

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

2009

CHAPTER 4

The Foundational Abstractions

C. R. Gallistel

4.1 A short history of the mind

By way of prelude, I make a rapid – and necessarily superficial – tour of familiar philosophical terrain, because the material on animal cognition that I then review has substantial bearing on long-standing philosophical issues of relevance to contemporary cognitive science.

4.1.1 Empiricist epistemology

In this epistemology, the newborn mind knows nothing. But it has the capacity to experience elemental sensations and to form associations between those sensations that recur together. Thus, all representation derives from experience: “There is nothing in the mind that was not first in the senses” (Locke 1690). The mind’s capacity to associate sensations makes it possible for experience to mold a plastic mind to reflect the structure of the experienced world. Thus, concepts derive their form from the form of experience. The farther removed from sensory experience a concept is, the more derived it is.

In this epistemology, our concepts of space, time, and number are maximally derivative. They are so far removed from sensory experience that they do not seem to have sensory constituents at all. Nor is it clear how their highly abstract, essentially mathematical form can be derived from experience. Neither the nature of the relevant experience, nor the inductive machinery necessary to derive them from that experience are in any way apparent. And yet these abstractions seem to play a foundational role in our representation of our experience.

4.1.2 Rationalist epistemology

Kant famously responded to this puzzle by arguing that the empiricists were wrong in attempting to derive our concepts of space, time, and number from our experience of the world. On the contrary, Kant argued, these organizing concepts are a precondition for having any experience whatsoever. We always represent our experiences, even the most elementary, as ordered in time and localized in space. The concepts of time and space are not derivable from our experience; rather, they are the foundation of that experience.

4.1.3 Cartesian dualism and human exceptionalism

Descartes famously argued that the machinery of the brain explains unmindful behavior. But, he argued, some behavior – behavior informed by thought – is mindful. He further argued that the operations of thought cannot be the result of mechanical (physically realizable) processes. He was among the originators of a line of thought about mind in human and non-human animals that continues to be influential, not only in popular culture but in scholarly and scientific debate. In its strongest form, the idea is that only humans have minds. In its weaker form, it is that humans have much more mind than non-human animals. A corollary, often taken for granted, is that the farther removed from humans an animal is on the evolutionary bush, the less mind it has. The most popular form of this idea in contemporary thought is that animals, like machines, lack representational capacity. Therefore, abstractions like space, time, number, and intentionality do not inform the behavior of non-human animals.

The popularity of the view that non-human animals know nothing of time, space, number, and intentionality owes much to the lingering effects of the behaviorism that dominated scientific psychology until relatively recently, and that still dominates behavioral neuroscience, particularly those parts of it devoted to the investigation of learning and memory. The more extreme behaviorists did not think that representational capacity should be imputed even to humans. Radical behaviorism fell out of favor with the rise of cognitive psychology. The emergence of computers, and with them, the understanding of the physics and mathematics of computation and representation played an important role in the emergence of contemporary cognitive psychology. The fact that things as abstract as maps and goals could demonstrably be placed into the indubitably physical innards of a computer was a fatal blow to the once widespread belief that to embrace a representational theory of mind was to give up the hope of a material theory of mind. The realization that a representational theory of mind was fully compatible with a material theory of mind was a

critical development in scientific thinking about psychology, because, by the early twentieth century, a theory of mind that made mind in principle immaterial was no longer acceptable in scientific circles.

By the early twentieth century, the progress of scientific thought made Descartes's concept of an immaterial mind that affected the course of events in a material nervous system unacceptable to the great majority of scientists committed to developing a scientific psychology. The widespread belief in a uniquely human mind did not, however, die with the belief in a materially effective immaterial mind. Rather, the belief in a uniquely human form of mental activity came to rest largely on the widely conceded fact that only humans have language. If one believes that language is *the* (or, perhaps, *a*) medium of thought, then it is reasonable to believe that language makes possible the foundational abstractions. One form of this view is that it is language itself that makes possible these abstractions. Alternatively, one may believe that whatever the unique evolutionary development is that makes language possible in humans, that same development makes it possible to organize one's experience in terms of the foundational abstractions.

