6. Doubt, Deference, and Deliberation: Understanding and Using the Division of Cognitive Labor

Frank Keil

In 'The "Meaning of Meaning"' Hilary Putnam (1975) famously suggested, as part of a more general defense of an externalist account of semantic content, that our referential practices are upheld by a 'division of linguistic labor'. A speaker who lacks the individual capacity to identify a term's referent may nonetheless use the term successfully, so long as she belongs to a linguistic community where some (group of) expert(s) have that capacity. Over the past thirty years, there has been a lively discussion about the implications of such a theory for questions about wide vs. narrow content, internalism vs. externalism about meaning, and the like. In the spirit of labor-division, I will leave discussion of these matters to others (see e.g. Burge 1979; Fodor 1998; Prinz 2002), focusing instead on a question that has received relatively little attention in either the philosophical or the psychological literature, namely how laypeople understand the nature and character of the division of cognitive labor.

In particular, I will consider how non-experts understand the ways in which knowledge might cluster in other minds. I will describe four distinct ways that people might think about the division of cognitive labor and say something about how those different ways are used to make sharply contrasting inferences about domains of expertise. Although there is evidence that all four ways are available quite early in cognitive development, there are also striking differences in how they are used at various ages. The kind of expertise that Putnam implied as guiding deference for the meanings of natural kind terms, namely that of the natural sciences, gradually comes to hold a privileged status during middle childhood. This pattern of developmental change in turn sheds light on the everyday value that attaches to having insight into the division of cognitive labor.

DIVISIONS OF LABOR

It is hardly news that cultures divide up chores in ways that create different areas of expertise. As cultures became less nomadic, crafts and skills emerged with distinctive experts in each. Economists and sociologists have long argued that divisions of labor are an essential part of increasing productivity in a culture (Smith 1776; Durkheim 1947; Hume 1739). In most human cases, divisions of physical labor carry with them implications for divisions of cognitive labor. A person who achieves greater skill in an area is likely to have distinctive cognitive capacities that support that skill. In addition, most divisions of cognitive labor in humans reflect different paths of learning, different experiences, and immersion in different local communities of knowledge. Given its pervasiveness across cultures, it is surprising that there has been relatively little work in the field of cognitive anthropology devoted to the cognitive bases of divisions of labor (Hutchins 1995).

Important psychological questions arise concerning the division of cognitive labor. How do most collective enterprises, such as the basic sciences, engineering, legal systems, and medicine, function when each individual only has a fraction of the necessary knowledge and understanding to make the whole enterprise work? In particular, how does one access a domain of knowledge in other minds when one is largely ignorant about that domain? If we know already that an individual has one piece of knowledge, how do we decide what else that person is likely to know? How do we decide which of two competing experts is more likely to be a source of correct information?

The answers to such questions open up several topics that overlap with the field of 'social epistemology' (Goldman 1999, 2001). For the most part, they are also beyond the scope of this paper, as are questions about how members of a scientific community divide up their labor (Kitcher 1990). Instead, the more narrow goal of this paper is to consider the psychological heuristics that people use to think about how knowledge might be clustered in other minds. What do we need to know outside our own areas of expertise to be able to expand on our knowledge in those unfamiliar areas?

There are several distinct ways of thinking of how knowledge might be clustered in other minds, ways that draw on different sorts of cognitive requirements and which can be explored through experimental studies. Since detailed descriptions of those studies are under way elsewhere in journals with a more experimental focus (Danovitch and Keil 2004; Lutz and Keil 2002; Keil 2003*a*), the focus of this paper will be on elaborating four distinct ways of thinking about expertise, summarizing the main findings of the experimental studies conducted by my laboratory group with adults and children, and considering how our developing understanding of the division of cognitive labor might be used in everyday life.

I will consider four ways of thinking about expertise: by category association, by privileged access, by goal implementation, and by underlying causal structure. These four possibilities do not exhaust the set of ways of thinking about knowledge clusters but they are the four most commonly used by laypeople. Moreover, they each suggest quite different heuristics for figuring out who knows what.

CATEGORY ASSOCIATION DIVISIONS OF KNOWLEDGE

Expertise can be understood as about anything normally associated with a category, providing that the categories involved are at the basic level of categorization or below. The basic level of categorization is the highest level at which categories seem to bristle with correlated properties not found at the next level up (Rosch *et al.*, 1976; Murphy 2003). These levels can vary somewhat across individuals and cultures, but normally would be at a level of chairs, tables and sofas and below that of furniture. Similarly, shirts, pants and sweaters are the basic level below that of clothing, and cars, trucks and motorcycles from a basic level below that of vehicles. The basic level is also the level of categorization at which children also tend to use their first words to pick out sets of things in the world (Mervis and Crisafi 1982).

The category association heuristic assumes that people have knowledge clusters consisting of all pieces of information normally associated with members of a low-level category. Thus, one might plausibly think of people who are chair, or motorcycle, or pants experts. Even more plausibly, one can think of experts at levels below the basic level, such as Hitchcock chair, off road motorcycle, and ski pants experts. The basic level is the highest level at which we might normally employ the category association strategy. It is less plausible, however, to think of thorough experts on all kinds of furniture, or vehicles, or clothing. The lower the level, the more one might plausibly think that a person knows most anything associated with a category. Thus, a ski pants expert might be expected to know the history of ski pants design, the costs of ski pants, which celebrities and racers wear what kinds of pants, and so on. A clothing expert could hardly be expected to have comparable diversity and detail of knowledge about all clothing.

