



cambridge.org/bbs

Elizabeth S. Spelke

Department of Psychology, Center for Brains, Minds, and Machines, Harvard University, Cambridge, MA, USA

## Précis

**Cite this article:** Spelke ES. (2024) Précis of *What Babies Know*. *Behavioral and Brain Sciences* **47**, e120: 1–81. doi:10.1017/S0140525X23002443

Target Article Accepted: 14 April 2023  
Target Article Manuscript Online: 30 May 2023  
Commentaries Accepted: 23 October 2023

### Keywords:

cognition in infancy; cognitive development; core knowledge; developmental cognitive science; origins of knowledge

**What is Open Peer Commentary?** What follows on these pages is known as a Treatment, in which a significant and controversial Target Article is published along with Commentaries (p. 18) and an Author's Response (p. 72). See [bbsonline.org](https://bbsonline.org) for more information.

### Corresponding author:

Elizabeth S. Spelke;  
Email: [spelke@wjh.harvard.edu](mailto:spelke@wjh.harvard.edu)

### Abstract

Where does human knowledge begin? Research on human infants, children, adults, and non-human animals, using diverse methods from the cognitive, brain, and computational sciences, provides evidence for six early emerging, domain-specific systems of core knowledge. These automatic, unconscious systems are situated between perceptual systems and systems of explicit concepts and beliefs. They emerge early in infancy, guide children's learning, and function throughout life.

Young children may be the most effective learners on earth. Over their first 6 years, they learn their native language and develop a commonsense understanding of the places through which they move, the objects they manipulate, and many of the actions, customs, habits, beliefs, and values of the people around them. They also learn new concepts of number, geometry, and mental states that expand their reasoning in these domains. Cultures and technologies vary greatly, but children's learning succeeds in diverse environments. Very little of this learning comes through explicit teaching. How do children do this?

The question is open, but research within the increasingly connected fields of experimental psychology, systems and cognitive neuroscience, and computational cognitive science suggests some answers. In *What Babies Know* (Spelke, 2022), I focus on the emergence of knowledge from birth to 1 year, and I offer one partial answer to this question: Children learn fast and flexibly, because they are endowed with at least six cognitive systems that capture fundamental properties of the things they learn about. Core knowledge of places, objects, agents, social beings, number, and geometry supports children's learning in these domains, both in infancy and beyond. It also supports their language learning and the development of new concepts that span the domains.

Core knowledge systems apply to diverse things: A place may lie on a mountain or a shopping mall; an object may be a car or grape; an agent may be a person or a hen; a social being may be the infant's father or an animated ball with a cartoon face. But each core knowledge system centers on the abstract, persisting and interconnected properties possessed by all the entities in its domain: The distances and directions of places from one another; the solidity and continuous motions of inanimate objects; the causal powers, efficiency and goal-directedness of agents and their actions; and the shareable experiences of social beings who engage with one another and form enduring relationships. Core knowledge therefore supports learning in any habitable environment.

Core knowledge systems have further properties in common. First, all occupy a middle ground between perception and belief. Like perceptual systems, they are functional at birth, and they operate automatically and unconsciously when one attends to entities in their domain. Like belief systems, they center on abstract concepts that support actions on, and inferences about, properties of the world that cannot be perceived directly, such as the solidity of an untouched object, the direction of a far-away destination, or the intentions of an actor. Because these properties are useful in all environments and at all ages, core knowledge is present and functional throughout life and provides our species' most basic common ground. Because it is unconscious, however, people rarely are aware of the universal foundations of our diverse beliefs and opinions.

Second, where tested, core knowledge systems have been found to function in the same ways, and to activate homologous brain systems, in humans and diverse animals. Many discoveries concerning the neural mechanisms of human navigation, for example, were sparked by findings from studies of rodents, who navigate in similar ways. These findings suggest that core knowledge systems are ancient: They emerged in ancestors common to many animals, from humans to monkeys, rodents, birds, fish, and possibly beyond. Research revealing core knowledge therefore overturns the common view that cognitively simpler creatures, like guppies or newborn mice, can sense their surroundings but lack our abstract concepts. Contrary to that view, the evidence reviewed in this book suggests that animals are more likely to share our most important abstract concepts than our more specific sensory experiences. Research using animal models therefore affords deeper study of the endogenous processes and prenatal experiences that give rise to human knowledge.

In this précis, as in the book, I first describe the insights, from the study of visual perception, that have led to the discovery of core knowledge (sect. 1). Then I turn to research probing the origins of knowledge of objects, places, and number (sects. 2–4), followed by a general discussion of the properties these systems share (sect. 5). I turn next to research on people as agents with causal powers and as social beings with shareable experiences (sects. 6 and 7). I end with research on infants' language learning (sect. 8) and their changing conceptions of the people who use language to share their experiences with others, including the infant (sect. 9). The developments described in these two sections may herald the emergence of a uniquely human learning process that carries infants beyond core knowledge.

## 1. Vision

When experimental psychologists began to probe the minds of infants in the late 1950s and 1960s, most believed that knowledge was the product of innate sensory systems, learned belief systems, and nothing more. They focused, therefore, on infants' capacities for perception and learning, and they developed the approaches and methods that have brought other cognitive systems into view. The book's first chapter introduces these approaches and methods by reviewing centuries of thinking and research on the nature, origins, and development of visual perception, from Descartes and Berkeley to Helmholtz to the present. It devotes most space to three twentieth century scientists – Eleanor Gibson, Richard Held, and David Marr – whose work provided the perspectives and methods (respectively, from experimental psychology, cognitive neuroscience, and computational cognitive science) that research on cognition in infancy builds on. Because insights into the mind tend to come first from studies of behavior, I limit my discussion in this précis to Gibson's transformative behavioral experiments probing the emergence and nature of the mechanisms by which human adults, human infants, and adult and infant animals of diverse species perceive the stable, three-dimensional (3D) visible surface layout.

To explore surface perception in adults, Gibson brought classical psychophysical experiments out of the laboratory and tasked adults, with no training in psychology or the vision sciences, with judging the absolute and relative distances and sizes of objects at diverse locations on open fields (Gibson & Bergman, 1954/1991). Their judgments were strikingly accurate, and they depended on patterns of optic flow, produced as the participants moved. To extend these tests to animals and to crawling human infants, Gibson invented the “visual cliff” (Fig. 1a): A structure with a raised central platform bordering two visible surfaces. Both to protect the participants and to remove all tactile information for depth, these surfaces were viewed through a glass plate that covered the array, positioned slightly below the platform on both sides and allowing

participants, with no training, to move at will in any direction. With clever lighting, the glass was invisible: When participants looked through it, from the platform, they saw one flat, textured surface, positioned far below the glass, and one flat, textured surface positioned a single safe step down. Gibson relied on the propensity of mobile animals, including crawling infants, to explore their surroundings spontaneously, and she observed their choices of visible surfaces to step on.

Gibson found that all the tested animals that walk on land avoided the apparently deep side of the cliff and crossed readily onto the apparently shallow side. For all these animals, moreover, locomotor choices were guided by patterns of optic flow, as they were in the open-air psychophysics experiments with human adults (Fig. 1b and caption). Gibson's findings provided evidence that depth perception is present in untrained animals and suggested that it depends on a common mechanism, attuned to changes in the light at the eyes as animals move over the ground.


With these tools, Gibson and her collaborators studied the origins of depth perception and found that it is innate. It was exhibited not only by newborn goats, who began to move immediately after birth when placed on the ground but not when placed on a small, elevated platform (Gibson, 1980), but also by dark-reared rats and cats whose only experience of visible surfaces occurred on the cliff device itself, where the protective glass ensured that the apparently deep and shallow sides were in fact equally close, safe, and traversable. The rats avoided the cliff on first exposure to the light, whereas the cats required a few days to adjust to a lighted environment, and initially, they walked on both sides of the platform. As their vision improved, however, the cats began to avoid the deep side, even though all their experiences implied that it was safe (Walk & Gibson, 1961/1991; Walk, Gibson, & Tighe, 1957/1991; Fig. 1c). Crawling human infants also refused to cross on the side of the visually distant surface, even when called by a parent, who tapped on the glass surface and assured the infant that it was safe (Fig. 1d).

Gibson's studies of the development of surface perception also focused on other modes of exploration. Because young human infants cannot locomote independently, she and her students tested their perception of 3D surfaces by measuring their head movements in response to optic flow fields specifying either an approaching window or an approaching surface whose borders were matched to the window in size, shape, and movement. As the window approached, 3-month-old infants, who cannot yet reach for objects, leaned forward and to the sides to explore the scene that it progressively revealed. As the surface approached, in contrast, they moved their heads backward to keep more of the scene in view and avoid a possible collision (Gibson, 1982/1991).

Like others before her, Gibson also measured the duration and direction of young infants' visual exploration of events, taking advantage of the fact that infants, like their elders, tend to look longer when visible events undergo informative changes. Her experiments showed that 5-month-old infants reacted with longer looking when a surface's visible movement changed from rigid to nonrigid than when it changed from one rigid motion to another (Gibson, Owsley, & Johnston, 1978/1991). Moreover, 1-month-old infants reacted to this change in the motion of a visible object when their encounter with the original motion occurred only by touch, as they sucked on a rigid or nonrigid nipple (Gibson & Walker, 1984/1991).

With this research, Gibson opened a door to the modern cognitive, brain, and computational sciences. She showed that rigorous psychophysical methods that had been argued, since

DR. ELIZABETH SPELKE is the Marshall L. Berkman Professor of Developmental Psychology at Harvard University (2001–present) and an Investigator at the Center for Brains, Minds, and Machines, Cambridge, MA. She holds a BA from Radcliffe College and a PhD from Cornell University. She has previously held positions at the University of Pennsylvania (1977–1986), Cornell University (1986–1996), and MIT (1996–2001).



**Figure 1.** (a) The original visual cliffs used for testing dark-reared rats (left) and goats (right). (b) Multiple cues distinguished the two sides of the apparatus (including texture density, pictured on the left), but animals and human infants navigated primarily by optic flow: the depth-dependent displacements of edges projected to the eye of a walking animal or crawling infant. (c) Dark-reared cats crossed equally onto the deep and shallow sides on first exposure to light but came to avoid the deep side as their vision improved, despite the consistent and equal safety of the two sides. (d) Human infants typically refused to crawl onto the deep side, even when a parent encouraged them to do so. All figures are reprinted from Walk & Gibson (1961/1991).

Helmholtz, to be applicable only to vision scientists and their highly trained graduate students, testing themselves in a laboratory, could be conducted with equal rigor on humans and other animals of diverse ages, skills, and experiences as they explored rich, natural environments. Participants in psychophysical experiments need not undergo training in systematic introspection or be capable of verbal reports, if experiments leverage their intrinsic motivation to explore the things and places they encounter.

Because visual exploration begins at the time of eye opening, it provides an especially useful window on the innate foundations of perceptual development, the course of perceptual learning, and the functioning of the developing human brain. Indeed, astute psychophysical studies, focused on the emergence of diverse visual functions including color vision and motion sensitivity, have used the simplest possible indicator of visual exploration in human infants: As each infant explores pairs of visual displays, projected side by side, on each of a series of brief trials, stimulus-blind observers judge which display the infant looks at more (Teller, 1979). Using this method, Richard Held and collaborators presented infants with displays viewed through stereo glasses while systematically altering the sizes and orientations of the resulting binocular disparities (to dissociate effects of perceiving depth from effects of detecting disparate images). Their research

generated beautiful psychophysical functions marking the onset, development, and acuity of stereopsis for each individual infant. They found that most infants began to look longer at the surfaces that they perceived to vary in depth between 10 and 20 weeks of age, marking the onset and rapid subsequent development of this cortical function (Held, Birch, & Gwiazda, 1980). Research by Rachel Keen Clifton, focused on auditory perception and exploratory head turning, provided evidence for the emergence of the cortical functions underlying auditory localization over the same ages (Clifton, Morrongiello, & Dowd, 1984).

Recent experiments in visual neuroscience, using optical imaging to chart the activity of large populations of cells in fetal mice, have found patterns of activity, generated in retinal ganglion cells and projected to the superior colliculus, a subcortical brain structure, that simulate the optic flow patterns that guide the locomotion of the adults, infants, and animals in Gibson's experiments (Ge et al., 2021). As in earlier research documenting spontaneous activity in the visual system prior to the onset of vision (e.g., Katz & Shatz, 1996), these activity patterns arise before the onset of visual experience. They suggest a mechanism by which innate capacities for depth perception might either emerge or be sharpened, prior to an animal's first encounters with a visible environment.




With this research, investigators of the development of visual perception provided the methods that fill this book. Throughout her work, Gibson argued that people and animals of all ages are motivated to explore their surroundings, and their exploratory behavior both reveals what they perceive and forecasts what they will learn (Gibson, 1969, 1991). Active looking and listening are likely the most important means for exploring and learning about the environment for young human infants who cannot yet grasp objects or locomote; since Gibson's seminal studies, exploratory looking has been extensively used to probe what infants see (Gredebäck, Johnson, & von Hofsten, 2010). Indeed, even unsuccessful attempts at exploration, such as newborn infants' failed attempts to contact a visibly moving object that stands just beyond their reach, reveal their perception of the object's distance, direction, and motion (von Hofsten, 1982). These behaviors also shed light on the workings of cognitive systems beyond perception.

