms and therefore are at risk of misunderstanding them or even dropping the course. We know from research that such knowledge is the kind attributed to experts and we know that it takes a very great deal of work over many, many years to acquire expertise for any non-core domain. A characterization of non-core domains is presented below (see section 15.2.2). I now return to considerations regarding core domains from the perspective of very early learning.

Consider the domain of natural number arithmetic as an example of a core domain. Importantly, the principles of arithmetic (addition, subtraction, and ordering) and their entities (numerons and separate, orderable quantities) do not overlap with those involved in the causal principles and their link to separably moveable animate and inanimate objects. As a result examples of relevant entities and their properties are distinctly different. For no matter what the conceptual or perceptual entities are, if you think they constitute a to-be-counted collection of separate entities, you can count them. It is even permissible to decide to count the spaces between telephone poles (a favorite game of many young American children) or collect together for a given count every person, chair, and pair of eyeglasses in a room. This is because there is no principled restriction on the kinds of items counted. The only requirement is that the items be taken to be perceptually or conceptually separable.

In contrast, when it comes to thinking about causality, the nature and characteristics of the entities really do matter. One's plans about interactions with an object will be constrained by the kind of entity it is and its environments. If the entity is an animate object, I will take into account its size, whether it can bite, its posture, how fast it can move, and so on. If I want to lift two chairs, I certainly will take into account their size and likely weight. I will do the same should I be asked to also lift the two men sitting in those chairs. I know that I do not have the kind of strength it takes to transfer the relevant energy to lift and move the men in the chairs. I might be able to lift the chairs by themselves. So when it comes to considering the conditions under which objects move, their material, weight, and size do matter. This contrast accomplishes what we want – an a priori account of psychological relevance. If the learner's goal is to engage in counting, then attention has to be paid to identifying and keeping as separate the to-be-counted entities, but not the particular attributes of these, let alone their weight.

Similarly, if the learner's goal is to think about animate or inanimate objects, then attention has to be given to the information that provides clues about animacy or inanimacy: for example, whether the object communicates with and responds in kind to like objects, moves by itself, and is made up of what we consider biological material. Food surely is another core domain. We care about the color of a kind of food, even if we rarely care about the color of an artifact or countable entity. In this regard, it is noteworthy that children as young as 2 years of age also take the color of food into account (Macario 1991).

15.2.1 What are core domains?

(1) They are *mental structures*. However skeletal, they actively engage the environment from the start. This is a consequence of their being biological, mental organizations. As a result they function to collect domain-relevant data and hence provide the needed memory "drawer" for the build-up of knowledge that is organized in a way that is consistent with the principles of the domain.

(2) They help us solve the problem of *selective attention*. This avoids the common circular argument that selective attention is due to salience and salience directs attention. To repeat, potential relevant candidate data are those that fit the equivalence class outlined by the principles of the domain. It is the principles of the domain that offer the definition of the relevance dimensions.

(3) They are *universal*. To say that a core domain is universal is not to say that everyone will have the exact same knowledge or that learning about the domain will occur in one trial. It is well to keep in mind that linguists who assume that there are universal principles that support language acquisition do not expect children to learn their language in one trial. Further, variability across languages is taken for granted. Still, the assumption is that there are innate principles that help the child solve the learnability problem. My appeal to the universality of some small set of core domains should be thought of as being in the same vein. The principles serve to outline the equivalence classes of possible data. Since the kind of data a given culture offers young children varies as a function of geography, urbanization, etc., it follows that the range of knowledge about a domain will vary, just as do languages.

To appeal to universal innate principles is *not* to assume that learning does not take place. Instead, it forces us to ask what kind of theory of learning we need to account for early learnings and the extent to which these serve as bridges or barriers to later learnings. For a discussion of why the terms "innate" and "learned" are not opposites, given our theoretical perspective on learning, see Gelman and Williams (1998).

(4) They are akin to *labeled and structured memory drawers* into which the acceptable data "are attached." This provides an account of how it is possible to build up understanding of a coherent knowledge domain.

(5) They support *learning on the fly*. They do so because of the child's active tendencies to search for supporting environments – be these in the physical, social, or communicative worlds represented in the environment. The fact that learning occurs on the fly and is very much a function of what the child attends to is why many students of young children's early cognitive development have moved in this direction.