Clustering of knowledge by category association might employ a simple cognitive heuristic. One merely needs to think of all bits of information that are normally associated with most members of that category. If I want to know something more about off road motorcycles, I might look for a person who demonstrates detailed knowledge about a few aspects of motorcycles and assume all other details will be known as well. This knowledge is perhaps best captured by the idea of people who are 'fans' or 'fanatics'. Elvis fans might be thought to know everything about Elvis, ranging from his songs, to his personal life, to the places he lived. Train fanatics might know everything about the history of trains, the ways trains worked, and the economic factors associated with trains. At a sufficiently low level of categorization, we might think it plausible that expertise could consist in having exhaustive knowledge of members of the category.

Where does this heuristic come from? It may arise from a social motivational hypothesis that people develop intense likes and dislikes for some categories; and, as a result, are deeply interested in everything frequently associated with most members of that category. We infer a drive to know 'everything' about a category either because it is highly valued or because it is a source of morbid fascination. The category association heuristic may also arise from the apparent ease of using a related strategy involving common lexical items. If John knows that 'Poodles' are F_1, where F_1 is an unusually detailed fact about poodles, simply assume that John is likely to know that 'Poodles' are F_n for any fact about poodles. Without knowing anything more about John or poodles one can blindly use the strategy of assuming that John is likely to have greater than average knowledge of the truth of virtually any sentence that makes a statement about 'poodles'. It would also be trivial to implement this strategy in a simple computer program that is fed text strings the size of sentences. If the category is low enough, a person's knowledge can be considered as exhausting everything that is typically associated with members of those categories or mentioned in discourse about lexical items that refer to that category.

Category exhaustion is interesting because it seems to be the simplest and most straightforward way of figuring out who knows what. The seductive simplicity of this heuristic makes it especially attractive to young children and to adults in cognitively loaded tasks. Thus, if one puts individuals under powerful time pressures, has them do several things at once, or inserts salient distracters in a task, these cognitive 'loads' tend to cause people to abandon more difficult cognitive heuristics in favor of simpler ones. Though subjects may not reveal their reliance on these heuristics in less pressured settings, cognitive load tasks can help experimenters identify which simple heuristics play a role in their everyday cognitive processing.

The category association approach, however, can be seriously misleading for one straightforward reason. It is virtually never the case people have exhaustive knowledge of members of a category, no matter how low the level. Moreover, as seen shortly, this strategy fails to predict other sorts of important elements of knowledge that can be reliably inferred from a few things that a person knows.

'ACCESS-BASED' DIVISIONS OF KNOWLEDGE

The socio-economic or subcultural practices of a society can often be used to think of divisions of cognitive labor that are 'access-based'. Thus, we can assume that different groups of people have different forms of expertise because they have been in proximity to a particular form of information that others have not by virtue of their station in life. For example, one might infer that a person who knows more than average about fine wines, resort spas, and charter jets, has that knowledge by virtue of being wealthy and therefore one might also expect that person to have greater knowledge about designer clothing, plastic surgeons, and home security systems. A person who knows more than average about soup kitchens, friendly police precincts, and warm heating vents may have that knowledge by virtue of being homeless and therefore is expected to have greater knowledge about homeless shelters and places with low and high rates of pedestrian traffic.

An understanding of access-based knowledge requires some sense of how people cluster in stable or semi-stable groups in a culture and what bits of information might be distinctive to those groups. This knowledge is not based on a category or category label but rather on an understanding of the distinctive environments of subgroups and of the experiences offered by those environments. It might be based on simple associations of activities with members of that group, or it might involve induction of totally novel forms of knowledge based on an understanding of the group and why it coheres as such. Thus, if one believes that the wealthy tend to pick activities that are exclusive by virtue of the expenses associated with engaging in those activities, one can induce that wealthy people are more likely to know about some novel but highly expensive product. In this way, an understanding of the division of cognitive labor on the basis of access can have a generative quality

This generative property helps illustrate why access-based models of expertise are not variants of the category exhaustion strategy applied to the special case of social categories. When one relies on beliefs about why and how a group of people choose activities, the ability to then induce a large set of new forms of expertise contrasts with a mere list of all facts associated with the members of the category. Moreover, access-based strategies also exclude some forms of knowledge that might be associated with a category but which do not follow from causal explanatory beliefs about a particular kind of access. For example, wealthy individuals in the United States are more likely to know about local Republican politicians because of a strong association between wealth and support of Republicans (Green *et al.* 2002); but the access-based heuristic of expertise described earlier for wealth relies on the notion of increased knowledge of expensive goods and activities and might not see the relevance of party affiliation.

'GOAL-CENTERED' DIVISIONS OF KNOWLEDGE

Different people have different relatively long-term goals. One person may want to play professional soccer, another to heal the sick, and another run a successful fish charter business. Knowledge of another's goals, plus some knowledge of how those goals are normally achieved, can also be used as a basis for inferring clusters of knowledge in other minds. Thus, a person whose goal is to run a successful fish charter business might be expected to know more than average about topics that would further the goal of having a large number of customers in a financially viable manner. That person is likely to know more than average about fish seasonal migration and foraging patterns, about marine weather, about diesel engines, about marine navigation, and about financing and insurance for commercial boats.

Goal-centered ways of understanding the division of cognitive labor are far-ranging and potentially powerful. They tend to go far beyond mere association of bits of information with individuals to a kind of problem-solving about what it takes to be successful in an endeavor and how the structure of a situation, such as a boat in a marine environment with customers, imposes certain requirements that in turn call on specific forms of expertise. The more one knows about the environment and about human capacities in such environments the more one can generate inductions about likely bits of knowledge in that area. An account of goal-based heuristics requires both a first order analysis of the knowledge of the goal-directed agent (e.g. the fishing boat skipper) and a second order analysis of the knowledge that one might have of goal-directed agents and their likely knowledge.