4.2 The birds and the bees

The history of thought abounds in ironies. One of them is that Sir Charles Sherrington's enormously influential book *The Integrative Action of the Nervous System* (Sherrington 1906) did as much as any work to persuade many scientists that a purely material account of mental activity – an account couched in neuroanatomical and electrophysiological language – was possible. The irony is that Sherrington, who died in 1952, was himself strongly committed to a Cartesian dualism. He believed that when he severed the spinal cord he isolated the purely physical neural machinery of the lower nervous system from the influence of an immaterial soul that acted on levels of the nervous system above his cut.

Sherrington placed the concept of the synapse at the center of thinking about the neurobiological mechanisms of behavior. His student, Sir John Eccles (1903–1997), further enhanced the centrality of the synapse in neuroscientific thinking by confirming through intracellular recordings of postsynaptic electrical processes Sherrington's basic ideas about synaptic transmission and its integrative (combinatorial) role. Eccles, too, was a Cartesian dualist, even though he secured the empirical foundations on which contemporary connectionist theories of mind rest. The irony is that a major motivation for connectionism is to found our theories of mind not only on physically realizable

processes but more narrowly on the understanding of neuroanatomy and neurophysiology that Sherrington and Eccles established. Indeed, the neurobiology commonly mentioned as a justification for connectionist theorizing about the mind is exactly that elaborated by Sherrington a century ago. Discoveries since then have made no contribution to the thinking of contemporary modelers.

A similar irony is that the empirical foundations for the now flourishing field of animal cognition were laid by behaviorist psychologists, who pioneered the experimental study of learning in non-human animals, and by zoologists, who pioneered the experimental study of instinctive behavior in birds and insects. Both schools were to varying degrees uncomfortable with representational theories of mind. And/or, they did not believe they were studying phenomena in which mind played any role. Nonetheless, what we have learned from the many elegant experiments in these two traditions is that the foundational abstractions of time, space, number, and intentionality inform the behavior of the birds and the bees – species that last shared an ancestor with humans several hundred million years ago, more than halfway back in the evolution of multicellular animals.

Some years ago (Gallistel 1990a), I reviewed the literature in experimental psychology and experimental zoology demonstrating that non-human animals, including birds and insects, learn the time of day (that is, the phase of a neurobiological circadian clock) at which events such as daily feedings happen, that they learn the approximate durations of events and of the intervals between events, that they assess number and rate (number divided by time), and that they make a cognitive map of their surroundings and continuously compute their current location on their map by integrating their velocity with respect to time. Here, in this paper, I give an update on some further discoveries along these lines that have been made in recent years.

4.2.1 *Birds and time*

The most interesting recent work on the representation of temporal intervals by birds comes from a series of brilliant experiments by Nichola Clayton, Anthony Dickinson, and their collaborators demonstrating a sophisticated episodic memory in food-caching jays (Clayton et al. 2006; Clayton et al. 2003, and citations therein; see also Raby et al. 2007). In times of plenty, many birds, particularly many species of jays, gather food and store it in more than ten thousand different caches, each cache in a different location, spread over square miles of the landscape (Vander Wall 1990). Weeks and months later, when food is scarce, they retrieve food from these caches. Clayton and Dickinson and their

collaborators took this phenomenon into the laboratory and used it to show that jays remember what they hid where and how long ago and that they integrate this information with what they have learned about how long it takes various kinds of food to rot.

The experiments make ingenious use of the fact that jays are omnivores like us; they'll eat almost anything. And, like us, they have pronounced preferences. In these experiments, the jays cached meal worms, crickets, and peanuts. Other things being equal, that is the order of the preference: they like meal worms more than crickets, and crickets more than peanuts. In one experiment, hand-reared jays, with no experience of decaying food, were given repeated trials of caching and recovery. They cached two different foods in two different caching episodes before being allowed to recover their caches. In the first of each pair of caching episodes, they were allowed to cache peanuts on one side of an ice-cube tray whose depressions were filled with sand. In the second episode of each pair, they were allowed to cache either mealworms or crickets on the other side of the same tray. Thus, on some caching trials, they hid peanuts in one half of the trays and mealworms in the other, while on other trials, they hid peanuts in one half and crickets in the other.