## 2. Objects

Chapter 2 focuses on the research of many investigators probing infants' knowledge of objects, beginning with my personal trajectory. When my research began, I believed I was studying an

intriguing aspect of visual perception: The origins of perception of the complete shapes of objects which – unless they are transparent – are never fully in view at any given time. Guided both by Gibson's work on surface perception and by the work of the neo-Gestalt psychologist, Albert Michotte (Michotte, Thines, & Crabbé, 1964), experiments in my lab probed infants' perception of objects that are partly hidden by other objects, objects that stand adjacent to or in front of other objects, and objects that move in and out of view. These studies focused on two exploratory behaviors – selective looking at and reaching for objects – that yielded converging findings. Their findings, however, did not accord either with the perceptual theories they were designed to test or with the research on surface perception that Gibson had pioneered. Here I give two examples.

First, studies led by Philip Kellman used Gibson's looking time method to probe 4-month-old infants' perception of the complete shapes of objects that are partly occluded. After infants' looking time to a center-occluded, straight rod had declined, the occluder was removed to reveal one connected rod or two short rods separated by a gap, on alternating trials (Fig. 2a). To our surprise, infants' looking rose equally for the two test displays, suggesting that they had not perceived the original rod either to end at the edge of the occluder or to be connected behind it. Further studies



**Figure 2.** Some strengths and limitations of young infants' representations of objects. (a) After a decline in looking time at an occlusion display (top figures), infants' looking times to two unoccluded displays (bottom figures) were compared. Red symbols indicate the direction of longer looking, indicative of a larger experienced change in the array. (b) While moving in a chair (i, bottom arrows), infants saw a stationary or moving rod that maintained a stable position in the infant's visual field (top dots and arrow). Their looking time (ii) was measured on the last six habituation trials and the first test trial presenting one connected rod (closed circles) and two rod fragments (open circles), undergoing the same motions as in habituation. Infants perceived a connected object when the rod itself moved and not when it was stationary, regardless of whether its image was displaced or stabilized in the infant's visual field. After Kellman and Spelke (1983) and Kellman et al. (1987). (c) When two familiar objects differing in shape, colors, texture, and affordances alternately appeared from behind a single screen (top four images) and then the screen was removed, infants treated the reappearance of one and of both objects (bottom two images) as equally novel or surprising. After Xu and Carey (1996).

presenting partly occluded triangles, spheres, and cubes confirmed these findings, which were at odds both with Michotte's experiments on adults and with the Gestalt theories that inspired them.

In contrast, infants did perceive center-occluded objects as complete, connected bodies when their visible surfaces moved horizontally back and forth together, even as their centers remained hidden (Kellman & Spelke, 1983). In these cases, however, infants' behavior still did not in accord with the Gestalt principle of "common fate," because infants inferred a connection between the visible ends of the rod when the ends of the rod maintained a constant position in the infant's two-dimensional (2D) visible field and moved only in depth. Infants also inferred a connection between the ends of the rod when they and the rod were moved together, such that the rod's image remained centered in their visual field, while global changes in optic flow specified the changing positions of the infant and the object (Fig. 2b). These findings suggest that infants' perception of the object depended on the perceived 3D displacement of its perceived visible surfaces within the scene, not on the 2D displacement of its sensed image in a succession of static arrays (Kellman, Gleitman, & Spelke, 1987; Fig. 2b).

These findings soon were joined by the findings of many other experiments on human infants, including newborns, performed by diverse investigators including Renee Baillargeon, Rachel Keen, Francesca Simion, Arlette Streri, and Claes von Hofsten. They were joined, as well, by experiments on newly hatched or controlled-reared chicks, conducted by Giorgio Vallortigara and his collaborators. Among other findings, young infants represented the persistence of objects that moved fully out of view, the solidity of visible or hidden objects that collided with other objects, and the continuity of object motion through space and time (see Baillargeon & Carey, 2012, for a review).<sup>1</sup> Newborn infants and newly hatched chicks also showed these abilities, in experiments suggesting a subcortical origin to their processing of object motion. For example, newborn infants perceived the complete shape of a partly occluded, moving rod only when the rod underwent rapid stroboscopic motion: A signature of processing in the superior colliculus (Regolin & Vallortigara, 1995; Valenza, Leo, Gava, & Simion, 2006). I return to this finding below.

In all these studies, young infants failed to track objects under conditions in which older infants, children, and adults succeed. My second example focuses on one such situation. Fei Xu and Susan Carey presented 10-month-old infants with a large screen that hid two familiar objects – for example, a toy duck and a child's shoe – that alternately appeared and disappeared at opposite sides of the screen but were never visible at once (Fig. 2c, top four images). After viewing this event repeatedly, the screen was removed to reveal either both objects or just one object, on alternating trials (bottom images). Infants' looking times suggested no expectation that two objects would appear rather than one (Xu & Carey, 1996). Their lack of surprise was striking, because infants showed signs of detecting and remembering the differing forms and functions of the two objects during the occlusion events: They looked longer when the two objects differed (like the duck and shoe) than when they were indistinguishable (two shoes). Moreover, when two objects appeared in alternation from behind two different screens with a visible gap between them, infants looked longer when the screens were removed to reveal just one object (Spelke, Kestenbaum, Simons, & Wein, 1995). Infants track objects based on their spatiotemporal properties – the

continuous existence and motion of each object – but not in accord with the objects' differing forms and functions.

All these experiments suggest that infants' knowledge of objects is limited and depends on mechanisms that are distinct from those underlying perception of visible surfaces or object forms. Faced with this evidence, I once proposed, wrongly, that objects are not grasped by a perceptual system but by the only alternative of which I could conceive: A system of central cognition, like our systems of explicit reasoning about objects and their mechanical interactions (Spelke, 1988). Research by Brian Scholl and others provided decisive evidence against this proposal (Scholl, 2001): Adults were found to share the representational system found in infants, and we use it both unconsciously and under constraints on attention and working memory that do not limit our conscious reasoning about objects. Scholl concluded that object representation depends on visual mechanisms; I inched toward the notion of core knowledge, but with little idea how core systems, interposed between perception and thought, might operate.

Recent research by the computational cognitive scientists, Tomer Ullman and Joshua Tenenbaum (2020), provides a useful way of thinking about the core system of object representation in relation both to vision and to explicit thought. They propose that physical reasoning depends on a model of objects and their interactions like that of the physics engines – computer programs – used in interactive video games. At each time step, physics engines use an approximation to Newtonian mechanics to transform a representation of one 3D array of objects, each with a particular size, coarse shape, mass, position, and motion, into the 3D positions and motions of the objects at the next time step. By running such a process forward, infants may predict the future states of objects, including objects that move out of view. By inverting this process, using Bayesian inference, infants may recover object properties that cannot be directly seen, including an object's solidity, weight, and occluded location and motion (Smith et al., 2021; Ullman, 2015; Ullman, Spelke, Battaglia, & Tenenbaum, 2017; Ullman & Tenenbaum, 2020).

In contrast to core knowledge of object physics, visual perception may depend on a model of the world like that in the graphics engines used in many animated films. Graphics engines begin with a 3D representation of a scene and its light sources, and they generate 2D images of the scene from particular vantage points. Armed with a graphics engine and confronting the optic array that meets the eye, the visual system may invert this process by sampling different 3D scenes, simulating the images they would project at the eye, and using Bayesian inference to infer the most likely 3D surface arrangement that gave rise to the present 2D array, as Helmholtz once proposed. Graphics engines function together with physics engines in the production of animated video games, and cognitive systems with similar properties may function together in animals and humans to support the emergence, growth, and use of knowledge.

If Ullman and Tenenbaum are correct, core knowledge of objects takes the form of an internal model of the physical world, and it operates by simulating the sorts of interactions that occur as objects move and collide. Such a system would differ both from the visual system, whose internal models focus on the properties of light and of light-reflecting surfaces, and from the explicit systems of concepts and equations that students learn in physics classes. To my knowledge, Ullman and Tenenbaum have not discussed the processes by which the first model of objects might grow in the minds of inexperienced chicks and


newborn infants, but I hazard a suggestion: Like the spontaneous, prenatal, subcortical neural activity that simulates patterns of optic flow, preparing fetal mice for their future encounters with visible surface layouts, there may be prenatal, subcortical neural activity that simulates the movements and interactions of spatio-temporally continuous, solid 3D bodies, preparing infants for their first encounters with the objects.

### 3. Places

As studies of infants' knowledge of objects proceeded, Barbara Landau and Henry Gleitman piqued my interest in another ability that likely builds on, and goes beyond, visual perception of the surrounding 3D layout. While studying the language learning of a young blind child, Landau observed that the child appeared to know the spatial layout of her home. To investigate the sources of this ability, she presented the child, and groups of sighted but blindfolded children, with a set of small-scale navigation tasks. First, the child was led from her mother to each of three objects occupying different locations in a single room, returning to her mother after visiting each object. Then the child was taken to one of the objects and encouraged to walk directly to another object. On arriving at the first object, she turned herself to face in the approximate direction of the next object and walked the approximately correct

distance to get to it: a novel path for her (Fig. 3a). Landau concluded that the child's acts of navigation depended on a representation of the geometry of the 3D array: Like explorers, the child represented the directions and distances that she had traveled and inferred the final distance and direction to be traveled, in accord with basic theorems of Euclidean geometry (Landau, Gleitman, & Spelke, 1981; Landau, Spelke, & Gleitman, 1984).

While we were studying the blind child, Cheng and Gallistel (1984) and Cheng (1986) showed that the representations that guide navigation by animals (in their studies, rats) are limited, relative to those guiding the explicit planning of human explorers. When rats were disoriented after observing the location of food within a rectangular room containing multiple landmarks, they subsequently reoriented themselves and searched for the food in accord with its distance and direction from the walls bordering the room, but not in accord with the room's most obvious landmarks, such as a single white wall in otherwise black surroundings or an odor emanating from a source near the food's location. Faced with these findings, Cheng and Gallistel proposed that the rats' reorientation depended on a "geometric module." Their hypothesis was stunningly corroborated by later research on mice, who use a nongeometric landmark to determine which of two rooms they are in but not to reorient themselves within that room, all within the same trial (Julian, Keinath, Muzzio, & Epstein, 2015).



**Figure 3.** Overhead view of displays from experiments testing the information used by young navigating children. (a) An oriented but blind child first walked between her mother and each of three objects (left figure, dotted arrows) and then was led to one object and encouraged to move independently to another object (left figure, solid arrows). The smaller figures show her paths of unaided (solid lines) and guided motion (dotted lines; Landau et al., 1981). (b) Children searched for a hidden object after they were turned slowly with eyes closed in a rectangular chamber, either repeatedly but become disoriented (top figures) or for one partial turn so as to remain oriented (bottom figures). In both conditions, children's eyes were closed while the two corner boxes were moved to new positions, dissociating the boxes' features from their relations to the bordering walls. Blue arrows indicate the hiding location and red stars indicate where children searched for the object. Disoriented children used room geometry to relocate the object, whereas oriented children used features of the boxes. Children showed these effects on the first trial, with no expectation that they might be disoriented, indicating that they encoded the box features in both conditions but used them to guide their search only when oriented. Drawing created by Kirsten Condry. In (c), dots of different sizes and spacing (left) influenced reorientation in a square room, though other pattern differences (e.g., dots vs. uniform gray) did not, as indicated by the figure on the right. Views of the room's corners suggest that differences in dot size and spacing produced an illusion of depth (left image), whereas other differences in patterning did not (e.g., right image). A further experiment confirmed the depth interpretation by revealing that the same differences in dot size and spacing did not elicit reorientation in a slightly elongated room with the larger dots on the more distant walls. After Lee, Winkler-Rhoades et al. (2012).

Young human infants do not navigate on their own, but infants of other species do, and experiments shed light on the origins and nature of the mechanisms that guide them. Chapter 3 discusses this research, as well as studies of young children. Linda Hermer and I adapted Cheng and Gallistel's methods for studies of 18-month-old toddlers, in hopes of discovering why humans navigate so much more flexibly than rats. To our astonishment, toddlers behaved like the rats: They navigated only by the distances and directions of the room's borders when they were disoriented, ignoring brightly colored walls, attractive toys, or distinctive decorations at the room's corners, though they used such landmarks when they were oriented (Hermer & Spelke, 1994, 1996; Fig. 3b). Later studies by Sang Ah Lee revealed that children's reorientation process was guided by the perceived distances and directions of the walls that bordered the floor: Children successfully reoriented in a square room whose patterning on the walls induced an illusion that the room was slightly rectangular, whereas they failed to reorient in a slightly rectangular room in which the same patterning induced an illusion that the room was square (Lee, Sovrano, & Spelke, 2012; Lee, Winkler-Rhoades, & Spelke, 2012; Fig. 3c).

In contrast to children and rats, human adults used both geometry and landmarks when tested under similar conditions, and they mentioned the blue wall when asked why they had searched in a particular location. When the room was devoid of landmarks, however, many adults reported that they simply guessed which of the four corners contained the hidden object, even though they searched only corners at the appropriate distances and directions. Behavioral and neuroimaging experiments on navigation in virtual environments revealed that adults activate separate systems for navigating using layout geometry, and for piloting to a previously visited location using landmark objects. Only the latter system is associated with the deployment of attention to features of the room (Doeller & Burgess, 2008; Doeller, King, & Burgess, 2008), although the system for representing places is activated only when people and animals plan and carry out their own acts of navigation, not when they are passively moved (e.g., Javadi et al., 2017).