Goal-based ways of clustering of knowledge would seem to be those most closely associated with how the division of physical labor evolved in societies. Weavers, potters, farmers, and healers all developed expertise that furthered their relatively straightforward, and usually very public, goals. To infer who knows what in the world, one needs to keep track of different goals of groups of individuals and note how those goals are normally achieved. Even knowledge of a completely novel goal can often yield quite fertile inductions about knowledge. There is, for example, a group of individuals known as 'disk recovery specialists', whose goal is to recover data from computer hard drives that have become inoperable. I had never heard of that group until quite recently, yet a simple knowledge of their goal allowed me to induce what those professionals are likely to know about: how hard drives work, a huge array of software and disk operating systems, market rates for data recovery, and legally binding contracts between specialists and clients.

Sometimes, the goal becomes subordinate to a causal understanding of a set of closely related phenomena associated with that goal. For example, suppose one's goal is to treat cancer. As one pursues that goal, the biology of cancer starts to loom larger than the goal itself, which depends largely on an in-depth understanding of the relevant biology. Indeed, many of the sciences as we know them today grew out of goal-based practices, in which a rich pattern of causal regularities became far more the focus of knowledge than the goal itself. The goal of transforming base elements into gold or silver was unattainable but led to an increased understanding of chemistry. The goals of breeding better crops, livestock, and pets led to a greater understanding of the biology of genetics. Some goals, such as those of the fishing charter business, intrinsically draw on many domains at once and continue to do so in their most advanced and refined forms; but others bring into relief the causal patterns and regularities of a particular science, which leads to the last way of understanding the clustering of knowledge.

CAUSAL-CLUSTER, OR DISCIPLINE-BASED, DIVISIONS OF KNOWLEDGE

For many academics, especially those in the natural and social sciences, the most obvious ways of clustering knowledge is by academic disciplines, with the additional assumption that such disciplines arise because of distinctive patterns of causal regularities in the world. Departments of biology, chemistry, physics, and psychology are often said to exist because there are special causal patterns that are signatures of each of those areas. We tend to assume that there is a relatively small set of core principles that govern much of what happens in a domain and that, by virtue of knowledge of those principles, people have greater than average knowledge of the indefinitely large number of phenomena arising from such principles. The canonical case is knowledge of Newtonian mechanics. We assume that a person who knows Newton's laws of motion and a certain level of mathematics is likely to understand virtually any set of interactions between bounded physical objects. (We may mistakenly underestimate the complexity of some multi-bodied systems, but the assumption as described is quite common.) Many scientists similarly assume there are comparable sets of principles underlying chemistry, biology, and other disciplines with further subdivisions within that form hierarchies of subdisciplines.

Understanding knowledge clusters in terms of underlying causal patterns might seem to be a rarified way of thinking about the division of cognitive labor. Perhaps it is a recent cultural invention that is only within the strong grasp of scientists. Wouldn't one need relatively sophisticated exposure to those causal patterns to be able to appreciate how they might be used as a way of understanding of the organization of knowledge in other minds? A brief consideration of some different versions of realism and how each might influence the division of cognitive labor helps one see why it might be otherwise and how psychological studies are relevant.

All forms of realism embrace the idea that there are enduring patterns of regularities in the world independent of human activities on that world. They differ, however, in the extent to which they see a world of fundamentally distinct sorts of regularities. Consider the contrast between the view that all of nature is reducible to an account couched in terms of the laws of physics and the view that there are distinct levels of explanation such that the laws of economics, for example, cannot be reduced to those physics (Fodor 1974). Antireductionist views would seem most naturally associated with a division of cognitive labor corresponding to each of the levels of explanation they embrace. Reductionist approaches, in contrast, need not make such commitments.

Even at the same level of explanation, realists can debate about the extent to which the world should be seen as a relatively homogeneous network of causal links between properties, or whether it should be seen as more heterogeneous, consisting of distinct causal patterns with their own architectural principles. Should the world be seen as 'dappled' with different clusters of regularities or as more consistently all of the same type (Cartwright 1999)?

A dappled world-view offers a natural way of explaining how different realms of expertise might emerge, especially one that endorses 'thick' causal relations in which the causal relations such as 'compress', 'support', 'allow', 'feed', and 'prevent' are thought to be intrinsically different from each other and not reducible to a generic notion of cause (Cartwright 2003). Different realms might have different clusters of thick causal relations typically associated with them as well as different ways of describing the interactions between those relations. Perhaps one domain, such as evolutionary theory, uses intrinsically statistical arguments while another, such as the mechanics of macroscopic bounded objects, does not. Expertise in one of these domains might therefore be compartmentalized and not easily generalized to another.

Realists can also debate the extent to which there is a privileged way of carving up the world as opposed to an indefinitely large number of alternative ways, each of which might be based in a different form of real world structure and process. For example, laypeople often assume that there are two distinct natural kinds corresponding to 'trees' and 'flowers'. In most of the biological sciences, however, the tree/flower contrast is meaningless. Daisies and apple trees are much more similar to each other in terms of microstructural properties, evolutionary ancestors, and DNA structure than apple trees and pine trees. Pine trees, in turn, are more similar to ferns than they are to oak trees (Dupre 1981). At the same time, the layperson is picking up on a real physical difference between trees and flowers. Indeed, computer simulations of how ancient plants would solve the problem of growing taller to get more light all tend to converge on structural solutions similar to modern trees with stout reinforced trunks and root structures and certain overall shapes that maximize light exposure to surfaces (Niklas 1996). There are two different sets of causal regularities that give rise to different sets of stable kinds, each of which might be stable because of its own form of causal homeostasis (Boyd 1999).