Either 4 hours, 28 hours, or 100 hours (4 days) after each pair-of-caching episode, they were allowed to recover food from both sides of the trays. On trials with only a 4-hour delay, both the mealworms and the crickets were still fresh and tasty when retrieved. At that delay, the jays preferred to retrieve from the caches where they had hidden either mealworms or crickets (depending on whether they had cached peanuts-and-mealworms or peanuts-and-crickets). On trials where a 28-hour delay was imposed between caching and recovery, the experimenters replaced the cached mealworms with mealworms that had been artificially rotted. Thus, on the first few peanuts-and-mealworms trials with a 28-hour delay before retrieval, the jays found inedible "rotten" mealworms where they had cached tasty fresh mealworms. By contrast, on peanuts-and-crickets trials, they found crickets that were still fresh after 28 hours in their caches. On trials with a 4-day delay before recovery, both the mealworms and the crickets had rotted; the peanuts alone remained fresh.

Control birds that never encountered rotted caches preferred the caches where mealworms and crickets had been hidden no matter how long the delay between caching and recovery. The experimental birds preferred those caches when only four hours had elapsed. When twenty-eight hours had elapsed, their preference after a few trials of each type depended on whether it was mealworms or crickets that they had hidden on the "better" side of the tray. If it was mealworms, they preferred the peanut caches, but if it was crickets, they preferred the cricket caches. When four days had passed, their preference after

a few trials (during which they learned about rotting) was for the peanut caches, whether it was mealworms or crickets that they had hidden on the "better" side of the tray.

In an ingenious extension of these experiments, Clayton, Yu, and Dickinson (2001) showed that the birds would adjust their retrieval preferences on the basis of information about rotting time acquired after they had made their caches. At the time the caches were made, they did not yet know exactly how long it took the meal worms to rot.

It appears from these experiments that the remembered past of the bird is temporally organized just as is our own. The birds compute elapsed intervals and compare them to other intervals in memory. They compare the time elapsed since they cached a cricket to what they have since learned about the time it takes a cricket to rot. Like us, birds reason about time.

4.2.2 *Birds reason about number*

There is an extensive literature showing that pigeons and rats can base behaviorally consequential decisions on estimates of the approximate number of events (Brannon and Roitman 2003; Dehaene 1997; Gallistel 1990a). In many of the experiments, the animal subjects make a decision based on whether the current number is greater or less than a target number in memory. Thus, these experiments give evidence that animal minds reason about number as well as about time. Brannon and her collaborators (Brannon et al. 2001) extended this evidence using a task that required pigeons to first subtract the current number from a target number in memory and then compare the result to another target number in memory.

In their experiment, the birds pecked first at the illuminated center key in a linear array of three keys on a wall of the test chamber. Their pecking produced intermittent flashes (blinks) of the light that illuminated the key. The ratio of the number of pecks made to the number of flashes produced varied unpredictably, for reasons to be explained shortly. After a number of flashes that itself varied unpredictably from trial to trial, the two flanking keys were illuminated, offering the bird a choice.

Pecking either of the newly illuminated side keys generated further intermittent flashes. Eventually, when the requisite number of further flashes on the side key they first chose had been produced, the bird gained brief access to a feeding hopper. For one of the side keys the requisite number was fixed. This number was one of the target numbers that the birds had to maintain in memory. For the other side key, the number of flashes to be produced was the number left after the flashes already produced on the center key were subtracted from a

large initial number. This large initial number was the other number that had to be maintained in memory. The greater the number of flashes already produced on the center key, the smaller the difference remaining when it was subtracted from this large initial number; hence, the more attractive the choice of the "number-left" key relative to the "fixed-number" key. The pigeons' probability of their choosing the number-left key in preference to the fixed-number key depended strongly and appropriately on the magnitude of the number left relative to the fixed number.

The random intermittency of the flashes partially deconfounded the duration of pecking on the center key from the number of flashes produced by that pecking, allowing the authors to demonstrate that the pigeons' choices depended on number, not duration.