Thus, the geometric representations guiding reorientation are modular, in Fodor's (1983) sense, for humans as well as rodents. Even as adults, we reorient ourselves unconsciously, guided by a tiny subset of the environmental features that we perceive and remember. Modular cognitive systems, Fodor argued, are the antithesis of belief systems: He thought they were "input systems" like vision, although navigation draws on capacities for learning, memory, and action planning. The second half of the chapter focuses on the remarkable convergence of these findings with the findings from studies of navigation and place representations in the brains of animals and people (e.g., O'Keefe & Burgess, 1996), and from studies assessing the efficacy and limits of artificial navigation systems in autonomously moving robots with no navigation aids (e.g., Thrun, 2002).

The studies from cognitive and computational neuroscience strengthen the evidence for an encapsulated navigation system that operates automatically and unconsciously, but beyond the limits of perception. Navigation depends on the hippocampus: A structure that is intimately involved in action planning, learning, and conscious episodic memory. It is fostered, moreover, by generative processes for simulating and comparing different possible routes through an environment as an aid to action planning (Foster, 2017). As adults, we sometimes plan actions consciously, but research using functional brain imaging reveals

processes of mental simulation that are far more rapid and unconscious, both in animals and in human adults. When adults rest in the middle of a real or virtual navigation task (or, indeed, a non-spatial task requiring "navigation" through a complexly structured array of task conditions), they report no awareness of the rapid simulation processes that neuroimaging reveals, but those processes are predictive of improved performance on subsequent task sessions (e.g., Javadi et al., 2017; Liu, Dolan, Kurth-Nelson, & Behrens, 2019; Shuck & Niv, 2019). Research on navigation and spatial memory provides the strongest evidence for a system that combines some of the features of perceptual systems with some of the features of belief systems. It also provides evidence for unconscious processes of mental simulation, both in adults and in inexperienced animals (e.g., Farooq & Dragoi, 2019).

#### 4. Number

Chapter 4 focuses on infants' sensitivity to number. One source of numerical information comes from the core object system: The representations that support tracking objects over occlusion and reasoning about their interactions are leveraged by older children and adults to support rapid determination of the exact number of objects in an array, up to a capacity limit of about three (Alvarez & Franconeri, 2007). This chapter focuses primarily on an earlier emerging source of numerical information: A system for representing, imprecisely, the relative numerosities of sets of objects or events, encountered in any perceptual modality, and for combining or dividing these sets in accord with the operations of arithmetic. Because this system operates on sets with ratio-limited accuracy, it has been dubbed "the approximate number system," or ANS.

The ANS is present and functional in newborn infants, who selectively look at visual-spatial arrays that roughly match the number of sequential sounds in a simultaneous auditory sequence (Izard, Sann, Spelke, & Streri, 2009; Fig. 4a). At birth, it is highly imprecise: Newborn infants match arrays of 4 or 12 visible objects to auditory sequences of the same number of syllables, but not arrays and sequences that contrast four with eight objects and syllables. Six-month-old infants succeed with the latter 2:1 ratio (Xu & Spelke, 2000) and can mentally transform visual arrays in accord with the operations of addition and subtraction: If an array of 10 objects is hidden behind a screen and then five of the objects move into view, infants expect approximately five objects to remain behind the screen, looking longer if the raising of the screen reveals 10 objects (McCrink & Wynn, 2004). The precision of the ANS increases throughout the first year (e.g., Xu & Arriaga, 2007) and beyond (Halberda, Mazzocco, & Feigenson, 2008), sharpens over the course of math instruction (Piazza, Pica, Izard, Spelke, & Dehaene, 2013), and varies with children's and adults' proficiency at mathematical learning and reasoning (Halberda et al., 2008; Halberda, Ly, Wilmer, Naiman, & Germine, 2012).

It is difficult to determine either the innateness of the ANS or its core function, because it applies to sets of diverse entities. Research by Fei Xu, conducted with considerably older infants, suggests that it functions, in part, to support infants' learning about the statistical properties of things and events. In one series of studies (Xu & Garcia, 2008), infants who viewed a transparent box containing mostly white but a few red balls looked longer if a person who blindly fished for balls in the box retrieved mostly red ones. Numerical estimation processes may foster infants' learning and reasoning in uncertain and variable environments.







experiments, functional magnetic resonance imaging (fMRI) and EEG, provided extensive evidence for representations of number that are unconscious yet demanding of attention to the entities (in their case, briefly presented number symbols; Naccache, Blandin, & Dehaene, 2002).

Like core object and place representations, therefore, ANS representations appear to arise from a core cognitive system that is distinct from our explicit integer concepts. Does the ANS support children's learning of mathematics? To date, evidence for effects of short-term training with numerical tasks that activate the ANS has shown some positive effects on children's subsequent symbolic math performance (Hyde, Khanum, & Spelke, 2014; Khanum, Hanif, Spelke, Berteletti, & Hyde, 2016; Park, Bermudez, Roberts, & Brannon, 2016). Nevertheless, preschool activities that exercise only the ANS have led to no long-term enhancement of children's subsequent learning of mathematics in school (Dillon, Kannan, Dean, Spelke, & Duflo, 2017). The core number system therefore is not the only cognitive system needed for learning of mathematics, but it likely contributes to mathematical reasoning and learning.

## 5. Core knowledge

Studies of early-emerging knowledge of objects, places, and number provide the clearest evidence for the existence and properties of core cognitive systems whose abstract content supports exploration and learning. For objects, the physical properties of cohesion, continuity, and solidity govern objects' movements and interactions and support learning about objects' forms and functions. For places, the geometric properties of distance and direction support learning about the navigable paths over the ground that connect out-of-view places. For number, the properties of order and composition likely support statistical learning about predictable objects, actions, and events, as well as children's learning of primary school mathematics.

Core systems have further properties: They are ancient, emerge early in life, and are invariant over later development, as evidenced by research revealing the same domain-specific abilities, limits, and signature patterns of neural activity in each core domain, across diverse species and ages. Core systems also are impervious to our explicit beliefs and are activated automatically and unconsciously when we attend to entities in their domain. All may place demands on attentional resources, if the core systems emerge and function as generative models that simulate either the behavior of entities in their domain (for objects and numbers) or the actions that can be performed in the domain (for places, acts of navigation).

In Chapter 5, I propose that these properties are related: Any cognitive system that has some of them is likely to have them all. An ancient system that first emerged in highly distant ancestors is likely to center on abstract content, because it had to be applicable to the diverse environments that the descendants of that last common ancestor came to inhabit. Moreover, a system of abstract concepts that supports exploration and learning in a broad range of habitable environments, for animals that vary in size and behave in different ways, is likely to be useful for people of all ages, whatever their circumstances and however those circumstances change with age and experience. It is likely, therefore, to function in all human cultures. To preserve its functionality over evolutionary time scales, unimpeded by mechanisms of top-down inhibition from later-emerging brain systems, such a system should be activated automatically, independently of volitional


control or conscious access. To function in the service of goal-directed behavior, it should be activated when a person or animal attends to entities in its domain.

Based on these considerations, I consider, in the next three chapters, whether the human mind contains other systems with this constellation of properties. I propose that three more core systems emerge in infancy and guide learning about entities that many animals must contend with. The system discussed in Chapter 6 has been least studied: It captures the forms and functions of objects of specific kinds, perhaps especially the living, inanimate beings (like trees) that grow in all habitable environments, whose distinctive forms allow for their recognition and categorization, and whose distinctive affordances for food and shelter bear on people's and animals' survival. For brevity, I won't discuss that system here. The other proposed systems focus, respectively, on animate beings that cause their own motion and plan efficient actions to achieve valued goal states, and on social beings who engage with one another, share experiences of attention and emotion with their social partners, and form enduring bonds.

## 6. Agents

Chapter 7 focuses on infants' knowledge of beings who sense their surroundings and generate their own movements. Like the movements of objects, the movements of agents are physically constrained: Agents cannot pass through walls or teleport. Unlike objects, however, agents' actions have unique and interconnected properties. Infants are sensitive to highly abstract properties of the actions of people, animals, and animated characters: They view their actions as intentional, goal-directed, perceptually guided, efficient, and causal. These properties are connected: Given evidence that an agent has caused a change in an inanimate object on contact, 3-month-old infants expect its future actions to be efficient (Liu, Brooks, & Spelke, 2019; Skerry, Carey, & Spelke, 2013; after Gergely, Nádasdy, Csibra, & Bíró, 1995; Fig. 5a), goal-directed (Woo, Liu, & Spelke, 2021; after Woodward, 1998), and perceptually guided (Choi, Mou, & Luo, 2018; after Luo & Johnson, 2009). Older infants infer that such agents value more highly the goals for which they undertake more costly actions (Liu, Ullman, Tenenbaum, & Spelke, 2017; after Jara-Ettinger, Gweon, Schulz, & Tenenbaum, 2016; Fig. 5b).

As in other domains, infants' knowledge of agents is limited; I mention two limits here. First, although newborn infants expect agents to engage in biological motion, they exhibit no specific knowledge of the attributes that distinguish the bodies of their own species from those of others: Newborn human infants are equally attentive to the biological motion of a human and a hen (Simion, Regolin, & Bulf, 2008; Vallortigara, Regolin, & Marconato, 2005). Young infants also have no knowledge of the most likely goals of the actions that human agents perform. Although older infants expect acts of reaching to be directed to objects rather than places (Woodward, 1998), in contrast to acts of locomotion (Hamlin, Wynn, & Bloom, 2007), 3-month-old infants lack these expectations (Sommerville, Woodward, & Needham, 2005), though they view agents' actions as goal-directed and quickly learn an actor's specific goal when given appropriate evidence (Woo et al., 2021). These findings suggest that an ancient system, common to humans and other animals and centering on abstract, general properties that apply to all actions, supports young infants' learning about people's distinctive appearance, actions, and goals.




**Figure 5.** Displays for experiments testing infants' expectations that agents will act efficiently to achieve valued goals. (a) Three-month-old infants repeatedly viewed an agent who reached for an object over a barrier and, in one condition, caused the object to light up on contact (top right image), until their looking time to this event declined. After removal of the barrier (middle and bottom right images), infants looked longer at the same indirect motion, now inefficient. This effect was abolished when the same videos were altered so that the object changed its state with no contact with the finger (top left images and data figure, after Liu, Brooks et al. (2019)). (b) Ten-month-old infants were presented with an agent who took a higher-cost action for one target than for the other (top four images). After looking time to these events had declined, they viewed a display in which the costs of reaching the two targets were equal. Infants looked longer when the agent approached the target for which it had taken a lower-cost action (bottom right image), providing evidence that they expected the agent to approach the target for whom it had taken the more costly action. After Liu et al. (2017).

Second, young infants display little understanding of social actions, like cooperation or shared attention to objects. A wealth of evidence (to which I return below) suggests that this understanding arises at the end of the first year, but for younger infants, it is conspicuously absent. For example, if 10-month-old infants first view two people who stand side by side and turn to face each other as they converse, they, like adults, expect that if the people's lateral positions reverse during the conversation, each person will alter their movement so as to face the other again. Younger infants, in contrast, exhibit no such expectation on viewing this common, goal-directed social act: They are equally unsurprised if a person answers a friendly overture by turning toward or away from the person to whom he is speaking (Beier & Spelke, 2012). This finding and others suggest that core knowledge of agents applies to people's causal actions on objects but not to their social engagements with other people. Consistent with that suggestion, core knowledge of agents may take the form of a generative model that simulates the behavior of others who act efficiently to bring about high-value goal states in partially

observable physical, but not social, environments (Baker, Jara-Ettinger, Saxe, & Tenenbaum, 2017; Liu et al., 2017; Ullman & Tenenbaum, 2020). Only later in the first year might infants incorporate social motives into their analysis of agents' actions (Brune & Woodward, 2007).

One set of situations, however, does appear to elicit young infants' understanding of social actions: Situations in which one agent helps another agent to achieve its goal (Hamlin et al., 2007). When a person or puppet repeatedly tries and fails to complete an action (such as climbing a hill or opening a box) in view of two other agents, one who helps the actor and one who does not, infants as young as 3 months subsequently tend to look at the agent who provided help, and 6-month-old infants tend to reach for that agent. Infants' reaching and looking have been interpreted as reflecting a social preference for the helpful character, but I believe the research leaves other interpretations open. First and foremost, looking and touching are exploratory behaviors reflecting interest or curiosity and guiding learning: Indeed, their exploratory function may be primary even in social contexts,



**Figure 6.** Experiments testing infants' use of social imitation to identify new members of their social world. Twelve-month-old infants, who viewed either their own parent or a stranger (the parent of another infant) imitating one of two puppets, subsequently looked to the puppet who was imitated by their parent, but not by the stranger, when a voice that was synchronized with both puppets called to them by name. After Thomas, Saxe et al. (2022).

like bars and parties, where they signal desires to get to know someone. Thus, young infants may touch or look more to a previously helpful character not because its motives were prosocial, but because it produced a more interesting and potentially informative outcome: It ended the sequence of failed actions and freed the protagonist to do something new. Future research on young infants, focused on more specific behavioral or neural signatures of exploration and social engagement, could distinguish these accounts.

## 7. Social cognition

Chapter 8 introduces the last system of core knowledge, focused on social beings, engagements, and relationships. Social cognitive development has long been studied, but no clear consensus has emerged regarding the origins and nature of social knowledge. In this chapter, I hypothesize that a core system like those described in previous chapters underlies infants' developing knowledge of the mental states, attributes, and social relationships that connect people to one another and to the infant.