One can take the tree/flower case as suggesting a 'promiscuous realism' in which there are indefinitely many realities that can be articulated over the same class of entities (Boyd 1999). This view can, in turn, devolve into a form of social constructionism in which real world structure becomes arbitrary and where human convention and invention fully explain the domains of scientific inquiry (Hacking 1999; Kukla 2000). A more nuanced view sees the sciences as akin to the making of maps (Kitcher 2001). Maps are correct, or true, by virtue of their correspondence with some set of relations in the world; but even given that strong commitment to realism, there are many such correspondences. (Just consider all the different kinds of maps one can have of a large city.) Thus, the map metaphor illustrates how the relationship between the causal structure of the world and domains of expertise, while quite varied, is not arbitrary. Intuitions about domains of expertise may also arise from social constructions or from innate biases about domains of inquiry; but there are versions of realism in which persistent causal regularities give rise to families of maps corresponding to domains of expertise. Studies on the psychological mechanisms people use to ascertain the division of cognitive labor therefore not only have the potential to inform how we access and rely on knowledge in other minds but also to shed some light on how our knowledge of the world might be related to the structure of that world. Moreover, if laypeople and children are able to pick up on those patterns of causal regularities, they might have some insight into domains of expertise roughly corresponding to the natural and social sciences without ever having direct exposure to those sciences.

INTUITIONS ABOUT WHO KNOWS WHAT

One way to explore intuitions about clusters of knowledge would be to simply ask people for their intuitions of what the scientists and other experts know; but such free-form questioning tends to yield a diverse and unstructured body of information about all activities associated with scientists. In our laboratory we have taken a more focused strategy, (Danovitch and Keil 2004; Lutz and Keil 2002). I describe phenomena that a person understands well and then ask what other phenomena that person also understands by virtue of understanding those initial phenomena. Most often this technique has been done as a triad task in which a person is described as knowing a great deal about a particular phenomenon and is then asked which of two other phenomena the person is also likely to know about. By presenting a forced choice between two alternatives it is possible to create various contrasts, or minimal pairs, that allow one to explore the relative 'pulls' of different dimensions. Thus, the format is typically of the form:

> John knows a great deal about why P1.
> What else is he likely to know about?
> Why P2?
> *or*
> Why P3?

For example:

> John knows a great deal about why water is transparent to light.
> What else is he likely to know about?
> Why gold conducts electricity so well?
> *or*
> Why gold prices rise in times of high inflation?

This sort of technique arguably reflects a common, real-life, way in which we attempt to rely on the division of cognitive labor. When trying to understand which of several possible people to approach so as to acquire a better understanding of a phenomenon, we will take as important data what each of those people already know, seeking out the relevant dimension of similarity between their known knowledge and the new phenomenon.

Several questions arise with respect to people's intuitions about triads of this sort. To what extent do people need explicit access to the causal mechanism themselves to be able to make a judgment of knowledge clustering? Is the structure of scientific disciplines in the modern university related in any way to laypeople's intuitions about knowledge clusters? How successful are laypeople at using underlying causal principles to cluster knowledge as opposed to clustering by surface objects, access, or goals? Finally, how do such patterns of judgment vary across development and across cultures? A series of studies have begun to provide answers to these questions.

In several studies, we described eight domains: physical mechanics, chemistry, adaptive/ecological biology, physiological biology, cognitive psychology, social psychology, economic and political science.[1] The divisions we chose correspond to distinct departments in at least some universities, although the two subareas of biology and psychology are often collapsed together. This particular group of eight was chosen because it could be placed into a neat symmetrical hierarchy of the natural and social sciences which are then further divided into the physical and biological sciences and the psychological and 'social system' sciences. This hierarchy allowed us to ask if items that were 'closer' together at the bottom of the hierarchy, such as physics vs. chemistry, would be harder to distinguish as knowledge clusters than those that were 'further' apart, such as physics vs. psychology. This hierarchy does reflect some Procrustean distortion of the disciplines into a more neatly ordered structure than really exists, but if it captures some degree of real-world structure, it should be reflected in patterns of judgments.

Expert informants who generated the items were asked to list phenomena that could easily be recognized and understood as such by both adults and elementary school children and would not involve any technical terms or exotic relations. Thus, the path of bouncing of a ball would be a better item than the nature of precession in gyroscopes. From a large set of generated items, the experimenters then selected a set that seemed clearest and least ambiguous and most likely to be accessible to children as well as adults. The items were further edited to make sure that various lexical cues to clustering were unlikely to be useful. Thus, if one physical mechanics item asked about the bouncing

[1] We avoided the humanities as it is much less plausible that those domains are organized around a set of core processes or causal relations that generate surface phenomena. Thus, the areas of study of an English department are more likely to be organized around various periods of literature and particular authors or regions and not around mechanisms of irony or production of imagery.

of balls the other physical mechanics item that it might be pitted against would not include a reference to a ball or bouncing, but might instead refer to the speed at which a pendulum swung.