4.2.3 *Birds and intentionality*

Jays are not above stealing the caches of others (Bednekoff and Balda 1996). Experienced jays are therefore reluctant to cache when another jay is watching. They remember which caches they made while being watched and which jays were watching them (Dally et al. 2006). When no longer watched, they selectively re-cache the food that others observed them cache (Emery and Clayton 2001). "Experienced" jays are those who have themselves pilfered the caches of other jays; those innocents who have not succumbed to this temptation are not yet wary of being observed by potential thieves while caching (Emery and Clayton 2001). Thus, nonverbal animals represent the likely intentions of others and reason from their own actions to the likely future actions of others (see also Raby et al. 2007).

4.2.4 *Bees represent space*

The zoologist Karl von Frisch and his collaborators discovered that when a foraging bee returns to the hive from a rich food source, it does a waggle dance in the hive out of sight of the sun, which indicates to the other foragers the direction (bearing) and distance (range) of the source from the hive (von Frisch 1967). The dancer repeatedly runs a figure-8 pattern. Each time it comes to the central bar, where the two circles join, it waggles as it runs. The angle of this waggle run with respect to vertical is the solar bearing of the source, the angle that a bee must fly relative to the sun. The number of waggles in a run is a monotonic function of the range, that is, the distance to the source.

It is somewhat misleading to say that the dance communicates the solar bearing, because what it really communicates is a more abstract quantity,

namely, the compass bearing of the source, its direction relative to the north-south (polar) axis of the earth's rotation. We know this because if the foragers that follow the dance and use the information thus obtained to fly to the source are not allowed to leave the nest until some hours later, when the sun has moved to a different position in the sky, they fly the correct compass bearing, not the solar bearing given by the dance. In other words, the solar bearing given by the dance is time-compensated; the users of the information correct for the change in the compass direction of the sun that has occurred between the time when they observed the dance and the time when they use the directional information they extracted from it. They are able to do this, because they have learned the solar ephemeris, the compass direction of the sun as a function of the time of day (Dyer and Dickinson 1996). Man is by no means the only animal that notes where the sun rises, where it sets, and how it moves above the horizon as the day goes on.

Knowledge of the solar ephemeris helps make dead reckoning possible. Dead reckoning is the integration of velocity with respect to time so as to obtain one's position as a function of time. Successful dead reckoning requires a directional referent that does not change as one moves about. That is, lines of sight from the observer to the directional referent must be parallel regardless of the observer's location. The farther away the point of directional reference is and the more widely perceptible from different locations on the earth, the better it serves its function. In both of these respects, the sun is ideal. It is visible from almost anywhere, and it is so far away that there is negligible change in its compass direction as the animal moves about. The problem is that its compass direction changes as the earth rotates. Learning the solar ephemeris solves that problem.

Dead reckoning makes it possible to construct a cognitive map (Gallistel 1990a: Chapter 5) and to keep track of one's position on it. Knowledge of where one is on the map makes possible the setting of a course from wherever one currently is to wherever one may suddenly wish to go. The computation involved is simple vector algebra: the vector that represents the displacement between one's current location and the goal location is the vector that represents the goal location minus the vector that represents one's current location. The range and bearing of the goal from one's current location is the polar form of that displacement vector.

There is a rich literature on navigation in foraging ants and bees, which make ideal subjects, because they are social foragers: they bring the food they find back to the communal nest, then depart again in search of more. In this literature, one finds many demonstrations of the subtlety and sophistication of the spatial reasoning that goes on in these miniature brains, which contain only on the order of 1 million neurons. For some recent examples, see Collett and

Collett (2000); Collett et al. (2002); Collett and Collett (2002); Harris et al. (2005); Narendra et al. (2007); Wehner and Srinivasan (2003); Wittlinger et al. (2007); Wohlgenuth et al. (2001). For a review of the older literature, see Gallistel (1990a: Chapters 3–6). Here, I have time to recount only two of the most important recent findings.