I first consider the overt social behaviors that engage young infants. When a real or pictured person looks at them, infants tend to look back (Farroni, Csibra, Simion, & Johnson, 2002) and mirror the person's movements of attention (Field, Woodson, Greenberg, & Cohen, 1982; Hood, Willen, & Driver, 1998). More strikingly, infants also tend to imitate the person's oral gestures and expressions of emotion (Field et al., 1982; Meltzoff & Moore, 1977; Fig. 6a). These behaviors emerge in the first months and are connected: For example, infants imitate and follow the gaze shifts of a face that begins with direct gaze, but not a face that begins with closed eyes or averted gaze. All three behaviors also are exhibited by infant apes and monkeys (e.g., Deaner & Platt, 2003; Ferrari, Paukner, Ionica, & Suomi, 2009; Mendelson, Haith, & Goldman-Rakic, 1982; Myowa, 1996), who tend to affiliate with people who imitate them (Paukner, Suomi, Visalberghi, & Ferrari, 2009). Attention following, imitation, and preferences for imitators also are exhibited by human adults (e.g., Chartrand & Bargh, 1999; Driver et al., 1999), although our imitative actions tend to be unconscious, and we only respond positively to people who imitate us when we are unaware that we are being imitated. Finally, these behaviors support learning: 6-week-old human infants, who attend to a person who looks at them and either opens his mouth or protrudes his tongue,

will attempt to reproduce the behavior a day later if the same person (but not a different person) faces them without moving (Meltzoff & Moore, 1994/2002), as do rhesus macaques (Paukner, Ferrari, & Suomi, 2011).

These findings suggest a hypothesis concerning the interconnected, abstract concepts on which a system of core social knowledge might center. Infants may respond in kind to the oral gestures, movements of attention, and emotional states of another person to signal their own engagement with the person and their motivation to share actions and experiences. Core social knowledge therefore may center on a conception of people as individuals with mental experiences of attention and emotion, who engage with one another and with the infant to share these experiences, and whose relationships persist through time.

The spontaneous behaviors of very young infants are hard to study, however, and harder to interpret. Although young infants move their attention in the direction of a potential social partner's changing gaze, this shift of attention guides their own eye movements only if the face that elicits it immediately disappears: An event that never occurs in real interactions. Young infants' own acts of imitation are subtle and sometimes appear after long and variable delays, making them difficult to leverage for further systematic study (see Meltzoff et al., 2018; Oostenbroeck et al., 2016). Although neonatal imitation has now been documented in newborn humans, apes, and monkeys, only oral movements – mouth opening, lip pursing, and tongue protrusion – have provided reproducible evidence for imitation in very young members of these species, possibly because the youngest primate infants, human and nonhuman, tend to engage primarily with close relatives, and oral movements resulting in saliva sharing are associated with close relationships at all ages (Thomas, Woo, Nettle, Spelke, & Saxe, 2022). Finally, neonatal imitation disappears after the first few months, possibly because infants with 3 or 4 months of social experience have learned who their close social partners are, reducing their receptiveness to a new, wholly unfamiliar person who purses their lips at them.

To overcome these limits, most of the research discussed in this chapter probes the early development of social knowledge by placing infants in the role of third-party observers of social interactions, and it analyzes infants' patterns of looking at the interactions as evidence for their expectations, surprise, or interest

in the events and participants. In this précis, I focus only on infants' responses to events involving two socially interactive behaviors – imitation and comforting – and their learning about the social connections between the individuals who exhibit them.

Experiments by Lindsey Powell and Heather Kosakowski presented 4-month-old infants with videos of two live people or animated characters who responded to the action of a third person or character, one by imitating that character and the other by responding with a different behavior. When the responding characters subsequently appeared without the target character, infants looked longer at the imitator, suggesting greater visual interest in characters who engage with others by sharing their actions (Kosakowski, Powell, & Spelke, 2016; Powell & Spelke, 2018). Infants also expected imitators to approach those whom they had imitated (Powell & Spelke, 2017).

Further research provides evidence that older infants learn about the relationships connecting the social beings whose interactions they observe. In one study, Annie Spokes presented infants with animated events involving two small characters who emitted baby cries, placed side by side below three large characters with adult voices who responded to those cries with comforting. The adult characters on the left and right responded only to the baby on their side of the display, and the adult in the center responded to one of the two babies as well. At test, only the adult characters appeared, and the central character alternately approached and danced with each of the side characters. Although the adults had not previously interacted directly, infants expected interactions between the two adults who had comforted the same baby: They inferred a social connection between two characters based on their interactions with a third party (Spokes & Spelke, 2017).

In recent studies, infants leveraged these abilities to learn about the members of their own social world. Ashley Thomas presented infants with videos of their own parent or an unfamiliar adult (the parent of another infant in the study) who responded to two puppets by imitating one puppet and not the other. These events were followed by videos in which the two puppets faced the infant in the parent's absence and simultaneously moved their mouths while a centrally located voice called to the baby by name (Fig. 6b). On hearing their name, infants looked to the puppet who had been imitated by their parent. Because no such effect occurred in the session with the unfamiliar adult, these findings provide evidence that infants inferred that the call to them came from the puppet whom their parent had imitated. Variations on this method showed that infants also expected a puppet who first demonstrated a relationship with them (by calling to them by name) to respond to the distress of their parent, but not to the distress of a parent of a different infant (Thomas, Saxe, & Spelke, 2022).

These findings provide evidence, I believe, for a core system of social knowledge. Like core knowledge in other domains, core social cognition is limited; I note two limits here. First, although newborn infants are acutely sensitive to people who look at them with direct gaze, they fail to distinguish people's faces from the faces of other animals, for they attend equally to the direct gaze of a human, a monkey, and a sheep (Pascalis, de Haan, & Nelson, 2002). Second, core social knowledge is suppressed if the parties to a social interaction attend to and act on objects. At 9 months, for example, infants smile at people who imitate their otherwise purposeless gestures but not at people who imitate their actions on objects (Agnetta & Rochat, 2004; Stern, 1985).

Moreover, 10-month-old infants accept objects more readily from people whose speaking or singing suggests that they are known social partners (Kinzler, Dupoux, & Spelke, 2007; Mehr & Spelke, 2017), but younger infants accept objects offered by anyone, regardless of their social identity or intentions, and show no understanding of gift-giving (Gordon, 2003). This pattern is striking, because the people whom the infant knows best both act on objects and engage with one another, and these events frequently occur together. Young infants, however, appear to view any given movement by a person either as a goal-directed action or as a social gesture, but not as simultaneously social and object-directed.

Several suggestions follow from these findings. First, core representations of people as social beings and as agents likely compete for attention in young infants, and perhaps in adults as well (Gray, Gray, & Wegner, 2007; Knobe & Prinz, 2008; Weisman, Dweck, & Markman, 2017). If so, then core social cognition may take the form of a generative model that simulates the shareable actions and experiences of known, individual people and the relationships that connect them. Infants' learning about their own social network has barely begun to be studied, however, and, to my knowledge, no detailed model of core social knowledge has been proposed. Second, infants' knowledge of people appears to undergo considerable changes toward the end of the first year. What might these changes be, and how and why might they occur? Such questions will take center stage in the successor to this book, but the last chapter of *What Babies Know* offers a preview of the ideas that I find most promising. To get there, however, the penultimate chapter turns to research on infants' learning of their native language.

## 8. Language

Infants' knowledge of language differs in critical ways from their knowledge of things, places, and people. Language is unique to humans and is learned slowly: Although infants begin to distinguish the sounds and partial meanings of highly frequent words in the first months (Bergelson & Swingley, 2012; Tincoff & Jusczyk, 1999), they don't confidently master the meanings of common words until about 14 months of age (Bergelson, 2019; Bergelson & Aslin, 2017). At birth, infants' knowledge of objects, places, and people is the same everywhere, but even newborn infants, drawing on experiences of language in the womb, respond differently to speech in different languages (Mehler et al., 1988). Core knowledge endures throughout life, but capacities for learning new languages decline with age (Hartshorne, Tenenbaum, & Pinker, 2018; Johnson & Newport, 1989). Finally, infants' core knowledge is manifest in their exploration, actions, and interactions with others. During most of the first year, in contrast, infants' language learning has little effect on their overt behavior. Much of the evidence for this learning comes from experiments that test for neural signatures of attention or surprise, elicited by changes in the structures or meanings of phrases (e.g., Friederici, Friedrich, & Christophe, 2007), or from behavioral experiments finding tiny increases in infants' looking to named objects or events (e.g., Bergelson & Swingley, 2015) or their head-turning toward speech with natural prosody and recognizable words (e.g., Jusczyk, 1997) and away from repetitive speech with no prosody or meaning (e.g., Saffran, Aslin, & Newport, 1996; see Black & Bergman, 2017).

Chapter 9 covers these and other aspects of infants' language learning, but in this précis I consider just one distinction that



infants master, between content words (e.g., the nouns, verbs, and adjectives that refer to things, events, and properties) and function words (e.g., the pronouns, determiners, auxiliaries, and part-words, like the past tense *-ed*, that signal relations between the things or events that content words designate). Function and content words differ in frequency (function words are more frequent), number (languages have many more content words), and phonological properties (content words tend to be spoken with greater stress and function words with shorter, reduced vowels).

Based on these properties, infants distinguish function from content words from the beginning, independently of their prenatal language exposure. In one study, for example, newborn infants were played a string of English function (or content) words, followed by new function and content words that were played in alternation. Based on the phonological properties that distinguish these categories in English, infants generalized to new words in the same category, and they reacted with interest to words in the opposite category, regardless of whether they had been exposed, in utero, to English or to a different language (Shi, Werker, & Morgan, 1999). Infants therefore distinguish between words in these two highly abstract categories. In contrast, infants do not learn the meanings or syntactic categories of any particular content or function word in their own language until the second year: For example, French- and German-learning infants below 14–18 months fail to expect that a native-language phrase beginning with a determiner (like the French or German counterpart of *the*) will include a noun, whereas a phrase beginning with a pronoun (like the counterpart of *he*) will include a verb (de Carvalho, He, Lidz, & Christophe, 2019; Hohle, Wiessenborn, Kiefer, Schulz, & Schmitz, 2004).

Although infants master the specific function words in their language slowly, they use their ability to recognize function words in any language to learn the patterns in which these words or part-words co-occur. In English, we say “The girl is dancing” and “The girl has danced,” but not “The girl is danced” or “The girl has dancing.” Remarkably, German-learning 4-month-old infants picked up on these relationships in a single session, when listening to sentences like these in an unfamiliar language (Italian): After a brief period of familiarization with the grammatical Italian phrases, EEG recordings showed a characteristic incongruity response to ungrammatical combinations of the functional morphemes, as did the EEGs of adult speakers of Italian. In contrast, adult speakers of German, tested under the same conditions as the German-learning infants, failed to respond to the relationships between the Italian function words (Friederici, Mueller, & Oberecker, 2011). The identification and proper use of function words is a formidable task for adult language learners.

Although the distinction between function and content words is universal, the typical ordering of these words varies across languages in ways that reflect a fundamental grammatical property of the language: The ordering of heads (like the subject of a sentence) and complements (like its predicate) in phrases. In some languages, like English and French, heads precede their complements (e.g., *he eats* and *a steak*). In some languages, like Japanese, this ordering is reversed, and in others, like German, both orderings occur in different sorts of phrases; children therefore must learn which ordering(s) their own language uses. The pairing of heads and complements with function and content words is not perfect: No function words appear in some phrases, and no content words in others (e.g., *John eats* and *another one*).

By 6–8 months, however, infants have learned how function and content words typically are ordered in their native language, and they quickly learn, using diverse cues, to find and order function and content words in a new, artificial language.

Evidence for this ability comes from a spectacular series of experiments, begun by Jacques Mehler and pursued by his collaborators and descendants, who have studied infants learning diverse languages (Gervain, Nespor, Mazuka, Horie, & Mehler, 2008). The experiments used a complex artificial language learning paradigm that I lack the space to describe here, based on the assumption that infants who have learned how heads and complements are ordered in their native language will expect a new artificial language to follow the corresponding ordering of function and content words, if no perceptible cues to their ordering are available. Indeed, infants quickly learned, in a single session, to categorize the function and content words in the artificial language based only on the differing numbers and frequencies of words in each category, and they ordered the function and content words in the artificial language in accord with the ordering of heads and complements in their native language. This finding provides evidence that infants have learned a central aspect of the syntax of their native language by about 7 months of age.

The words in the artificial language were spoken with uniform timing and intonation in the studies just described, but in further studies, additional information for the boundaries of phrases, based on subtle variations in pitch or syllable length that occur in some, but not all, of the world’s languages, was added to the sequences in the artificial language. Infants used this information to determine the head–complement order of the artificial language, even when that order was opposite to the order of heads and complements in their native language, and even when the cues to phrasal boundaries in the artificial language were not present in the infant’s own language (Bernard & Gervain, 2012). In further studies, infants whose family members spoke two languages with opposite head–complement orders used these cues to infer the ordering of function and content words in each of their languages (Gervain & Werker, 2013).

In brief, long before infants have confidently mastered the sounds or meanings of any words, they have begun to learn about the fundamental abstract structure and ordering of the words and phrases in their native language or languages. Language learning does not proceed from surface properties, like specific phonemes or words, to deeper, abstract properties, like syntactic structure and morphology; indeed, it appears to proceed in the opposite direction: Younger infants are more attuned to the abstract and general properties of language – properties that adults are not aware of – and only later learn what particular sounds distinguish one word from another (Werker, 1989). In these respects, the representations that support language learning resemble the representations of core knowledge.