Most of our adult subjects have been college students in North America, a limitation addressed partly by our developmental studies. These adults performed in a manner that was nearly 'error' free, meaning that they would cluster items that were in the same disciplines as more likely to be known in common. For example, if told that one person 'knew all about why a basketball bounces better on the sidewalk than on the grass', they would judge that the same person was much more likely to know 'why a big, heavy boat takes a really long time to stop' than 'why laundry soap cleans kids' dirty clothes'. The basketball and boat cases are both in the domain of mechanics while the soap case is in the domain of chemistry. Because their performance was so high, there was not a strong distance effect in which items further apart in the hierarchy were more easily seen as distinct. Nonetheless, there was a modest effect along these lines. An equally important finding was that many adult subjects were unable to actually explain the phenomena that they clustered together. For example, an adult might judge that a person who knew a great deal about 'why people sometimes fight more when they are tired' would be more likely to know 'why people smile at their friends when they see them' than 'why salt on people's icy driveways makes the ice melt sooner'. In many cases, adults would report that they had no idea of why the phenomenon occurred but were highly confident of their clustering judgment. Similarly, most adults easily judged that a person who knows a great deal about 'why sugar gives us energy to run around and do things' is more likely to know 'why bug spray in the water hurts fish' than 'why policemen can't put people in jail without a reason'—yet many of those same adults were unable to provide even the simplest explanations for those phenomena.

The coupling of a strong confidence in judgments with a frequent inability to explain the basis for such judgments suggested developmental studies. Children might also have a sense of the division of cognitive labor based on discipline-like principles even if they were unable to articulate those principles. Several studies with children ranging in age between 5 and 10 years have now shown that quite young children do have intuitions about the division of cognitive labor that map roughly onto those corresponding to the academic disciplines. There is also evidence that the 'distance' between the disciplines, as represented by their hierarchical relations, influences performance. Thus, even 5 year olds were at above chance levels on contrasts such as physics vs. cognitive psychology or economics vs. adaptive/ecological biology, approaching almost 70 per cent correct response rates. By contrast 5 year olds were unable to distinguish cognitive from social psychology as in the following example:

> This expert knows all about why some people act like leaders and some people act like followers.
> Do they know more about why people forget things when they get interrupted by the telephone ringing?
> *or*
> Do they know more about why people help each other when they're in trouble?

Nine year olds, on the other hand, immediately saw the contrast and clustered like adults.

Thus, by 5 years of age, children are showing some ability to cluster knowledge in a manner that seems to correspond to the ways in which phenomena are generated by common underlying sets of causal relations. Although the children rarely mentioned such causal relations directly, they do seem to have some implicit sense of broad patterns of causation distinctive to different domains of the natural and social sciences. These might include notions that mechanics is a domain with immediate causal consequences between objects that are monotonically related to the causal force of the first object. By contrast, in the social psychological realm, interactions are often non-monotonic and can occur with considerable delays.

Because the younger children so rarely explained their answers we had to use more indirect methods to assess what causal schemas they might be using. In one follow-up study, we tested the presence of such simple causal schemas by using cases that were technically in a domain such as mechanics but which did not contain a salient causal schema and others that were not in mechanics but had a component that was similar to a causal schema in the domain of mechanics. For example, it appeared that young children saw a coherent domain consisting of bounded objects in motion where consequences were monotonically related to the speed of the initial object mentioned. It was quite easy for them to cluster together these cases. However, when asked about problems of static mechanics, such as bridge structures, the children were less sure about how to cluster that piece of knowledge. Conversely, when presented with a phenomenon in psychology that involved a salient bounded object in motion ('John knows why you cannot see a bullet moving by you'), some younger children erroneously clustered that knowledge with mechanics.

To what extent could children be solving these problems by simply noting word co-occurrence patterns in roughly paragraph-sized chunks of text? Perhaps children don't need any sense of the causal patterns that exist in the world; they merely need to keep mental tabulations of how often terms such as ' ball', 'bounce', 'fall', and ' hit' co-occur. Then, they cluster phenomena based on their mental tabulations of how much the words in two phenomena have been noted to co-occur in bodies of text. The more powerful co-occurrence methods also tabulate how often words co-occur with an intervening word as a measure of conceptual similarity (Landauer and Dumais 1997). Thus, if 'ball' and 'bounce' co-occur frequently and 'bounce' and 'spring' co-occur frequently, even if 'ball' and 'spring' rarely co-occur, 'ball' and 'spring' will be judged as more similar because of the intervening relationship with 'bounce'. This procedure has been automated and strings of words can be put into programs based on large bodies of text, which then calculate conceptual relatedness.

Such frequency-based cues may help see knowledge clusters of various sorts, but they cannot be the sole basis. In the studies with children just described, the sentences describing the phenomena were fed into a popular frequency-based computer program (Landauer and Dumais 1997). As the sentences were constructed with an eye towards minimizing influences of frequency, it was expected that the program could not cluster the phenomena on discipline-based grounds. Indeed, it was at chance. Even in a study where preschoolers engaged at above chance levels of sorting, word frequency cues were ruled out (Lutz and Keil 2002). Another possible cue might be the clustering together of certain phenomena in instructional curricula. This alternative is more difficult to definitively test, but a look at elementary school curricula in the natural sciences (there is virtually none in the social sciences) suggests that very little information is imparted that would convey the appropriate clusterings.