For many years, researchers in the insect navigation field have questioned whether ants and bees make an integrated map of their environment (e.g., Collett and Collett 2004; Dyer 1991; Wehner and Menzel 1990; but see Gould 1990). The alternative generally proposed is that they have memorized range-bearing pairs that enable them to follow by dead reckoning routes back and forth between familiar locations. They have also memorized snapshots of the landmarks surrounding those locations (Collett et al. 1998; Collett et al. 2002; Collett 1992; Collett and Baron 1994) together with the compass directions of those landmarks, and they have memorized snapshots of landmarks passed en route between these locations (Fukushi and Wehner 2004). But, it is argued, all of this information is integrated only with regard to a particular route and summoned up only when the ant or bee is pursuing that route (Collett and Collett 2004).

Part of what has motivated skepticism about whether the information from different routes is integrated into an overall map of the environment is that bees often appear to fail a key test of the integrated-map hypothesis. The question is, can a bee or ant set a course from an arbitrary (but recognizable!) location on its map to an arbitrary goal on its map? One way to pose this question experimentally is to capture foraging bees when they are leaving the hive en route to a known goal and displace them to an arbitrary point within their foraging territory. When released at this arbitrary new location, do they reset their course, or do they continue to fly the course they were on when captured? Under some conditions, they do reset their course (Gould 1986; Gould and Gould 1988; Gould 1990), but in most experiments, most of the bees continue to fly the course they were on (Dyer 1991; Wehner and Menzel 1990). This suggests that they cannot recompute the course to their old goal from their new location.

Against this conclusion, however, is the fact, often reported in footnotes if at all, that the bees who take off for the wild blue yonder on a course inappropriate to their goal (given their release location) are nonetheless soon found either at the goal they had when captured or, more often, back at the hive. They do not go missing, whereas bees released in unfamiliar territory do generally go missing, even if that territory is quite close to the hive.

The problem has been that we had no idea what happened between the time the bees disappeared from the release site flying on the wrong course to the

time they reappeared, either at their intended goal or back at the hive. Menzel and his collaborators (2005) have taken advantage of the latest developments in radar technology to answer the question, what do misdirected bees do when they discover that they have not arrived at their intended goal? Radar technology has reached the point where it is possible to mount a tiny reflector on the back of a bee and track that bee at distances up to a kilometer. Thus, for the first time, Menzel and his collaborators could watch what misdirected bees did. What they did was fly the course they had been on when captured more or less to its end. This brought them to an equally arbitrary location within their foraging terrain. They then flew back and forth in a pattern that a sailor, aviator, or hiker would recognize as the sort of path you follow when you are trying to “get your bearings,” that is, to recognize some landmarks that will enable you to determine where you are on your map. At some point this flying back and forth hither and yon abruptly ended, and the bee set off on a more or less straight course either for the goal they had been bound for when captured or back to the hive. In short, they can set a course from an arbitrary location (the location where they find themselves when they realize that they are not getting where they were going) to another, essentially arbitrary location (the location of the feeding table they were bound for). This result argues in favor of the integrated map hypothesis.

The final result I have time to report (Gould and Gould 1988; Tautz et al. 2004) moves the level of abstraction at which we should interpret the information communicated by the waggle dance of the returned bee forager up another level. These little-known results strongly suggest that what the dance communicates is best described as the map coordinates of the food source. Moreover, it appears that before acting on the information, potential recruits consult their map for the additional information that it contains.

In these experiments, a troop of foragers was recruited to a feeding table near the hive, which was then moved in steps of a few meters each to the edge of a pond and then put on a boat and moved out onto the pond. At each step, the table remained where it was long enough for the troop foraging on it to discover its new location and to modify appropriately the dance they did on returning to the hive. So long as the table remained on land, these dances garnered new recruits. But when the table was moved well out onto the water, the returning foragers danced as vigorously as ever, but their dances did not recruit any further foragers – until, in one experiment, the table approached a flower-rich island in the middle of the pond, in which case the new recruits came not to the boat but to the shore of the island, that is, to the nearest plausible location. In short, bees’ past experience is spatially organized: like the birds, they remember

where they found what, and they can integrate this spatially indexed information with the information they get from the dance of a returning forager.