Over the course of the first year, infants’ language learning interacts with core knowledge in interesting ways. First, infants only learn languages that are spoken by members of their social world: They don’t learn from the radio or from videos of an adult speaking directly to a different infant (Kuhl, Tsao, & Liu, 2003). Nevertheless, infants learn language even in cultures in which adults do not speak to them, likely based on the conversations they overhear (Cristia, Dupoux, Gurven, & Stieglitz, 2019). Second, infants’ earliest word learning tends to focus on entities captured by core knowledge: Words that name objects (like

*apple*; Bergelson & Swingley, 2012), accompany social gestures (like *bye-bye*; Bergelson & Swingley, 2013), or refer to body parts associated with actions (like *hands* and *feet*; Tincoff & Jusczyk, 2012), or close social engagements (like *mouth* and *nose*; Bergelson & Aslin, 2017).

Third, the function words that activate representations from a single core knowledge system can be processed rapidly and automatically, so they tend to be short and unstressed. This includes the English prepositions *in* and *on* that apply to mechanical relationships in the domain of the object system (Hespos & Spelke, 2004), and the plural markers that signal entities in the domain of the number system and are spontaneously invented by isolated deaf children (Coppola, Spaepen, & Goldin-Meadow, 2013). In contrast, representations that require activation of two or more core knowledge systems take longer to process, and the function words that express them tend to be longer and spoken with stress, like the English prepositions that designate spatial relationships between objects (e.g., *along*, *between*, *beside*, and *above*; Landau, 2017; Strickland, 2017). Nevertheless, words activate core knowledge systems automatically and seemingly effortlessly in adults and infants, suggesting that language, from the beginning, serves to represent objects and events more economically than do the core knowledge systems that simulate the behavior of these entities.

At about 9 months, infants begin to use the phrases spoken by others, containing both content words and words that are social (such as “Look!”): as invitations to share experiences with the speaker: invitations that transcend the limits of core knowledge. In one set of studies, infants viewed the events presented in Xu and Carey’s experiments in which two different objects alternately moved in and out of view, but with accompanying phrases announcing the appearance of each object with one content word (e.g., “Look, a toy. Look, a toy.”), two content words (e.g., “Look, a truck. Look, a duck.”), or no content words (e.g., “Look at this. Look here.”). Infants inferred that the screen hid two objects when a different content word announced each object but not otherwise (Dewar & Xu, 2007; Xu, 2002). In studies by Sandra Waxman, infants learned a new, subtle category of objects (e.g., diverse vehicles) when the same content word heralded each object (e.g., “Look, an auto.”) but not when the same function words were spoken (“Look at this one.”; Waxman & Braun, 2005; Waxman & Markow, 1995). Thus, infants infer, from a speaker’s choice of content words together with words that convey social intentions, whether the speaker aims to share experiences of one object, experiences of two objects, or experiences of the commonalities among a larger set of objects.

These findings suggest that infants have come to expect that speakers will be efficient, informative, and relevant to the situation they are speaking about: They have become sensitive to the pragmatics of language and other forms of communication. Once infants have developed this expectation, language will provide them with a remarkably effective tool for learning about the world by observing what people choose to talk about and what they choose to say. How might this expectation arise? By 3 months, infants expect people’s object-directed actions to be efficient and directed to things within their field of view: A possible basis for expectations that people will speak efficiently and relevantly to the current context. Infants also expect that people will share their experiences in states of social engagement: A possible basis for an expectation that they will speak informatively. If core agent and social representations compete for attention, however, how do older infants combine these notions and grasp the

intentions behind a single act of speaking? The last chapter focuses on this question.

## 9. Beyond core knowledge

Between 9 and 14 months, changes occur in infants’ social exploration and communication. At about 10 months, infants begin to understand social actions like the offering of an object (Gordon, 2003). Over the next months, infants begin to point to objects and to follow the points of others (Liszkowski, Carpenter, Henning, Striano, & Tomasello, 2004), to share attention to objects by looking back and forth between the object and the infant’s social partner (Bruner, 1974; Tomasello, 2008), and to view others’ pointing and looking at objects as goal-directed social actions (Brune & Woodward, 2007). In the last chapter of this book, I hypothesize that these developments stem from the emergence of a new understanding of people’s actions, engagements, and mental states. At about 10 months, infants begin to view people’s actions as guided by intentions that are both social and goal-directed: People are social agents. At about 12 months, infants begin to view people’s mental states as diverse, shareable experiences of a given situation: States that represent the world more finely than do perceptual systems, action systems, or systems of core knowledge (see Tomasello, 2018).

Both these changes, I hypothesize, are underpinned in part by infants’ developing mastery of language. By 9 months, infants likely interpret phrases like “Look, a duck,” as an invitation to share an experience of an object. This change in their understanding of speech may bring a new understanding of people, whose speech conveys intentions that are both social and object-directed. By 12 months, infants may notice that words like *animal*, *dog*, and *Fido* can be applied to the same object but convey different experiences of that object. This change may enrich children’s understanding of their own and other speakers’ mental states as simultaneously intentional (they refer to things outside themselves) and phenomenal (they convey experiences that the social agent wishes to share).

Does language learning prompt the emergence of these conceptions of other people’s actions and mental states, or does it reflect changes in mental state reasoning that have other causes? The question is open, but I lean toward the first possibility for several reasons. First, representations from different core systems compete for attention, preventing young infants from understanding social actions like pointing to objects. Nevertheless, young infants appreciate that words refer in some way to things, as revealed by their early steps in learning word meanings. As symbols, words package information more economically than do systems of core knowledge, which appear to function in part through processes of mental simulation. Words also combine to express a wide array of thoughts. Language therefore provides a medium in which representations from different systems of core knowledge could be represented economically, allowing distinct representations to combine and be called on for further learning and reasoning.

Second, language carves up the world more finely than does perception, action, or core knowledge. The latter systems allow us to perceive, attend to, and act on objects, but they don’t single out the experiences of an object that we might wish to share: Are we pointing to indicate to someone that this is a duck, the family pet, an intruder on the pond, or an event like the coming of spring? Our intentions to share experiences are conveyed best by language. As children’s understanding of their native language

grows, they become better placed to understand the diverse mental states of the people who speak to them.

With this understanding, children can begin to use language to learn new perspectives on the world – for example, that stars with the appearance of tiny points of light in the sky are distant suns, and the apparently flat ground on which we walk is a planet. These advances suggest a third reason to favor the hypothesis that language learning plays a causal role (among other factors) in children’s learning of new concepts. Most of what we know as adults is learned from other people, and almost everything that children learn in school is conveyed in part by language. Neither children nor adults learn history, politics, science, or mathematics by pointing at things and exploring them in isolation: Other people, from friends and teachers to authors and public speakers, guide them through the vast conceptual space that is available to minds that are endowed with a combinatorial and symbolic natural language, informed by core knowledge.

A further reason to believe that language learning aids conceptual development stems from the tendency of languages to change with the changing lives and needs of their speakers: Languages evolve to increase their scope and efficiency as cultures change. When new tools of uncertain usefulness are invented, speakers may describe them by means of phrases, like “computing machinery,” in Turing’s seminal paper (1950), or “self-driving car” today. If their referents prove to be widely useful and become ubiquitous, such phrases are likely to be replaced by single words, like “computer” or “laptop.” Languages thus provide efficient ways of capturing information that is useful to their speakers, increasing the economy of the phrases that convey their thoughts. Thus, the language that the child is learning provides a treasure trove of information about the culture in which he lives: It signals, by the frequency and brevity of its words and the contexts in which the words are spoken, the concepts that people find most useful and the circumstances in which they call on them.

A final reason to assign a causal role to language is that language can reverse what I call the curse of a compositional mind: Creatures who are endowed with a productively combinatorial language of thought can form a plethora of concepts, complicating the task of finding the right concept to use on any given occasion (see Piantadosi, Tenenbaum, & Goodman, 2012, for related ideas). Because natural languages are learned from speakers who aim to be informative, however, the most frequent content words that children hear will refer to concepts that speakers consider broadly useful, and children will learn these words and corresponding concepts before the less frequent ones. Because people aim to speak economically, moreover, child learners won’t be flooded with too much information. Finally, because people aim to speak relevantly to their own and their listener’s current concerns, children will entertain the concepts that others express primarily in contexts in which the speakers deem them to be useful guides to thought.

Prior to the onset of formal schooling, I suggest, 1-year-old children who have come to grasp the basic syntax, semantics, and pragmatics of their language will begin to develop new perspectives on the world, guided in part by the speech of those who share experiences with them. Children’s learning will vastly outpace that of other animals because the things people say, and leave unsaid, reflect not only the insights that the speaker has achieved on her own, but all that she has learned from others, directly and indirectly. The words and phrases of the language that people use to express their thoughts have made each natural language a valuable source of cultural information for child learners who are equipped

with core knowledge. This, however, is a story still to be told – I hope, in part, through this book’s successor.

**Acknowledgments.** The author thanks Tomer Ullman and Brandon Woo for help with this précis; the many investigators and students whose research it reviews, and all the parents and infants who have contributed to research on the origins of knowledge in infancy. The research findings described in the book and précis come from many countries, labs, investigators, and funding sources. I acknowledge the NSF-MIT Center for Brains, Minds, and Machines for supporting the recent research described in this précis, and the Institut Jean Nicod, Harvard University and the Paris Institute for Advanced Studies for supporting my efforts to pull together all the research that the book describes.

**Financial support.** This research received no specific grant from any funding agency, commercial, or not-for-profit sectors.

**Competing interests.** None.

## Note

1. Evidence for expectations of continuous object motion sometimes fail to replicate in experiments presenting videotaped events involving objects (e.g., Smith-Flores, Perez, Zhang, & Feigenson, 2022; Walco, 2022), likely because spatiotemporal continuity is violated during the editing of such events. In the animated events that children see on television, objects move discontinuously during cuts from one camera angle to another; during their video calls, people appear and disappear discontinuously at the call’s beginning and end. It remains to be seen how infants interpret these events, as their exposure to video continues to increase both in their ordinary lives and in experiments probing their knowledge. Most studies of object representation were conducted before the development of high-definition video or remote video conferencing, and they used real objects.

## References

- Agnetta, B., & Rochat, P. (2004). Imitative games by 9-, 14-, and 18-month-old infants. *Infancy*, 6(1), 1–36.
- Alvarez, G. A., & Franconeri, S. L. (2007). How many objects can you track? Evidence for a resource-limited attentive tracking mechanism. *Journal of Vision*, 7(13), 1–10.
- Amalric, M., & Cantlon, J. F. (2022). Common neural functions from children’s learning in naturalistic and controlled mathematics paradigms. *Journal of Cognitive Neuroscience*, 34(7), 1164–1182.
- Amalric, M., & Dehaene, S. (2016). Origins of the brain networks for advanced mathematics in expert mathematicians. *Proceedings of the National Academy of Sciences of the United States of America*, 113(18), 4909–4917.
- Baillargeon, R., & Carey, S. (2012). Core cognition and beyond: The acquisition of physical and numerical knowledge. In S. M. Pauen (Ed.), *Early childhood development and later outcome* (pp. 33–65). Cambridge University Press.
- Baker, C. L., Jara-Ettinger, J., Saxe, R., & Tenenbaum, J. B. (2017). Rational quantitative attribution of beliefs, desires, and percepts in human mentalizing. *Nature Human Behavior*, 1, 0064.
- Beier, J. S., & Spelke, E. S. (2012). Infants’ developing understanding of social gaze. *Child Development*, 83(2), 486–496.
- Bergelson, E. (2019). The nascent lexicon: Word learning in infants. Keynote address, Boston University Conference on Language Development, Boston, MA.
- Bergelson, E., & Aslin, R. D. (2017). Nature and origins of the lexicon in 6-month-old infants. *Proceedings of the National Academy of Sciences of the United States of America*, 114(49), 12916–12921.
- Bergelson, E., & Swingle, D. (2012). At 6 months, infants know the meanings of many common nouns. *Proceedings of the National Academy of Sciences of the United States of America*, 109(9), 3253–3258.
- Bergelson, E., & Swingle, D. (2013). The acquisition of abstract words by young infants. *Cognition*, 127(3), 391–397.
- Bergelson, E., & Swingle, D. (2015). Early word comprehension in infants: Replication and extension. *Language Learning and Development*, 11(4), 369–380.
- Bernard, C., & Gervain, J. (2012). Prosodic cues to word order: What level of representation? *Frontiers in Psychology*, 3, 451.
- Black, A., & Bergman, C. (2017). Quantifying infants’ statistical word segmentation: A meta-analysis. *Proceedings of the Cognitive Science Society*, 124–129.
- Brune, C. W., & Woodward, A. L. (2007). Social cognition and social responsiveness in 10-month-old infants. *Journal of Cognition and Development*, 8(2), 133–158.
- Bruner, J. (1974). From communication to language: A psychological perspective. *Cognition*, 3(3), 255–287.