In short, it appears that, by the age of 5, and possibly even in the later preschool years, when children are asked to cluster bits of causal explanatory knowledge (i.e. knowledge of 'why' for various phenomena) their judgments appear to be based on inferences about what kinds of causal patterns give rise to those phenomena. They seem to assume that a person who understood one phenomenon well must have done so by virtue of a good grasp of the causal principles that gave rise to that phenomena and therefore is likely to understand other phenomena arising from the same causal principles. On the few occasions where children did attempt to justify their responses, they often talked about the underlying basis for the phenomena and not about the experts or the knowledge itself. For example, one child clustered together two economics items because they both involved 'selling' (even though selling was never explicitly mentioned in the examples). That child said nothing about the experts themselves. Through their intuitions about knowledge clustering, these children are reflecting some of the divisions of knowledge that correspond roughly to natural and social science departments in the modern university. They see these clusters even though most of them have never heard of such departments.

FRAGILITY OF DISCIPLINE-BASED KNOWLEDGE CLUSTERS IN CHILDREN AND A CONTINUING TENSION

Even though young children do cluster knowledge in a manner that accords roughly with some academic disciplines, this ability is fragile when it is faced with competing ways of clustering knowledge. Thus, if a child is presented with a phenomenon that is caused by a certain set of causal relations but also has a salient goal, the goal may well dominate clustering decisions with other phenomena. For example, if a child is told that a person knows all about how marbles bounce off each other in the game of marbles and can use that to win a lot, the child might think the person is more likely to be an expert on another non-mechanics phenomena associated with winning at marbles (e.g. 'why different colored marbles help you keep track of who is winning?') than on a phenomenon that is mechanics but is unrelated to marbles (e.g. 'why yo-yos come back up?'). When goal-based clusters are pitted against discipline-based ones, the goal-based ways of clustering tend to dominate in younger children, with a dramatic shift occurring during the elementary school years such that discipline-based choices start to dominate by the age of 10 years (Danovitch and Keil 2004). When domains such as mechanics and psychology were pitted against salient goals, the goals won out in almost all 5 year olds and many 7 year olds. Discipline-based ways of clustering knowledge, although available to young children when presented with no competition from goal-based or category association heuristics, are not particularly salient or privileged early on. Instead, goal-based ways of clustering knowledge are more appealing to younger children.

Between roughly the ages of 5 and 10 years, however, a view develops in which underlying causal principles come to be seen as especially powerful ways of understanding the division of cognitive labor. We are currently exploring several mechanisms that might be helping to bring about this shift. One important influence may be the use of higher and higher level category labels with increasing age. We have recently shown that even kindergarteners are more likely to think that an 'animal' expert would have animal knowledge clustered on the basis of biological principles while a 'duck' expert might well be understood as having knowledge organized around goals or category labels (e.g. knowing everything and anything about ducks). The higher the category, the more implausible it is that knowledge would be clustered by goal or topic. For example, when told that a person knew a lot about ducks and asked if she would know more about 'why ducks need to sleep' or about 'why ducks are in a lot of cartoons' children chose roughly equally between these two alternatives. But when told that a person knew a lot about animals, children of all ages made the discipline-based choice ('why ducks need to sleep') by a huge margin. Since children's language reveals an increasing use of higher-order terms with age (Mervis and Crisafi 1982), it may well be that use of such terms helps reveal the special nature of discipline-based clusters.

The tension between discipline-based clusters and other forms remains in adolescence and on into adulthood. If one increases the cognitive load of the knowledge-clustering tasks, people may start to be influenced by topics or goals. For example, if instead of presenting people with triads, they are presented with a large set of say, forty-eight file cards with different phenomena on each and asked to cluster them into like kinds, roughly 35 per cent of adults will cluster them by category labels as opposed to underlying causal discipline (Keil and Rozenblit 1997).

In short, there are clear signs of a sensitivity to causal structure in very young children, a sensitivity that can be used as a way of thinking about the division of cognitive labor. This way of clustering knowledge, however, is just one of many for young children and seems to be cognitively more challenging than alternatives such as goals and surface topics. During the elementary school years there is a profound shift in which clustering knowledge by underlying causal structure comes to have a privileged status, at least in simple triad tasks. We are currently exploring more fully the basis for this shift and how it relates to other changes in how children understand the nature of knowledge and its distribution in other minds. We are interested in how changes in various patterns of language use might provide clues to the special status of knowledge clustered on the basis of causal principles. In addition, we are interested in whether richer understanding of underlying causal mechanisms in one domain can act as a kind of model that triggers a bias for that way of clustering knowledge in all domains.

FOCUSING THE LENS ON UNDERLYING CAUSAL STRUCTURE

Not all ways of asking about what others know shine an equally bright spotlight on underlying causal structure. Through a series of studies we have been able to show that certain factors highlight discipline-like relations.

The actual form of posing such questions makes quite a difference. For example, the 'why' part of the questions and the division of labor framing may collectively have a strong influence on judgments of clusters. In the tasks described earlier, the framing has usually been of the form:

> X knows why P1
> What else is X more likely to know?
> Why P2?
> *or*
> Why P3?

Consider now a triad that strips away both the 'why framing' and the question about expertise and simply presents the phenomena:

> P1
> Which is more similar to P1?
> P2
> *or*
> P3?

This second triad would seem to be simpler, and yet in tasks with both adults and children the tendency to cluster on disciplinary grounds drops considerably as other ways of clustering knowledge such as by goals or surface topics become more prominent. There are, of course, many different dimensions of similarity along which phenomena can be compared and when the raw phenomena are presented the discipline-based dimension is not especially salient. One can cluster on surface perceptual similarity of phenomena, on the basis of common lexical items or on the basis of any number of other dimensions. Embedding phenomena in frames that ask about people's 'why knowledge' tends to highlight the underlying causal principles. For example, if adults are presented with the following triad in stripped away form, they may be close to chance levels in clustering either P2 or P3 with P1. By contrast, when the same three phenomena are embedded in a 'X knows all about why' context, there is a strong preference to cluster P3 with P1. Knowing why a phenomenon occurs highlights the core causal processes responsible for that phenomenon in ways that most other contexts do not.