4.3 Conclusions

The findings I have briefly reviewed imply that the abstractions of time, space, number, and intentionality are both primitive and foundational aspects of mentation. Birds and bees organize their remembered experience in time and space. The spatio-temporal coordinates of remembered experience are accessible to computation. The birds can compute the intervals elapsed since they made various caches at various locations at various times in the past. And they can compare those intervals to other intervals they have experienced, for example, to the time it takes a given kind of food to rot. The bees can use the dance of a returning forager to access a particular location on their cognitive map, and they can use that index location to search for records of food in nearby locations. Birds can subtract one approximate number from another approximate number and compare the result to a third approximate number. And birds making a cache take note of who is watching and modify their present and future behavior in accord with plausible inferences about the intentions of the observer.

To say that these abstractions are primitive is to say that they emerged as features of mentation early in evolutionary history. They are now found in animals that have not shared a common ancestor since soon after the Cambrian explosion, the period when most of the animal forms now seen first emerged.

To say that they are foundational is to say that they are the basis on which mentation is constructed. It is debatable whether Kant thought he was propounding a psychology, when he argued that the concepts of space and time were a precondition for experience of any kind. Whether he was or not, these findings suggest that this is a plausible psychology. In particular, these findings make it difficult to argue that these abstractions arose either from the language faculty itself or from whatever the evolutionary development was that made language possible in humans. These abstractions appear to have been central features of mentation long, long before primates, let alone anatomical modern humans, made their appearance.

Discussion

RIZZI: I was wondering how far we can go in analogy between the foraging strategy that you described and certain aspects of language. I wondered whether there is experimental evidence about strategies of rational search of this kind:

first you go to the closer spots and later to more distant spots. A particular case that would be quite interesting to draw an analogy with language would be the case of intervention, presenting intervention effects in these strategies. For instance, just imagine a strategy description of this kind, that there is a direct trajectory for a more distant cache; there is one intervening spot with a less desirable kind of food (let's say nuts rather than peanuts, or rather than worms). Would there be anything like experimental evidence that this kind of situation would slow down somehow the search for the more distant spots – or anything that would bear on the question of whether there are distance and/or intervention effects in search strategies? Because that is very typical of certain things that happen in language – in long-distance dependencies.

GALLISTEL: As regards the second part of your question, on the interfering effect of an intervening, less desirable cache, I don't know of anything that we currently have that would be relevant, although it might very well be possible to do this. The setup that Clayton and Dickinson used, as I just said, doesn't lend itself at all to that because it's not like a natural setup where this situation would arise all the time. The birds are just foraging in ice-cube trays. However, some years ago we did a traveling salesman problem with monkeys, where they very much have to take distance into account, and where they have to take into account what they are going to do three choices beyond the choice that they are currently making. That is, the monkeys had to harvest a sequence, going to a number of cache sites. This was done by first carrying a monkey around and letting it watch while we hid food, before releasing it to harvest what it had seen hidden. The question was, would it solve the traveling salesman problem by choosing the most efficient route, particularly in the interesting cases where to choose the most efficient route, the least-distance route, you would have to, in your current choice, foresee or anticipate what you were going to do in a subsequent task. And they very clearly did do that. They clearly did show that kind of behavior, so I think that's relevant.

HAUSER: One of the puzzles of some of the cases that you brought up is that lots of the intimate knowledge that the animals have been credited with seems to be very specialized for certain contexts, which is completely untrue of so much of human knowledge. So in the case of the jays, it seems to be very, very located to the context of cache recovery. Now, maybe it will eventually show itself in another domain. We're taking advantage of natural behavior so maybe it will not. But in the same way that the bees seem to be one of the only species that externalize this knowledge in the communicative signal in a richness that is totally unparalleled in any other species but humans, so you get this kind of odd thing where the bees are only really sort of talking about one specific context.

You have rich social relationships, but there is no communicative signal outwards at all. So the question is – the way I’ve put it in the past is – animals have this kind of laser-beam intelligence and we have this kind of floodlight, and what happens? How do you get from this very, very selective specialization to probably a promiscuous system in humans?