- Cantlon, J. F., Libertus, M. E., Pineda, P., Dehaene, S., Brannon, E. M., & Pelphrey, K. A. (2009). The neural development of an abstract concept of number. *Journal of Cognitive Neuroscience*, 21(11), 2217–2229.
- Chartrand, T. L., & Bargh, J. A. (1999). The chameleon effect: The perception–behavior link and social interaction. *Journal of Personality and Social Psychology*, 76(6), 893–910.
- Cheng, K. (1986). A purely geometric module in the rat's spatial representation. *Cognition*, 23(2), 149–178.
- Cheng, K., & Gallistel, C. R. (1984). Testing the geometric power of an animal's spatial representation. In H. L. Roitblat, H. S. Terrace, & T. G. Bever (Eds.), *Animal cognition* (pp. 409–423). Erlbaum.
- Choi, Y., Mou, Y., & Luo, Y. (2018). How do 3-month-old infants attribute preferences to a human agent? *Journal of Experimental Child Psychology*, 172, 96–106.
- Clifton, R. K., Morrionello, B. A., & Dowd, J. M. (1984). A developmental look at an auditory illusion: The precedence effect. *Developmental Psychobiology*, 17(5), 519–536.
- Coppola, M., Spaepen, E., & Goldin-Meadow, S. (2013). Communicating about quantity without a language model: Number devices in homesign grammar. *Cognitive Psychology*, 67(1–2), 1–25.
- Cristia, A., Dupoux, E., Gurven, M., & Stieglitz, J. (2019). Child-directed speech is infrequent in a farmer-forager population: A time allocation study. *Child Development*, 90(3), 759–793.
- Deaner, R. O., & Platt, M. L. (2003). Reflexive social attention in monkeys and humans. *Current Biology*, 13(18), 1609–1613.
- de Carvalho, A., He, A., Lidz, J., & Christophe, A. (2019). Prosody and function words cue the acquisition of word meanings in 18-month-old infants. *Psychological Science*, 30(3), 319–332.
- Dehaene, S., & Cohen, L. (1997). Cerebral pathways for calculation: Double dissociation between rote verbal and quantitative knowledge of arithmetic. *Cortex*, 33(2), 219–250.
- Dehaene, S., Spelke, E. S., Pineda, P., Stancu, R., & Tsivkin, S. (1999). Sources of mathematical thinking: Behavioral and brain-imaging evidence. *Science (New York, N.Y.)*, 284(5416), 970–974.
- Dewar, K., & Xu, F. (2007). Do 9-month-old infants expect distinct words to refer to kinds? *Developmental Psychology*, 43(5), 1227–1238.
- Dillon, M. R., Kannan, H., Dean, J. T., Spelke, E. S., & Duflo, E. (2017). Cognitive science in the field: A preschool intervention durably enhances intuitive but not formal mathematics. *Science (New York, N.Y.)*, 357(6346), 47–55.
- Doeller, C. F., & Burgess, N. (2008). Distinct error-correcting and incidental learning of locations relative to landmarks and boundaries. *Proceedings of the National Academy of Sciences of the United States of America*, 105(15), 5909–5914.
- Doeller, C. F., King, G. A., & Burgess, N. (2008). Parallel striatal and hippocampal systems for landmarks and boundaries in spatial memory. *Proceedings of the National Academy of Sciences of the United States of America*, 105(15), 5915–5920.
- Driver, J., Davis, G., Ricciardelli, P., Kidd, P., Maxwell, E., & Baron-Cohen, S. (1999). Gaze perception triggers reflexive visuospatial orienting. *Visual Cognition*, 6(5), 509–540.
- Farooq, U., & Dragoi, G. (2019). Emergence of preconfigured and plastic time-compressed sequences in early postnatal development. *Science (New York, N.Y.)*, 363(6423), 168–173.
- Farroni, T., Csibra, G., Simion, F., & Johnson, M. H. (2002). Eye contact detection in humans from birth. *Proceedings of the National Academy of Sciences of the United States of America*, 99(14), 9602–9605.
- Ferrari, P. F., Paukner, A., Ionica, C., & Suomi, S. J. (2009). Reciprocal face-to-face communication between rhesus macaque mothers and their newborn infants. *Current Biology*, 19(20), 1768–1772.
- Field, T. M., Woodson, R., Greenberg, R., & Cohen, D. (1982). Discrimination and imitation of facial expression by neonates. *Science*, 218(4568), 179–181.
- Fodor, J. A. (1983). *The modularity of mind*. MIT Press.
- Foster, D. J. (2017). Replay comes of age. *Annual Review of Neuroscience*, 40, 581–602.
- Friederici, A. D., Friedrich, M., & Christophe, A. (2007). Brain responses in 4-month-old infants already are language-specific. *Current Biology*, 17(14), 1208–1211.
- Friederici, A. D., Mueller, J. L., & Oberecker, R. (2011). Precursors to natural grammar learning: Preliminary evidence from 4-month-old infants. *PLoS ONE*, 6(3), e17920.
- Ge, X., Zhang, K., Grizibis, A., Hamodi, A. S., Martinez Sabino, A., & Crair, M. C. (2021). Retinal waves prime visual motion detection by simulating future optic flow. *Science (New York, N.Y.)*, 373(6553), 412.
- Gergely, G., Nádasdy, Z., Csibra, G., & Bíró, S. (1995). Taking the intentional stance at 12 months of age. *Cognition*, 56(2), 165–193.
- Gervain, J., Nespore, N., Mazuka, R., Horie, R., & Mehler, J. (2008). Bootstrapping word order in prelexical infants: A Japanese–Italian cross-linguistic study. *Cognitive Psychology*, 57(1), 56–74.
- Gervain, J., & Werker, J. F. (2013). Prosody cues word order in 7-month-old bilingual infants. *Nature Communications*, 4(1), 1–6.
- Gibson, E. J. (1969). *Principles of perceptual learning and development*. Appleton Century Crofts.
- Gibson, E. J. (1980). Autobiography. In G. Lindzey (Ed.), *A history of psychology in autobiography* (Vol. 7, pp. 239–271).
- Gibson, E. J. (1982). The concept of affordances in development: The renaissance of functionalism. In W. A. Collins (Ed.), *The concept of development*. Minnesota symposium on child psychology (Vol. 15). Erlbaum. Reprinted, with new comments, in Gibson, E. J. (1991). *An odyssey in learning and perception*. MIT Press.
- Gibson, E. J. (1991). *An odyssey in learning and perception*. MIT Press.
- Gibson, E. J., & Bergman, R. (1954/1991). The effect of training on absolute estimation of distance over the ground. *Journal of Experimental Psychology*, 48(6), 473–482. Reprinted, with a new introduction, in Gibson, E. J. (1991). *An odyssey in learning and perception* (pp. 205–220). MIT Press.
- Gibson, E. J., Owsley, C. J., & Johnston, J. (1978/1991). Perception of invariants by 5-month-old infants: Differentiation of two types of motion. *Developmental Psychology*, 14, 407–415. Reprinted with a new introduction in Gibson, E. J. (1991). *An odyssey in learning and perception* (pp. 527–540). MIT Press.
- Gibson, E. J., & Walk, R. D. (1960). The “visual cliff”. *Scientific American*, 202(4), 62–71.
- Gibson, E. J., & Walker, A. S. (1984/1991). Development of knowledge of visual-tactile affordances of substances. *Child Development*, 55, 453–460. Reprinted with a new introduction in Gibson, E. J. (1991). *An odyssey in learning and perception* (pp. 541–555). MIT Press.
- Gordon, P. (2003). The origin of argument structure in infant event representations. In *Proceedings of the Annual Boston University Conference on Language Development* (Vol. 1, pp. 189–198). Cascadia Press.
- Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science (New York, N.Y.)*, 315(5812), 619.
- Gredebäck, G., Johnson, S. P., & von Hofsten, C. (2010). Eye tracking in infancy research. *Developmental Neuropsychology*, 35, 1–19.
- Halberda, J., Ly, R., Wilmer, J. B., Naiman, J. Q., & Germine, L. (2012). Number sense across the lifespan as revealed by a massive internet-based sample. *Proceedings of the National Academy of Sciences of the United States of America*, 109(28), 11116–11120.
- Halberda, J., Mazocco, M. M., & Feigenson, L. (2008). Individual differences in nonverbal number acuity correlate with maths achievement. *Nature*, 455(7213), 665–668.
- Hamlin, J. K., Wynn, K., & Bloom, P. (2007). Social evaluation by preverbal infants. *Nature*, 450(7169), 557–559.
- Hartshorne, J. K., Tenenbaum, J. B., & Pinker, S. (2018). A critical period for second language acquisition: Evidence from 2/3 million English speakers. *Cognition*, 77(2018), 263–277.
- Held, R., Birch, E. E., & Gwiazda, J. (1980). Stereoacuity of human infants. *Proceedings of the National Academy of Sciences of the United States of America*, 77(9), 5572–5574.
- Hermer, L., & Spelke, E. S. (1994). A geometric process for spatial reorientation in young children. *Nature*, 370(6484), 57–59.
- Hermer, L., & Spelke, E. S. (1996). Modularity and development: The case of spatial reorientation. *Cognition*, 61(3), 195–232.
- Hespos, S. J., & Spelke, E. S. (2004). Conceptual precursors to language. *Nature*, 430(6998), 453–456.
- Hohle, B., Wiessenborn, J., Kiefer, D., Schulz, A., & Schmitz, M. (2004). Functional elements in infants' speech processing: The role of determiners in the syntactic categorization of lexical elements. *Infancy*, 5(3), 341–353.
- Hood, B. M., Willen, J. D., & Driver, J. (1998). Adults' eyes trigger shifts in visual attention in human infants. *Psychological Science*, 9(2), 131–134.
- Huttenlocher, J., & Lourenco, S. F. (2007). Coding location in enclosed spaces: Is geometry the principle? *Developmental Science*, 10(6), 741–746.
- Hyde, D. C., Boas, D. A., Blair, C., & Carey, S. (2010). Near infrared spectroscopy shows right parietal specialization for number in pre-verbal infants. *NeuroImage*, 53(2), 647–652.
- Hyde, D. C., Khanum, S., & Spelke, E. S. (2014). Brief non-symbolic, approximate number practice enhances subsequent exact symbolic arithmetic in children. *Cognition*, 141(1), 92–107.
- Hyde, D. C., & Spelke, E. S. (2009). All numbers are not equal: An electrophysiological investigation of small and large number representations. *Journal of Cognitive Neuroscience*, 21(6), 1039–1053.
- Hyde, D. C., & Spelke, E. S. (2011). Neural signatures of number processing in human infants: Evidence for two core systems underlying numerical cognition. *Developmental Science*, 14(2), 360–371.
- Hyde, D. C., & Wood, J. N. (2011). Spatial attention determines the nature of nonverbal number representation. *Journal of Cognitive Neuroscience*, 23(9), 2336–2351.
- Izard, V., Sann, C., Spelke, E. S., & Streri, A. (2009). Newborn infants perceive abstract numbers. *Proceedings of the National Academy of Sciences of the United States of America*, 106(25), 10382–10385.
- Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The naive utility calculus: Computational principles underlying commonsense psychology. *Trends in Cognitive Sciences*, 20(8), 509–604.
- Javadi, A.-H., Emo, B., Howard, L. R., Zisch, F. E., Yu, Y., Knight, R., ... Spiers, H. J. (2017). Hippocampal and prefrontal processing of network topology to simulate the future. *Nature Communications*, 8(1), 14652.
- Johnson, J. S., & Newport, E. L. (1989). Critical period effects in second language learning: The influence of maturational state on the acquisition of English as a second language. *Cognitive Psychology*, 21(1), 60–99.
- Julian, J. B., Keinath, A. T., Muzzio, I. A., & Epstein, R. A. (2015). Place recognition and heading retrieval are mediated by dissociable cognitive systems in mice. *Proceedings of the National Academy of Sciences of the United States of America*, 112(20), 6503–6508.