(6) The *principles of the structure and entities within a domain are implicit*. There is no claim that an infant or young child can state them, and I would bet that most adults cannot do so either, any more than they can state linguistic principles.

(7) Learning in these domains is *highly motivated* by the child. They ask relevant questions, including how a remote control works, why a parent says the car battery is dead, and what number comes after 100, 1000, etc. I well remember a little girl in a schoolyard telling me she was too busy to talk. She had set herself to count to "a million." I asked when she thought she would get there. Her reply was, "A very, very, very long time." She pointed out that she needed to eat, sleep, and probably would be very old.

Many young children's online inclinations to self-correct and rehearse are part of their overall tendencies to put into place the competencies that are within their purview. Examples of young children self-correcting their efforts or even rehearsing what they have just learned are ubiquitous in the developmental literature. A common report from parents has to do with their children asking "What's that?" after they have answered the question what seems like more than fifty times. Such rituals can go on for days and, then, without a clue, drop off the radar screen. In a related way, we are finding that the children in the preschools where we work are eager to have us ask more questions about unfamiliar animate and inanimate objects, no matter what the socioeconomic class represented by their families.

(8) The *number of core domains is probably relatively small*. They are only going to be as large as is necessary for us to get universal shared knowledge without formal instruction. To repeat what was presented above: just as different language communities support the acquisition of different languages, different language/cultural communities will favor differential uptake of the relevant data that they offer. Nevertheless, the underlying structure should be common – at least to start.

15.2.2 What are non-core domains?

(1) They are not *universal*; they have no representation of the targeted learning domain, and therefore no understanding of the data to start.

(2) They involve the *mounting of new mental structures* for understanding and require considerable effort over a very extended period of time, typically about ten years.

(3) The *number of non-core domains is not restricted*. This is related to the fact that individuals make different commitments regarding the extensive effort needed to build a coherent domain of knowledge and related skills. Success at

the chosen goal depends extensively on the individual's ability to stick with the learning problem, talents and the quality of relevant inputs, be these text materials, cultural values, and demonstrations and/or the skills of a teacher. Some examples of non-core domains include: chess, sushi-making, sailing, orchestra conductor, master chef, CEO, golf pro, car mechanic, dog show judge, discrimination learning; algebra, Newtonian physics, theory of evolution, theory of probability, composer, linguist, military general, abalone diver, and so on.

Learning about a non-core domain also benefits extensively from a teacher or master of the domain – an individual who selects and structures input and provides feedback. Still, no matter how well-prepared the teacher might be, the learner often has a major problem if she is unable to detect or pick up relationships or at least parts of relationships that eventually will relate to other relevant inputs. The task can be even more demanding if one has to acquire a new notational system, which can be hard in its own right.

Finally, early talent in non-core domains does not guarantee acquisition of expertise. It will take around ten years of dedicated work to reach the level of expert for the domain in question, be this musical composition, x-ray reading, chess, or Olympic competition, as well as a host of other areas, including academic ones. See Ericsson et al. (1993) for a review and theoretical discussion.

15.3 Early learning mechanisms

For me, the queen learning mechanism is *structure-mapping*. Given an existing structure, the human mind will run it roughshod over the environment, finding those data that are isomorphic to what it already stores in a structured way. This kind of learning of the data in a given domain need not take place in one trial. It could be that one first identifies the examples of the relevant patterned inputs and then maps to the relevant structure. Subsequently, further sections of the pattern are put in place. In any case, the details that are assimilated fit into a growing set of the class of relevant data that fill in the skeletal structure.

Importantly, input data can vary considerably on the surface, as long as they represent examples of the same principles and therefore are considered examples within the equivalence class of data that are recognizable by the principles. This carries with it the implication that the input stimuli do not have to be identical; in fact, they are most likely to be variants of the same underlying structure. Multiple examples are good for all kinds of reasons – different ways of doing the same thing, or beginning to look, compare, and contrast analogically

to see if they belong together. Given an existing structure, it is possible to have online self-monitoring correction, by which I mean that the child can say "That's not right; try again." In fact, in our counting protocols, we have examples of children saying, "One, two, three, five – no, try dat again!" – for five trials, then getting it right and saying, "Whew!" Nobody told the child to do this; he or she just did it. We see a lot of this kind of spontaneous correction or rehearsal of learning that is related to the available structure.