GALLISTEL: Well, of course the competence–performance distinction is just as important in interpreting the behavior of animals as it is in interpreting the language of humans. They have a lot of competences that they don’t always choose to show us. But I agree with your basic point, and in fact it is something I have often emphasized myself. Animals show a lot of competence in a very sharply focused way. If I were to venture into perilous terrain and ask what language does for thought, one suggestion that one might offer is that, because it allows you to take these representations that arise in different contexts with, on the surface, different formal structure, and map them onto a common representational system, it may enable you to bring to bear the representational capacity of this module on a problem originally only dealt with by that module, and so this module can contribute something that the original module wouldn’t have been able to do on its own. And that would be where the floodlight quality of human reasoning came in perhaps. The idea that language didn’t really introduce new representational capacity, except perhaps insofar as it created a representational medium in which anything could be, to some extent at least, represented.

URIAGEREKA: At some point I would like to hear your opinion, Randy, on this *Science* report on the bees doing their dance also for the purpose of finding a new nest, so the behavior is apparently not fully encapsulated for the purposes of foraging. I had no idea that they also did that, find a viable nest with procedures akin to those involved in foraging. I don’t know how plastic that is. The point I’m trying to emphasize is this: would we find more of those apparently plastic behaviors if we knew where to look? That said, in the case of plasticity that we have seen in our system, my own feeling (and this is sheer speculation) is that generalized quantification – that is, the type of quantification that involves a restriction and a scope – is certainly central to much of human expression, but may be hard to find in other species. In fact, if Elena Herburger is right in her monograph on focus, this sort of full-fledged, crucially binary quantification may even be central to human judgment, especially the way Wolfram Hinzen is pushing that idea. It may be that the type of syntax you require for that type of quantification (which is one of the best understood systems in linguistics), however it is that we evolved it, might as well liberate, if you will, a kind of richly quantificational thought that I would actually be very

interested to see if animals exhibit. I mean, you know much more than I do about these things, Randy, but the experiments I have read do not get to generalized quantification. For example, in dolphin cases in the literature, it is reported that these animals get, say, *bring red ball*, *bring blue ball*, and so on; let’s grant that much. But apparently they do not get *bring most ball* or even *bring no ball*. So maybe that would be another way to push these observations, another thing to look for, constructing experiments to test for behaviors of that truly quantificational sort.

CHOMSKY: Randy’s comment sort of suggests Liz Spelke’s experiment,¹ i.e. using language for intermodal transfer (visuo-spatial, for instance).

GALLISTEL: You’re right, it does seem to, but in fact I’m not sympathetic to that. I don’t agree with Liz on the interpretation of those experiments, but what I said does seem to point in that direction.

GELMAN: I’d like to modify what Randy said, to say that what seems to be unique to humans is a representational capacity. Language is one that can be used for a wide range of activities, but notational capacities are also representations. Drawings can be representations, plans, and so forth – there are many options. And I have yet to see data that animals can go invariably from one representational format to another.

PARTICIPANT: It’s only a simple question. Do the systems of communication of bees and birds display feedback? For example, if they make a mistake and then realize that they’ve made a mistake, do they communicate it?

GALLISTEL: Ahhhh [scratches head; laughter]. That’s tough! Sort of implying that as a result, where the bees that are following the dance consult their map, sort of implying that they conclude that the dancer didn’t know what the dancer was talking about, right? [Chuckles to himself.] Because if the information conveyed by the dance is sufficiently inconsistent with the information on their map, they appear to discount the information in the dance. I’m not sure whether that isn’t correcting themselves, of course. I’m not sure this is relevant, but there are recent experiments by Laurie Santos,² one of Marc’s many good students, who has gone on to do work that Marc has also done on observing the mind sort of thing, where you have to represent whether the other animal knows what you know, in order to choose. This has been a big issue for a long, long while. But I thought her recent experiments, which I cannot reproduce (I’m sure Marc can, as they were partly or mostly undertaken with Marc)

¹ Lipton and Spelke (2003).