(P1) *A big, heavy boat takes a really long time to stop*
(P2) You can't understand two friends talking at the same time
(P3) You can bounce a basketball better on the street than on grass

Other factors can also enhance a focus on underlying causal processes. There is an advantage in posing the question as one of information-seeking, as in 'You want to know more about why P1: who would be a better person to ask, a person who knows why P2 or a person who knows why P3?' That way of framing the question, which seems to make it more immediately relevant to a participant, shifts children to even higher levels of discipline-based sortings (Danovitch and Keil 2004). As mentioned earlier, posing the question about higher-level categories, such as animals as opposed to ducks, also shifts participants more towards discipline-based clusters.

Thus, asking about the division of cognitive labor with a special focus on why-questions, using more high-level categories, and posing the questions in terms of personal information-seeking, all tend to focus the lens of similarity on the dimension corresponding to underlying causal relations. All these factors enhance performance in children at least as young as 5 years of age. Moreover, young children find it very natural to make judgments about who knows what based on an initial piece of knowledge. Many facets of meta-cognitive awareness, such as about the limits of one's memory and attentional processes, develop quite late; but a sense that knowledge is clustered into different domains in other minds emerges early and is robust.

TO WHAT END?

Why should young children be so adept at thinking about the division of cognitive labor and why should they show some ability to detect underlying causal relations and use them as a basis for thinking about expertise? Put differently, to what end do they use their sense of the division of cognitive labor? We do not yet know the full answer to this question; but there are some indications of potential uses that help us understand why children are sensitive to the different forms of expertise.

One use may be in evaluating the quality of potential experts. A series of studies in progress is exploring the idea that when children seek out new information, they use their notions of the division of cognitive labor to decide which individuals or sources to approach for new information. A child is told about two self-proclaimed experts. One claims to know a great deal about three phenomena, one from physics, one from economics, and one from psychology, while another claims to know a great deal about three phenomena from physics. Very preliminary evidence suggests that quite young children may know that the first 'expert' is much less plausible than the second. Thus, even young children may have doubts about the likely expertise of a 'Renaissance person'.

A second more direct use of divisions of cognitive labor is to know who to ask for further information or help on a topic. Even preschoolers may seek out different teachers for different problems, even when the problems are novel and don't simply match old ones that certain teachers have solved on prior occasions. When faced with several different adults to approach for information or for a problem solution, it can be very helpful to consider what proven areas of knowledge each of those adults already have. As we have seen, younger children might use different and sometimes misleading heuristics for seeking out the best experts, but in many cases they will do far better than chance. In a similar vein, when children hear bits of conflicting information from different adults they may use their sense of the legitimate division of cognitive labor to weigh the quality of the information that they hear. Thus, if a series of statements from one individual does not cohere as a natural domain of knowledge, a particular fact in that series may be discounted more than the same fact embedded in a series that is more coherent.

There may be a more important and subtler use, however, that is seen in groups at all ages. A sense of the division of cognitive labor provides confidence about one's current knowledge. The vast causal complexity of the natural and artificial worlds makes it impossible for any one person to have much more than the shallowest grasp of causal structure in a domain (Wilson and Keil 1998). Although there is evidence that people delude themselves in thinking that they understand such causal relations in far more detail than they really do (Rozenblit and Keil 2002), they are nonetheless also aware of at least some of the gaps in their knowledge.[2] A grasp of the division of cognitive labor enables them to feel that their knowledge is well grounded to the extent that there are legitimate experts who, collectively, could provide additional supporting information that could fill in the gaps. This form of support is closely related to how we might rely on the division of linguistic labor. If I believe that a panda bear is a particular kind of bear and label it as such, I may have considerable confidence about that belief because I have heard biologists state that DNA analyses show a clear pattern of commonality with other bears as opposed to other species.

My confidence arises from my sense of how knowledge in the science is distributed, a sense of the modern discipline of biology, and of the central role of microstructural properties such as DNA to understanding species. This idea of experts in biology is not restricted to those who encountered such concepts late in high school or college. It is accessible in a rough manner to surprisingly young children. Across a wide range

[2] There is a tendency to grossly overestimate one's causal explanatory understanding of both devices and natural phenomena. Whether it is everyday objects as simple as a zipper or a flush toilet or more complex objects such as a helicopter, adults and children alike think they have far more detailed understandings of the mechanism than they really do. People's initial ratings of what they know drop sharply after they are asked to actually provide explanations. This 'illusion of explanatory depth' is specific to estimates of how well one understands how things work. In contrast, people tend to be quite well calibrated in their estimates of how well they know facts, procedures, or narratives (Rozenblit and Keil 2002).

of ages, it may guide the strength of our beliefs and the extent to which we are willing to revise those beliefs and be persuaded by others.

Cross-cultural investigations of people's notions of the division of cognitive labor are just beginning and will be an important way of examining the extent to which the causal structure of the world drives intuitions about who knows what. The developmental studies suggest that there may be a striking universality of intuitions about clustering of why knowledge of everyday phenomena. Thus, even in traditional societies that have never had any exposure to the Western sciences, there may be a shared sensitivity to clusters of causal patterns that are used to infer clusters of knowledge. The causal patterns are relatively invariant across cultures; and if they are an important source of information for intuitions about expertise, they should cause a convergence on beliefs about relevant experts. Clustering of knowledge on the basis of category association, access, and goals, however, may show far more cultural variation. All three of those factors can be heavily influenced by culture and language. Discipline-based ways of thinking about expertise may therefore be the most robust and constant across cultures. This prediction poses a challenge to views that the domains of inquiry of the natural and social sciences are largely socially constructed.