15.4 More on core domains: the case of natural number

There is a very large literature now on whether babies or even preschoolers count or not. An ability that counts as one in the domain is arithmetic, or more precisely, natural number arithmetic on the positive integers. First of all, the meaning of a counting list does not stand alone. There is nothing about the sound "tu" that dictates that it follows the sound "won" and so on. Instead the requirements are that a list of count words follow:

(1) the *one-to-one principle*. If you are going to count, you have to have available a set of tags that can be placed one-for-one, for each of the items, without skipping, jumping, or using the same tag more than once;
(2) the *stable order principle*. Whatever the mental tags are, they have to be used in a stable order over trials. If they were not, you could not treat the last tag as
(3) the *cardinal value,* which is conserved over irrelevant changes.

The relevant arithmetic principles are ordering, add, and subtract. Counting itself is constrained by three principles. If you want to know if the last tag used in a tagging list is *understood as a cardinal number,* it is important to consider whether a child relates these to arithmetic principles; it helps also to determine how the child treats the effects of adding and subtracting.

It helps to see that count words behave differently than do adjectives, even if they are in the same position in a sentence. In Fig. 15.2, one can see that it is acceptable to say that each of the round circles is *round* or *a circle,* but one cannot say that each of the five circles is *five* or *a five circle.* The other thing we know is this: if we put several objects in front of 2-year-olds who are just beginning to speak, they are likely to label the object kind. Hence it is not clear that they are going to say "One," when there is one object. Of interest is whether it is possible to switch the child from interpreting the setting as a labeling one or one for counting. If we can switch attention, and therefore show the setting is ambiguous for the child, we might pick up some early

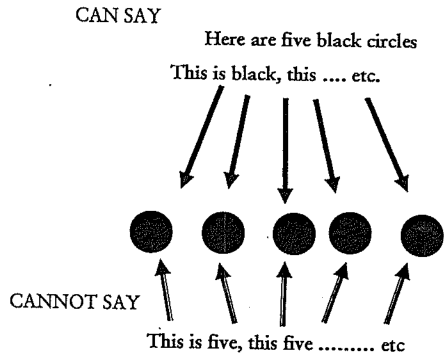

Fig. 15.2. A set of circles that can be labeled as five circles, black circles, or five black circles. Further, each can be called a black circle but not a five circle. This is because 'five' only refers to the set as a whole and not the individuals.

counting knowledge data. We accomplished this with a task that I call *What's on the Card?* (Gelman 1993).

We tested three age groups of children: those who ranged in age from 2 years 6 months to 2 years 11 months; 3 years 0 months to 3 years 2 months; and 3 years 3 months to 3 years 6 months. The following example of a protocol illustrates both the procedure and how our youngest children responded.

Experimenter: *See this card? What's on this card?*
Child: *A heart.*
Experimenter (feedback): *That's right. There is one heart on the card.*

Next two trials first show two hearts and then three hearts in a row:

Experimenter (with the 2-heart card): *See this card? What's on* this *card?*
Child (has now shifted and taken up the instruction to shift domain mindset): *Two hearts.*
Experimenter: *Show me.*
Child: *One, two.*
Experimenter: *So what's on the card?*
Child: *Two.*

And then we get a similar pattern for three hearts. There are several points to make about the procedure. As expected, the child first answered a wh- question with a label reply. However, when offered the option to treat differently subsequent examples that showed an increasing number of the item, the child took the bait. This was so for subsequent blocks of trials with new sets of cards, each set depicting different item kinds. Indeed, the youngest age group counted and indicated the cardinal value on 91 percent of their trials.

Thus, they understood our hint that they treat the display as opportunities to apply their nascent knowledge of the counting procedure and its relation to cardinality.