² Santos et al. (2002).

were very persuasive on that score. Part of Marc's genius has been to exploit naturalistic circumstances, and they exploited naturalistic circumstances in a way to make a much more compelling case that the animal knew that the other animal didn't know X.

PARTICIPANT: I was wondering if you have feedback when you have something similar to negation. It is usually claimed that negation is unique to human language...

GALLISTEL: Ohhhh, like where the catcher in a baseball game shakes off the signal? I can't quickly think of a clear example that one could regard as equivalent to negation. But negation is certainly a kissing-cousin of inversion, and animals invert all the time. I mean, they invert vectors, right? Not only do they calculate the home vector themselves when they are out there and they have found food, but when they get back, what they are dancing is not the vector they calculated coming home, but the inverse vector, the vector for going the other way. About negation, I always remember that tee-shirt that says, "What part of No don't you understand?" [Laughter]. It seems to me about as elementary as you can get.

PIATTELLI-PALMARINI: Concerning foraging, I have seen work by my colleague Anna Dornhaus, concerning some of the optimal criteria that honeybees meet in foraging,³ which is rather astounding, because they have constructed a graph of how many bees are proactive (they go out and look for food) versus the reactive foragers that wait for the dance. So they have calculated the percentages of proactive versus reactive, and the graph you get depends on how long the food is available. And you have a triple point like in second-order phase transitions in physics and chemistry. It's extraordinary. They have a number of predictions that sound very weird, but then they observe them in nature or in the laboratory. So it seems that, when we approach foraging in a quantitative way, among other things, it is one of those fields in which the species seem to be doing the best thing that they could possibly do. Have you any comments on that, because it is a question of great current interest in linguistics. It wouldn't be the only case in which you have biological systems that are doing the best that can be done.

GALLISTEL: Yes, this question of optimality is apt to provoke very long arguments in biological circles. I can give you sort of a general view, and then my own particular view. If you look on the sensory side, you see spectacular optimality. That is, sensory transduction mechanisms are, most of them, very near the limits of what is physically possible. So the threshold for audition, for

³ Dechaume-Moncharmont et al. (2005).

example, is just above the threshold set by physics – there's a slight vibration on the eardrum due to the fact that on a small surface there is stochastic variation in how many molecules of air hit that surface, and that produces a very faint vibration in the eardrum that is an ineliminable noise in the system. And the amount of additional vibration that you need from another source is just above that limit. The most essential thing is to calculate how much the eardrum is moving at that threshold. It is moving less than the diameter of an atom! So that's a lot better than you would have thought at the beginning.

Similarly with the eye. One of the proofs before it was directly demonstrated that the absorption of a single photon by a single rhodopsin molecule in a single rod generated a signal that could make its way all the way through the nervous system came from a famous experiment by Hecht, Shlaer, and Pirenne in which they showed that there was a clearly detectable effect.⁴ This was subsequently studied by Horace Barlow and Barbara Sakitt,⁵ and they showed that for every quantum or photon of light absorbed, there was a quite sizeable increase in the probability that a human would say that he had detected the flash. There are ten million rhodopsin molecules in the outer segment of a single rod, and there are a million rods in the retina. So it is a little bit like one of these huge soccer matches and someone burps and the referee says, "Who burped?" There are a hundred million spectators and somehow the burp is centrally detectable. That's pretty impressive.

There is wide agreement about this – the facts are extremely well established. When you come to computational considerations, that is where the arguments begin, but of course that reflects the fact that we, unlike the sensory things, don't know what's going on. Most neuroscientists think that the computations are just one spike after the next, right? But this seems to me nonsensical. Any engineer will tell you that the contradictions that follow the transduction of the signal are more important than the transduction in the first place. That is, if you've got a good signal but lousy signal processing, then you've wasted your time producing a good signal. So it seems to me that the pressure to optimize the computations is at least as great as the pressure to optimize the signal transduction, and we know that the signal transduction is very near the limits of what is physically possible. So I tend to think that the computations, or processing of the signal, are also at the limits of what is computationally possible. But since we know practically nothing about how the nervous system computes, it's hard to say.

⁴ Hecht et al. (1942).

⁵ Sakitt (1972).