In short, in all cultures, we come to depend on the knowledge of others. The division of cognitive labor is an essential infrastructure that allows us to transcend the very limited understandings that exist in the mind of any one individual. To benefit from the division of cognitive labor, however, we need ways of thinking about domains of expertise that can be used to tap into that expertise when needed. There are several distinct heuristics that can be used to figure out who knows what. Although there are major developmental changes in which heuristics are preferred, very young children are sensitive to many of these heuristics, including one that refers to the underlying causal patterns responsible for large classes of phenomena. At all ages, these heuristics provide a rudimentary sense of domains of expertise that can be used to evaluate the quality of new information. Thus, an important basis for doubt lies in our patterns of deference to others, patterns that heavily influence our deliberations throughout much of our development.[3]

[3] Much of the research reported on in this article was supported by NIH Grant R37-HD23922 to Frank Keil. Many thanks to Tamar Gendler for extensive comments on earlier drafts of this paper.

REFERENCES

Boyd, Richard (1999) 'Homeostasis, Species, and Higher Taxa', in R. Wilson (ed.), *Species: New Interdisciplinary Studies* (Cambridge), 141–85.

Burge, Tyler (1979) 'Individualism and the Mental', in Peter French (ed.), *Midwest Studies in Philosophy*, iv. *Studies in Metaphysics* (Minneapolis), 73–122.

Cartwright, Nancy (1999) *The Dappled World* (Cambridge).

—— (2003) 'Causation: One Word; Many Things', *Causality: Metaphysics and Methods Technical Report*. Centre for the Philosophy of Natural and Social Science, London School of Economics (London).

Danovitch, Judith, and Frank Keil (2004) 'Should you ask a Fisherman or a Biologist? Developmental Shifts in Ways of Clustering Knowledge', *Child Development*, 75: 918–31.

Dupre, John (1981) 'Natural Kinds and Biological Taxa', *Philosophical Review*, 90: 66–90.

Durkheim, Emile (1947) *The Division of Labor in Society*, tr. George Simpson (New York).

Fodor, Jerry (1974) 'Special Sciences', *Synthese*, 28: 97–115

—— (1998) *Concepts: Where Cognitive Science went Wrong* (New York).

Goldman, Alvin (1999) *Knowledge in a Social World* (Oxford).

—— (2001) 'Social Epistemology', *Stanford Encyclopedia of Philosophy* (http://plato.stanford.edu/entries/ epistemology-social/).

Green, Donald, Bradley Palmquist, and Eric Schickler (2002) *Partisan Hearts and Minds: Political Parties and the Social Identities of Voters* (New Haven).

Hacking, Ian (1999) *The Social Construction of What?* (Cambridge).

Hume, David (1739) *A Treatise of Human Nature* (Oxford).

Hutchins, Edwin (1995) *Cognition in the Wild* (Cambridge).

Keil, Frank (2003*a*) 'Folkscience: Coarse Interpretations of a Complex Reality', *Trends in Cognitive Sciences*, 7: 368–73.

—— (2003*b*) 'That's Life: Coming to Understand Biology', *Human Development*, 46: 369–77.

—— and Leonid Rozenblit (1997) 'Knowing Who Knows What', presented at the 1997 Meeting of the Psychonomics Society, Philadelphia.

Kitcher, Philip (1990) 'The Division of Cognitive Labor', *Journal of Philosophy*, 87: 5–22.

—— (2001) *Science, Truth and Democracy* (New York).

Kukla, Andre (2000) *Social Construction and the Philosophy of Science* (London).

Landauer, Thomas, and Susan Dumais (1997) 'A Solution to Plato's Problem: The Latent Semantic Analysis Theory of Acquisition, Induction, and Representation of Knowledge', *Psychological Review*, 104: 211–40.

Lutz, Donna, and Frank Keil (2002) 'Early Understanding of the Division of Cognitive Labor', *Child Development*, 73: 1073–84.
Mervis, C. B., and M. A. Crisafi (1982) 'Order of Acquisition of Subordinate, Basic, and Superordinate Categories', *Child Development*, 53: 258–66.
Murphy, Gregory (2003) *The Big Book of Concepts* (Cambridge).
Niklas, Karl (1996) 'How to Build a Tree', *Natural History*, 105: 48–52.
Prinz, Jesse (2002) *Furnishing the Mind: Concepts and their Perceptual Basis* (Cambridge).
Putnam, Hilary (1975) 'The Meaning of ''Meaning''', in Keith Gunderson (ed.), *Language, Mind, and Knowledge*, ii (Minneapolis), 131–93.
Rosch, E., C. Mervis, W. Gray, D. Johnson, and P. Boyes-Braem (1976) 'Basic Objects in Natural Categories', *Cognitive Psychology*, 8: 382–439.
Rozenblit, Leonid, and Frank Keil (2002) 'The Misunderstood Limits of Folk Science: an Illusion of Explanatory Depth', *Cognitive Science*, 26: 521–62.
Smith, Adam (1976) *An Inquiry into the Nature and Causes of the Wealth of Nations*, ed. R. H. Campbell and A. S. Skinner (Oxford).
Wilson, Robert, and Frank Keil (1998) 'The Shadows and Shallows of Explanation', *Minds and Machines*, 8: 137–59.