- Jusczyk, P. W. (1997). *The discovery of spoken language*. MIT Press.
- Katz, L. C., & Shatz, C. J. (1996). Synaptic activity and the construction of cortical circuits. *Science (New York, N.Y.)*, 274(5290), 1133–1138.
- Kellman, P. J., Gleitman, H., & Spelke, E. S. (1987). Object and observer motion in the perception of objects by infants. *Journal of Experimental Psychology: Human Perception and Performance*, 13(4), 586–593.
- Kellman, P. J., & Spelke, E. S. (1983). Perception of partly occluded objects in infancy. *Cognitive Psychology*, 15(4), 483–524.
- Khanum, S., Hanif, S., Spelke, E. S., Berteletti, L., & Hyde, D. C. (2016). Effects of non-symbolic approximate number practice on symbolic numerical abilities in Pakistani children. *PLoS ONE*, 11(10), e0164436.
- Kinzler, K. D., Dupoux, E., & Spelke, E. S. (2007). The native language of social cognition. *Proceedings of the National Academy of Sciences of the United States of America*, 104(30), 12577–12580.
- Knobe, J., & Prinz, J. J. (2008). Intuitions about consciousness: Experimental studies. *Phenomenology and the Cognitive Sciences*, 7(1), 67–83.
- Kosakowski, H. L., Powell, L. J., & Spelke, E. S. (2016). Preverbal infants' third party imitator preferences: Animated displays versus filmed actors. Talk presented at the International Congress on Infant Studies, New Orleans, LA.
- Kuhl, P. K., Tsao, F. M., & Liu, H. M. (2003). Foreign-language experience in infancy: Effects of short-term exposure and social interaction on phonetic learning. *Proceedings of the National Academy of Sciences of the United States of America*, 100(15), 9096–9101.
- Landau, B. (2017). Update on “what” and “where” in spatial language: A new division of labor for spatial terms. *Cognitive Science*, 41(2), 321–350.
- Landau, B., Gleitman, H., & Spelke, E. (1981). Spatial knowledge and geometric representation in a child blind from birth. *Science (New York, N.Y.)*, 213(4513), 1275–1278.
- Landau, B., Spelke, E., & Gleitman, H. (1984). Spatial knowledge in a young blind child. *Cognition*, 16(3), 225–260.
- Lee, S. A., Sovrano, V. A., & Spelke, E. S. (2012). Navigation as a source of geometric knowledge: Young children's use of length, angle, distance and direction in a reorientation task. *Cognition*, 123(1), 144–161.
- Lee, S. A., Winkler-Rhoades, N., & Spelke, E. S. (2012). Spontaneous reorientation is guided by perceived surface distance. *PLoS ONE*, 7(12), e51373.
- Liszkowski, U., Carpenter, M., Henning, A., Striano, T., & Tomasello, M. (2004). Twelve-month-olds point to share attention and interest. *Developmental Science*, 7(3), 297–307.
- Liu, S., Brooks, N., & Spelke, E. S. (2019). Origins of the concepts cause, cost, and goal in 3-month-old infants. *Proceedings of the National Academy of Sciences of the United States of America*, 116(36), 17747–17752.
- Liu, S., Ullman, T. D., Tenenbaum, J. T., & Spelke, E. S. (2017). Ten-month-old infants infer the value of goals from the costs of actions. *Science (New York, N.Y.)*, 358(6366), 1038–1041.
- Liu, Y., Dolan, R. J., Kurth-Nelson, Z., & Behrens, T. E. J. (2019). Human replay spontaneously reorganizes experience. *Cell*, 178(3), 640–652.
- Luo, Y., & Johnson, S. C. (2009). Recognizing the role of perception in action at six months. *Developmental Science*, 12(1), 142–149.
- McCrink, K., & Wynn, K. (2004). Large-number addition and subtraction by 9-month-old infants. *Psychological Science*, 15(11), 776–781.
- Mehler, J., Jusczyk, P., Lambertz, G., Halsted, N., Bertoncini, J., & Amiel-Tison, C. (1988). A precursor of language acquisition in young infants. *Cognition*, 29(2), 143–178.
- Mehr, S. A., & Spelke, E. S. (2017). Shared musical knowledge in 11-month-old infants. *Developmental Science*, 21(2), e12542.
- Meltzoff, A. N., & Moore, M. K. (1977). Imitation of facial and manual gestures by human neonates. *Science (New York, N.Y.)*, 198(4312), 75–78.
- Meltzoff, A. N., & Moore, M. K. (1994/2002). Imitation, memory, and the representation of persons. *Infant Behavior and Development*, 25(1), 39–61.
- Meltzoff, A. N., Murray, L., Simpson, E., Heimann, M., Nagy, E., Nadel, J., ... Ferrari, P. F. (2018). Re-examination of Oostenbroeck et al. (2016): Evidence for neonatal imitation of tongue protrusion. *Developmental Science*, 21(4), e12609.
- Mendelson, M. J., Haith, M. M., & Goldman-Rakic, P. S. (1982). Face scanning and responsiveness to social cues in infant rhesus monkeys. *Developmental Psychology*, 18(2), 222–228.
- Michotte, A., Thines, G., & Crabbé, G. (1964). *Les compléments amodaux des structures perceptives*. Universitaires de Louvain.
- Myowa, M. (1996). Imitation of facial gestures by an infant chimpanzee. *Primates*, 37(2), 207–213.
- Naccache, L., Blandin, F., & Dehaene, S. (2002). Unconscious mass priming depends on temporal attention. *Psychological Science*, 13(5), 416–424.
- Nieder, A., & Dehaene, S. (2009). Representation of number in the brain. *Annual Review of Neuroscience*, 32, 185–208.
- O'Keefe, J., & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, 381(6581), 425–428.
- Oostenbroeck, J., Suddendorf, T., Nielsen, M., Redshaw, J., Kennedy-Costantini, S., Davis, J., ... Slaughter, V. (2016). Comprehensive longitudinal study challenges the existence of neonatal imitation in humans. *Current Biology*, 26(10), 2334–2338.
- Park, J., Bermudez, V., Roberts, R. C., & Brannon, E. M. (2016). Non-symbolic approximate arithmetic training improves math performance in preschoolers. *Journal of Experimental Child Psychology*, 152, 278–293.
- Pascalis, O., de Haan, M., & Nelson, C. A. (2002). Is face processing species-specific during the first year of life? *Science (New York, N.Y.)*, 296(5571), 1321–1323.
- Paukner, A., Ferrari, P. F., & Suomi, S. J. (2011). Delayed imitation of lipsmacking gestures by infant rhesus macaques. *PLoS ONE*, 6(12), e28848.
- Paukner, A., Suomi, S. J., Visalberghi, F., & Ferrari, P. F. (2009). Capuchin monkeys display affiliation toward humans who imitate them. *Science (New York, N.Y.)*, 325(5942), 880–883.
- Piantadosi, S. T., Tenenbaum, J. B., & Goodman, N. D. (2012). Bootstrapping in a language of thought: A formal model of numerical concept learning. *Cognition*, 123, 199–217.
- Piazza, M., Pica, P., Izard, V., Spelke, E., & Dehaene, S. (2013). Education enhances the acuity of the nonverbal approximate number system. *Psychological Science*, 24(6), 1037–1043.
- Powell, L. J., & Spelke, E. S. (2017). Human infants' understanding of social imitation: Inferences of affiliation from third-party observations. *Cognition*, 170, 31–48.
- Powell, L. J., & Spelke, E. S. (2018). Third party preferences for imitators in preverbal infants. *Open Mind*, 2(2), 61–71.
- Regolin, L., & Vallortigara, G. (1995). Perception of partly occluded objects by young chicks. *Perception and Psychophysics*, 57(7), 971–976.
- Saffran, J., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science (New York, N.Y.)*, 274(5294), 1096–1098.
- Scholl, B. J. (2001). Objects and attention: The state of the art. *Cognition*, 80(1–2), 1–46.
- Shi, R., Werker, J. F., & Morgan, J. L. (1999). Newborn infants' sensitivity to perceptual cues to lexical and grammatical words. *Cognition*, 72(2), B11–B21.
- Shuck, N. W., & Niv, Y. (2019). Sequential replay of nonspatial task states in the human hippocampus. *Science (New York, N.Y.)*, 364(6447), eaaw5181.
- Simion, F., Regolin, L., & Bulf, H. (2008). A predisposition for biological motion in the newborn baby. *Proceedings of the National Academy of Sciences of the United States of America*, 105(2), 809–813.
- Skerry, A., Carey, S., & Spelke, E. S. (2013). First-person action experience reveals sensitivity to action efficiency in pre-reaching infants. *Proceedings of the National Academy of Sciences of the United States of America*, 110(46), 18728–18733.
- Smith, K. A., Mei, L., Yao, S., Wu, J., Spelke, E., Tenenbaum, J. B., & Ullman, T. D. (2021). Modeling expectation violation in intuitive physics with coarse probabilistic object representations. 33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, CA.
- Smith-Flores, A. S., Perez, J., Zhang, M. H., & Feigenson, L. (2022). Online measures of looking and learning in infancy. *Infancy*, 27(1), 4–24.
- Sommerville, J. A., Woodward, A. L., & Needham, A. (2005). Action experience alters 3-month-old infants' perception of others' actions. *Cognition*, 96(1), B1–B11.
- Spelke, E. S. (1988). Where perceiving ends and thinking begins: The apprehension of objects in infancy. In A. Yonas (Ed.), *Perceptual development in infancy: The Minnesota symposium on child psychology* (Vol. 20). Erlbaum.
- Spelke, E. S. (2022). *What babies know: Core knowledge and composition*. Oxford University Press.
- Spelke, E. S., Kestenbaum, R., Simons, D. S., & Wein, D. (1995). Spatio-temporal continuity, smoothness of motion, and object identity in infancy. *British Journal of Developmental Psychology*, 13(2), 113–142.
- Spokes, A. C., & Spelke, E. S. (2017). The cradle of social knowledge: Infants' reasoning about caregiving and affiliation. *Cognition*, 159, 102–116.
- Stern, D. N. (1985). *The interpersonal world of the infant: A view from psychoanalysis and developmental psychology*. Basic Books.
- Strickland, B. (2017). Language reflects “core” cognition: A new theory about the origin of cross-linguistic regularities. *Cognitive Science*, 41(1), 70–101.
- Teller, D. Y. (1979). The forced-choice preferential looking procedure: A psychophysical technique for use with human infants. *Infant Behavior and Development*, 2, 135–153.
- Thomas, A. J., Saxe, R., & Spelke, E. S. (2022). Infants infer potential social partners by observing the interactions of their parents with unknown others. *Proceedings of the National Academy of Sciences of the United States of America*, 119(32), e2121390119.
- Thomas, A. J., Woo, B., Nettle, D., Spelke, E. S., & Saxe, R. (2022). Early concepts of intimacy: Young humans use saliva sharing to infer close relationships. *Science (New York, N.Y.)*, 375(6578), 311–315.
- Thrun, S. (2002). Robotic mapping: A survey. In G. Lakemeyer & B. Nebel (Eds.), *Exploring artificial intelligence in the new millennium* (pp. 1–35). Morgan Kaufman.
- Tincoff, R., & Jusczyk, P. W. (1999). Some beginnings of word comprehension in 6-month-olds. *Psychological Science*, 10(2), 172–175.
- Tincoff, R., & Jusczyk, P. W. (2012). Six-month-olds comprehend words to refer to parts of the body. *Infancy*, 17(4), 432–444.
- Tomasello, M. (2008). *Origins of human communication*. MIT Press.
- Tomasello, M. (2018). How children come to understand false beliefs. *Proceedings of the National Academy of Sciences of the United States of America*, 115(34), 8491–8498.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind: A Quarterly Review of Psychology and Philosophy*, 59(236), 433–460.

- Ullman, T. D. (2015). On the nature and origin of intuitive theories: Learning physics and psychology. Doctoral dissertation, Department of Brain and Cognitive Sciences, MIT.
- Ullman, T. D., Spelke, E. S., Battaglia, P., & Tenenbaum, J. B. (2017). Mind games: Game engines as an architecture for intuitive physics. *Trends in Cognitive Sciences*, 21(9), 649–665.
- Ullman, T. D., & Tenenbaum, J. B. (2020). Bayesian Models of conceptual development: Learning as building models of the world. *Annual Review of Developmental Psychology*, 2(1), 533–558.
- Valenza, E., Leo, I., Gava, L., & Simion, R. (2006). Perceptual completion in newborn human infants. *Child Development*, 77(6), 1810–1821.
- Vallortigara, G., Regolin, L., & Marconato, F. (2005). Visually inexperienced chicks exhibit spontaneous preference for biological motion patterns. *PLoS Biology*, 3(7), e208.
- von Hofsten, C. (1982). Eye-hand coordination in the newborn. *Developmental Psychology*, 18(3), 450.
- Walco, E. (2022). Surprise-induced exploration as a tool for learning: A comparative approach with human infants and non-human primates. Doctoral dissertation, Department of Human Evolutionary Biology, Harvard University.
- Walk, R. D., & Gibson, E. J. (1961/1991). A comparative and analytical study of visual depth perception. *Psychological Monographs*, 75, 44.
- Walk, R. D., Gibson, E. J., & Tighe, T. J. (1957/1991). The behavior of light- and dark-reared rats on a visual cliff. *Science (New York, N.Y.)*, 126(3263), 80–81. Reprinted with a new introduction in Gibson, E. J. (1991). *An odyssey in learning and perception* (pp. 141–146).
- Waxman, S. R., & Braun, I. (2005). Consistent but not variable names as invitations to form object categories: New evidence from 12-month-old infants. *Cognition*, 95(3), 257–302.
- Waxman, S. R., & Markow, D. B. (1995). Words as invitations to form categories: Evidence from 12- to 13-month-old infants. *Cognitive Psychology*, 29(3), 257–302.
- Weisman, K., Dweck, C. S., & Markman, E. M. (2017). Rethinking people's conceptions of mental life. *Proceedings of the National Academy of Sciences of the United States of America*, 114(43), 11374–11379.
- Werker, J. F. (1989). Becoming a native listener. *American Scientist*, 77(1), 54–59.
- Woo, B. M., Liu, S., & Spelke, E. S. (2021). Open-minded, not naïve: 3-month-old infants encode objects as the goals of other people's reaches. *Proceedings of the 43rd Annual Meeting of the Cognitive Science Society*, 43, 1914–1920.
- Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, 69(1), 1–34.
- Xu, F. (2002). The role of language in acquiring object kind concepts in infancy. *Cognition*, 85(3), 223–250.
- Xu, F., & Arriaga, R. I. (2007). Number discrimination in 10-month-old infants. *British Journal of Developmental Psychology*, 25(1), 103–108.
- Xu, F., & Carey, S. (1996). Infants' metaphysics: The case of numerical identity. *Cognitive Psychology*, 30(2), 111–153.
- Xu, F., & Garcia, V. (2008). Intuitive statistics by 8-month-old infants. *Proceedings of the National Academy of Sciences of the United States of America*, 105(15), 5012–5015.
- Xu, F., & Spelke, E. S. (2000). Large number discrimination in 6-month-old infants. *Cognition*, 74(1), B1–B11.

## Open Peer Commentary

### What we don't know about what babies know: Reconsidering psychophysics, exploration, and infant behavior

Karen E. Adolph<sup>a\*</sup>  and Mark A. Schmuckler<sup>b</sup> 

<sup>a</sup>Department of Psychology, New York University, New York, NY, USA and

<sup>b</sup>Department of Psychology, University of Toronto Scarborough, Scarborough, ON, Canada

[karen.adolph@nyu.edu](mailto:karen.adolph@nyu.edu); [mark.schmuckler@utoronto.ca](mailto:mark.schmuckler@utoronto.ca)

<https://www.nyuactionlab.com/>; <https://www.utscc.utoronto.ca/labs/infant/>

\*Corresponding author.

doi:10.1017/S0140525X23003217, e121

#### Abstract

Researchers must infer “what babies know” based on what babies do. Thus, to maximize information from doing, researchers should use tasks and tools that capture the richness of infants' behaviors. We clarify Gibson's views about the richness of infants' behavior and their exploration in the service of guiding action – what Gibson called “learning about affordances.”