What about addition and subtraction? A rather long time ago I started studying whether very young children (2 1/2 years to 5 years) keep track of the number-specific effects of addition and subtraction. In one series of experiments, I used a magic show that was modeled after discussions with people in Philadelphia who specialized in doing magic with children. The procedure is a modification of a shell game. It starts with an adult showing a child two small toys on one plate vs. three on another plate. One is randomly dubbed the winner, the other the loser. The adult does not mention number but does say several times which is the winner-plate and which is the loser. Henceforth both plates are covered with cans and the child is to guess where the winner is. They pick up a can, and if it hides the winner plate they get a prize immediately. If they do not see a winner, they are asked where it is, at which they pick up the other one and then get a prize. The use of a correction procedure is deliberate: it helps children realize that we are not doing anything unusual, at least from their point of view. This set-up continues for ten or eleven trials, at which point the children encounter a surreptitiously altered display either because items were rearranged, or changed in color, kind, or number (more or less).

The effect of adding or subtracting an object led to notable surprise reactions. Children did a variety of things; such as put their fingers in their mouth, change facial expression, start searching, and even asking for another object (e.g., "I need another mouse"). That is, they responded in a way that is consistent with the assumption that addition or subtraction is relevant, and they know how to relate them. When we do this experiment on 2-year-olds, with 1 vs. 2 and then transfer to 3 vs. 4, we get a transfer of the greater-than or less-than relationship. That is, we have behavior that fits the description of the natural number operations.

Oznat Zur developed a new procedure that involved 4- to 5-year-olds playing a game that involved putting on different hats. Each hat signaled a new game for the child and either a repeat or variation of a condition. For example, children played at being a baker by selling and buying donuts. To start, a child was given nine donuts to put up on the bakery shelf and asked how many he had. Then someone came into the store with pennies and said, "I have two pennies." The child then handed over two donuts, at which point an adult experimenter asked him to predict, without looking or counting, how many were left. After making a prediction, the child counted to check whether it was right. This sequence of embedded predictions and checks continued. The children did very well. Their answers were almost all in the correct direction. And many of their answers fell within a range of $n \pm 1$ or 2. Further, the results were replicated in a class, the members of whom were about the same age but did not have an opportunity to play a comparable game before the experiment (Zur and Gelman 2004).

In yet another experiment, Hurewitz, Papafragou, Gleitman, and Gelman (2006) asked children ranging in age from 2 years 11 months to the late 3-year-old range to place a sticker either on a two- or four-item frame on one set of trials, or *some* vs. *many* on another set of trials. The children had an easier time with the request that used numerals as opposed to quantifiers. The word "some" gave them the most difficulties in this task, a finding that challenges the view that beginning language-learners find it harder to use numerals as compared to quantifiers.

15.5 Rational numbers are hard

I will conclude now with two contrasting numerical concepts: the successor principle and rational numbers. The successor principle captures the idea that there is always another cardinal number after the one just counted or thought about. This is because addition is closed under the natural numbers. As expected, when Hartnett and Gelman (1998) asked children ranging in age from about 6 years to 8 years of age if they could keep adding 1 to the biggest number that they could or were thinking about, a surprising number indicated that they could. Even when we suggested that a googol or some other very large cardinal number was the biggest number there could be, we were challenged by the child, who noted it was possible to add another 1 to even our number.

The successor principle is seldom taught in elementary school, whereas notions about fractions are. However, when it comes to moving on to considering rational numbers, and the idea that one integer divided by another is a rational number, we run into another example of a HoW domain. This perhaps is not surprising since there is no unique number between a pair of rational numbers. Formally, there is an infinite number of rational numbers between any two pairs of this kind of number. There is more to say about this, but I think that starts to give you the flavor that we really have moved into a different domain and that we may have a case of a conceptual change.

To end this presentation, I illustrate the kind of errorful but systematic patterns of responses we have obtained from school-aged children asked to place in order, from left to right, a series of number symbols, each one of which is on a separate card. Keep in mind that these children were given practice at placing sticks of different lengths on an ordering cloth; they were even told that it was acceptable to put sticks there of the same length but different colors and to move sticks, and then the test cards, until they were happy with their placement order. Careful inspection of the placements reveals that the children invented natural number solutions. For example, an 8-year-old started by placing each of three cards left to right as follows: 1/2, 2/2, 2 1/2, etc. The following interpretation captures these and all further placements. The child took the cards as an opportunity to apply his knowledge of natural number addition:

$$(1 + 2 = 3), (2 + 2 = 4), (2 + 1 + 2 = 5).$$

Other children invented different patterns but all invented some kind of interpretation that was based on natural numbers.

One might think that students would master the placement of fractions and rational number well before they enter college. Unfortunately, this is not the case. When Obrecht, Chapman, and Gelman (2007) asked whether undergraduates made use of the law of large numbers when asked to reason intuitively about statistics, they determined that students who could simply solve percent and decimal problems were reliably more able to do so. Those who made a lot of errors preferred to use the few examples they encountered that violated the trend achieved by a very large number of instances. This continues, unfortunately, through college. I will leave you with that. If you want to know now why your students are horrified and gasp when they are faced with a graph, it is probably because they do not understand rational numbers and measurement.

15.6 Conclusion

To conclude, humans benefit from core domains because these provide a structural leg-up on the learning problem. We already have a mental structure, albeit skeletal, to actively search the environment for relevant data – that is, data that share the structure of innate skeletal structures – and move readily onto relevant learning paths. The difficulty about non-core, HoW domains is that we have to both construct the structure and find the data. It is like having to get to the middle of a lake without a rowboat.

Discussion

HIGGINBOTHAM: There has been some interesting work in recent years by Charles Parsons on intuitions of mathematical objects – not intuitive judgment, but intuitions of the number 3.[2] What he observes is that, from some fairly simple premises, you start off making a stroke. You can envisage that it is possible that you can always add 1. If you have two sequences of strokes, then one of them is an initial segment of the other, and therefore if you took one off each one, they would be different. Now that is already all of the Peano axioms, except induction, and the question would be, when they have that, to check it by saying, "Look, here's this notation system. Can you reach any number that way?" If you can ask that question and get an answer, then you'll get the intuit, because Parsons is deliberately ambivalent or merely suggestive on this point.

GELMAN: Believe it or not, we haven't studied anything that is relevant. Before that, however, I do want to point out that I left out names of my collaborators on the study wherein young children correctly identified 2 and 4 but erred with the same arrays when their task was to identify *some* and *all*, one of whom is in the room, Lila Gleitman, and two of our post-docs at the time, Anna Papafragou and Felicia Hurewitz, who is the senior author of the paper that just came out.[3] As to your question, we ran another interview, where we said, "I am going to give you a dot-making machine that makes dots on paper and never breaks or runs out of paper. This is how many we have now. What happens if we push it (the button)? Will that be more dots on the paper?". Many children understood that the successive production of dots would never stop save for physical limits on themselves, i.e., "that that would never stop ... [except] if you died, had to eat or go to sleep." This is an example of the nonverbal intuition about the effect of an iterative process.

HIGGINBOTHAM: Yes, to get induction, you need something more. You need the idea that for any number x, if I make enough strokes, I can get to x.

GELMAN: Yes, we didn't ask that one, but there is another one where we asked the question in the Cantorial way. That is, children who were having no trouble with our initial infinity interview were engaged in a version of Cantor's proof. We had drawings of hands in a line, each of which was holding hands with a numeral in a parallel line placed in one-to-one correspondence. We then asked whether we could keep adding hands and numerals, one at a time. This done, we went on to ask whether there were as many hands as numerals. The children agreed. In fact, they agreed at first that equivalence would hold if each person was paired with an odd number. The kids would say yes, probably because they had said yes to the first questions. "You know, they had the same answer."

[2] Parsons (1990).
[3] Hurewitz et al. (2006).

But then when we pointed out the contradiction, that we were skipping every even number, the reaction was, "Oh no, this is crazy, lady. Why are you wasting my time?" It probably is the case that even these children did not understand the abstract notions that follow from one-to-one correspondence. However, it is not so easy to develop a task that is free of confounding variables. The trick is to figure out exactly how to ask what you want to get at. And it isn't that easy, because you have to tell them, "I want you to tell me what the induction is," without telling them that I want you to tell me that. My bottom line? Be careful about saying that there are groups of people who cannot count with understanding, who have only a few number words.

PIATTELLI-PALMARINI: You mentioned quantifiers versus numbers, and not surprisingly, numbers are easier than quantifiers. In fact, there is a dissertation in Maryland, by Andrea Gualmini, showing that children have a problem in understanding quantifiers until very, very late.[4] Do you have further data on the understanding of quantifiers?

GELMAN: The question of when quantifiers are understood is very much complicated by the task. I don't know that dissertation, but I know studies from the 1970s showing that the quantifier tasks (*all* and *some*, etc.) were not handled well until 6 years of age. We actually have been able to change the alligator task (Hurewitz et al. 2006) so that the kids do very well on *all* and *some* questions. The problem is, fundamentally, that we are talking about a set-theoretic concept. Once you make it easier, move them out of the full logic of class inclusion or one-on-one correspondence, the task does get easier, but that is in a sense the point of why I don't understand why anybody thinks the quantifiers are a primitive out of which come the count numbers. The formal rules for quantifiers, whichever formal system you go into – it is going to be different, because whatever that system is, it will have a different notation, there will be different rules about identity elements than there are in arithmetic, and the effect of adding we automatically know is different. I mean, if you add *some* to *some*, you get *some*. If you add 1 to 1, you don't get 1. So these are very different systems, and furthermore, the quantifiers are very context-sensitive. It depends on what numbers you are working with. So when we looked across the tasks, we could start doing task analysis, but we haven't done it completely.

URIAGEREKA: Just a brief follow-up on that. I think in principle it would be useful to bring in the notion of conservativity, which is quite non-trivial for binary quantifiers, as has been shown. So not only would you have numerals

[4] Gualmini (2003).

versus quantifiers, but among the quantifiers, you would have the ones where in effect you have an order restriction and a scope, versus the ones where you don't, and that probably can make a big difference too.

GELMAN: I totally agree. I should just say I have no argument with that. This is not an accidental combination of people working together. We have, now, two faculty members who specialize in quantifiers and their acquisition, and these are all issues they have written about, are going to work on, and so on. My interest was that this was a way to demonstrate experimentally what I have written about as a purely formal distinction. I had tried to show why arguments about development that involve the count words coming out of the quantifiers didn't make any sense. But that was the logical argument. It was now nice to be able to show that they do behave separately.

URIAGEREKA: This partly also relates to the claim of context sensitivity, because strictly speaking, it is when you do have to organize the part that the quantifier leaves on with regard to the scope that you need massive context sensitivity, but not the other way around.

GELMAN: Right.

CHAPTER 16

The Learned Component of Language Learning

Lila Gleitman

Isolated infants and children have the internal wherewithal to design a language : there isn't one around to be learned (e.g., Senghas and Coppola 2001 Such languages exhibit categories and structures that look suspiciously like thos of existing languages. There are words like *horse* and *think*. Not only that: th mapping between predicate type and complement structure is also quite orthodox as far as can be ascertained. For instance, even in very primitive instances of suc self-made languages, *sleep* is intransitive, *kick* is transitive, and *give* is ditransitiv (e.g., Feldman, Goldin-Meadow, and Gleitman 1978). This fits with recent dem onstrations – one of which I mentioned during the round-table discussio (see page 207) – that even prelinguistic infants can discriminate between certai two- and three-argument events in the presence of the (same) three interactin entities (Gordon 2003). All of this considerable conceptual and interface appar atus being in place, and ("therefore") language being so easy to invent, one migh wonder why it's hard to acquire an extant language if you are unlucky enoug to be exposed to one. For instance, only ten or so of the required 50,000 o so vocabulary items are acquired by normally circumstanced children on an single day; three or four *years* go by before there's fluent production of modestl complex sentences in all their language-particular glory. What takes so long?

The answer generally proposed to this question begins with the problem o word learning, and is correct as far as it goes: ultimately, lexical acquisitio is accomplished by identifying concepts whose exemplars recur with recurren phonetic signals in the speech or signing of the adult community. That is, w match the sounds to the scenes so as to pair the forms with their meanings Owing to the loose and variable relations between word use and the passin scene – the "stimulus-free property of language use," as Chomsky (1959c