#### How we know what babies know

Inferences about infant perception and cognition must be based on observable behaviors because babies can't talk, they don't follow instructions, and their physiological responses may reflect different psychological processes than those in adults (Blumberg & Adolph, 2023). But how to link observable behaviors with unobservable perception or cognition? A commendable feature of Spelke's (2022) book is her celebration of Eleanor Gibson and Richard Held for their pioneering behavioral methods to test infant depth perception. Gibson focused on exploratory behaviors and devised the “visual cliff” to test infants' responses to an apparent drop-off (Gibson, 1988; Gibson & Walk, 1960). Held used psychophysical functions to index the development of stereoscopic depth perception and devised the “kitten carousel” to test the role of self-produced visual feedback from locomotion (Held, Birch, & Gwiazda, 1980; Held & Hein, 1963).

Spelke's descriptions of Gibson's and Held's methods lay the foundation for her reliance on looking time to study infant cognition. Her arguments for innate avoidance of a drop-off lay the foundation for her claims about core knowledge. We take issue with these foundations on several counts. To really know what babies know requires more than rich interpretations based on lean looking-time behaviors. Standard looking-time procedures do not exploit the richness of infant behavior; infant exploration entails more than looking or not looking at a display; and standard looking-time procedures (and also Gibson's use of the visual cliff) are not psychophysical methods. Finally, Gibson's views about whether avoidance of the visual cliff is innate evolved over the course of her career.

Our comments are both biographical and intellectual. Like Spelke, we have unique insights into Gibson's work because all three of us were Gibson's doctoral students. However, we worked with Gibson at different periods in the evolution of her ideas – Spelke from 1973 to 1977, Schmuckler from 1983 to 1988, and Adolph from 1987 to 1993.

Spelke's cohort relied on looking-time methods (e.g., visual habituation, preferential looking). As in modern-day protocols, infants sit in a car seat or caregiver's lap while viewing a visual display, and researchers use group differences in looking duration to infer what infants know. The idea is that infants' “visual exploration” – here, defined by *whether*, not *where*, infants look at a display – can reveal perceptual sensitivity, learning over the course of multiple trials, or what infants knew prior to viewing the stimuli. Indeed, Spelke (1976, 1979) innovated an influential cross-modal looking-time procedure in which babies hear sounds or touch objects that match one of two visual displays.

Over the ensuing decades when Schmuckler and Adolph were doctoral students, Gibson returned to the questions that fascinated her in the early 1960s with the visual cliff. In contrast to Spelke's looking-time era, now infants were out of the car seat moving around. The focus was on sensitivity to optic flow for

balance and steering and perception of affordances for locomotion over challenging ground surfaces (e.g., Adolph, 1997; Adolph, Eppler, & Gibson, 1993; Gibson, 1997; Gibson et al., 1987; Schmuckler & Gibson, 1989). In these studies, apparatuses were not covered with safety glass; instead, babies' safety was ensured with netting or a researcher rescued infants when they fell. New apparatuses were built, new paradigms were developed, new tools were available for recording behavior, and thus new findings about perception, exploration, and action emerged. Some findings conflicted with Gibson's earlier work or called for new interpretations, including infants' behavior at the edge of a drop-off. Gibson willingly entertained these new findings and ideas, even if they conflicted with her prior interpretations. She would listen thoughtfully and say, "Well, dear, that's very interesting. Do more experiments to figure out what's going on."

### *Rich descriptions of behavior provide critical evidence about what infants know*

Behavior is infinitely rich. It is up to researchers – using the recording technologies and tools at their disposal – to decide how much of the richness to describe. For example, to demonstrate that conditioned responses need not be rigidly stereotypic, Gibson (1952, 1991) described nine reactions in infant goats to an aversive conditioned stimulus (shock to foreleg) – including both flexing and extending the leg, walking forward and backward, and wheeling in circles. Her descriptions of behavior on the visual cliff were equally rich (Gibson & Walk, 1960; Walk & Gibson, 1961). One insightful observation was that animals only show fear when forced onto the safety glass: When placed on the deep side, lambs and kids freeze in defensive postures with front legs rigid and hind legs limp. But when deciding for themselves whether to cross, animals calmly explore the shallow side of the apparatus and peer over the edge of the deep side. Later work showed that human infants spend most of their time near the brink of the deep side, regardless of whether they cross or avoid, showing neutral or positive – not negative – facial expressions in either case (Adolph, Tamis-LeMonda, Ishak, Karasik, & Lobo, 2008; Kretch & Adolph, 2017; Tamis-LeMonda et al., 2008).

Rich descriptions are not colorful frills. Rather, the suite of behaviors can provide converging evidence. For example, infants' increased hesitation, exploratory looking and touching, displacement behaviors, refusal to walk, and use of alternative strategies all point to infants' perception that a deformable waterbed or steep slope does not afford walking (Adolph, 1997; Gibson et al., 1987). Similarly, infants lean away from a looming object and toward an open aperture (Carroll & Gibson, 1981) and blink in response to a looming object but not to a looming aperture (Schmuckler & Li, 1998), providing converging evidence that infants perceive affordances for collision versus passage. Alternatively, behaviors can alter interpretations of the primary outcome measure, as when novice walkers explore the edge of a precipice, but walk over the brink nonetheless (Adolph, 1997; Adolph et al., 2008; Karasik, Tamis-LeMonda, & Adolph, 2016).

### *Looking time does not capture the richness of visual exploration*

In standard looking-time paradigms (e.g., preferential looking, visual habituation, violation of expectation), the richness of infants' visual exploration is reduced to a single metric – accumulated duration of looking. In the 1970s, researchers were limited

by available tools for recording duration of looks, and they could not record where on the display infants looked. But now eye-tracking technologies and video-annotation tools show that visual exploration is exceedingly rich – filled with looks of varying durations to various parts of a display. It is worthwhile to capture this richness. For example, Kellman and Spelke (1983) found that infants who are habituated to a partially occluded rod moving behind a box perceive one unified rod rather than two rod parts. But visual habituation relies on group data and is not sufficiently detailed or powerful to explain *why* infants perceive object unity for moving objects. Kellman and Spelke could only speculate. In contrast, richer eye-tracking data reveal that individual differences in visual exploration explain infants' perception of object unity: Infants who look at the visible parts of the moving rod during habituation subsequently dishabituate to two rod parts at test, whereas infants who visually explore the box or background do not (Johnson, Slemmer, & Amso, 2004).

Moreover, in standard looking-time studies, other behaviors babies emit are routinely ignored – pupil dilation, looks to caregivers, facial and manual gestures, vocalizations, postural changes, and so on. Looking-time researchers acknowledge that "surprise" at an "unexpected" event is merely shorthand for "longer looking" to displays researchers consider unexpected. But uninformed readers do not realize that behavioral indices of surprise such as infants' facial expressions – when included at all – often fail to provide converging evidence (Camras et al., 2002; Scherer, Zentner, & Stern, 2004).

Finally, when infants act in the world with all the behaviors in their repertoires, possibilities for exploration are limitless (Gibson, 1988). Babies explore objects with their eyes, hands, and mouths. They explore ground surfaces by looking, touching, and testing various postures and forms of locomotion. Their every movement generates perceptual information.

### *Looking time is not psychophysics: Open air or otherwise*

Spelke implies that looking time involves psychophysics. Since the 1960s, psychophysics is associated with methods that link systematic variations in a physical dimension of the environment with the accuracy of perception (Cornsweet, 1962; Green & Swets, 1966). Many psychophysical methods do not require self-report by adult humans (as Spelke attributes to Helmholtz). Notably, Teller's (1979) "forced-choice preferential-looking" procedure is a psychophysical method appropriate for use with babies and other nonverbal animals.

Psychophysical methods are extremely powerful because the ground truth is the known physical dimension. In addition, they yield psychometric functions for individual participants. In contrast, standard looking-time procedures are weak methods because they are based on the reliability of human judgments (or computer-vision algorithms) with no external ground truth. They must rely on group data and therefore cannot allow conclusions about individual infants. Moreover, because preferential and cross-modal preferential-looking paradigms lead researchers to accept both novelty and familiarity preferences as evidence of discrimination, these paradigms incorporate a fundamental ambiguity in determining and interpreting infants' perceptual experiences (Hunter & Ames, 1988). Contrary to Spelke's claims, psychophysics and looking time are different beasts.

Spelke correctly credits Held for conducting psychophysical experiments with infants, but incorrectly credits Gibson for doing the same. Gibson did conduct "open-air" psychophysical



studies with adults walking through fields judging distances among targets (Gibson & Bergman, 1954), and she admired Teller's ingenious forced-choice psychophysical method (Gibson & Pick, 2000). However, Gibson never used psychophysics with infants in her own studies. To be sure, change in a visual display from habituation to test or contrasts between shallow and deep sides of the visual cliff are experimental manipulations. But they are not psychophysics.

### *The evolution of Gibson's ideas from depth perception to perceiving affordances*

In her original studies, Gibson viewed the visual cliff as a test of depth perception, and she did indeed conclude that avoidance of the deep side was innate (Gibson & Walk, 1960; Walk & Gibson, 1961). Decades later, Gibson (1991) retained her assumption of innateness, but reinterpreted her work in terms of perception of affordances (Gibson, 1997), and by the 2000s, she believed that avoidance of a precipice is learned (Gibson, 1997, 2002; Gibson & Pick, 2000).

So why do infants avoid falling at the edge of an impossibly high drop-off, steep slope, narrow bridge, wide gap, or narrow ledge? As both Spelke and Gibson propose, sensitivity to depth information might be available at birth or shortly thereafter. However, depth perception is only a necessary condition, not a sufficient one. Altricial animals require learning to perceive the difference between a step and a cliff, an incline and a steep slope, a walkway and a bridge, and so on – that is, they must learn to perceive affordances for locomotion (Gibson & Pick, 2000).

What do infants learn? Negative feedback from falling is not necessary, as Walk and Gibson (1961) showed with dark-reared kittens that walked repeatedly onto the deep side of the visual cliff in the light and as Adolph (1995, 1997, 2000) showed with human infants on slopes and gaps. Likewise, experience with drop-offs, slopes, or other such obstacles is not necessary. Rather, infants must learn to perceive the relations between the physical features of the environment and the current status of their bodies and skills (Gibson, 1997). Such relations change from moment to moment, so exploration is needed to generate the requisite information. Spelke's innate knowledge and fetal dreams in precocial animals cannot supplant real-time exploration and learning.

### *Conclusions*

This commentary about Spelke's book is not merely a bunch of middle-aged, former students arguing about their mentor's legacy. It's about how to do developmental science. Spelke writes in her prologue, "To learn how infants think, we have to listen to what the infants in our studies tell us." However, overly simplified behaviors provide only a crude narrative of what babies might know. For Gibson, the richness of behavior is key. She advised her students to observe infants with an open mind, and to "let the behaviors speak to you." Thus, a critical lesson we learned in her lab is that the rich details of what babies do allow their voices to be heard more clearly.

**Acknowledgments.** The authors wish to thank Mark Blumberg and Catherine Tamis-LeMonda for their helpful comments on previous versions of this commentary.

**Financial support.** This commentary was supported by NICHD grant R01-HD33486 and DARPA grant N66001-19-2-4035 to K. E. A., and a Natural Sciences and Engineering Research Council of Canada Discovery Grant to M. A. S.

**Competing interest.** None.

### *References*

- Adolph, K. E. (1995). Psychophysical assessment of toddlers' ability to cope with slopes. *Journal of Experimental Psychology: Human Perception and Performance*, 21, 734–750.
- Adolph, K. E. (1997). Learning in the development of infant locomotion. *Monographs of the Society for Research in Child Development*, 62(3, Serial No. 251), 1–140.
- Adolph, K. E. (2000). Specificity of learning: Why infants fall over a veritable cliff. *Psychological Science*, 11, 290–295.
- Adolph, K. E., Eppler, M. A., & Gibson, E. J. (1993). Crawling versus walking infants' perception of affordances for locomotion over sloping surfaces. *Child Development*, 64, 1158–1174.
- Adolph, K. E., Tamis-LeMonda, C. S., Ishak, S., Karasik, L. B., & Lobo, S. A. (2008). Locomotor experience and use of social information are posture specific. *Developmental Psychology*, 44, 1705–1714.
- Blumberg, M. S., & Adolph, K. E. (2023). Protracted development of motor cortex constrains rich interpretations of infant cognition. *Trends in Cognitive Sciences*, 27, 233–245.
- Camras, L. A., Ujiie, T., Mayake, K., Wang, L., Murdoch, A., Meng, Z., ... Campos, J. (2002). Observing emotion in infants: Facial expression, body behavior, and rater judgments of responses to an expectancy-violating event. *Emotion*, 2, 179–193.
- Carroll, J. J., & Gibson, E. J. (1981). *Differentiation of an aperture from an obstacle under conditions of motion by three-month-old infants*. Paper presented at the meeting of the Society for Research in Child Development, Boston, MA.
- Cornsweet, T. N. (1962). The staircase-method in psychophysics. *American Journal of Psychology*, 75, 485–491.
- Gibson, E. J. (1952). The role of shock in reinforcement. *Journal of Comparative and Physiological Psychology*, 45, 18–30.
- Gibson, E. J. (1988). Exploratory behavior in the development of perceiving, acting, and the acquiring of knowledge. *Annual Review of Psychology*, 39, 1–41. doi:10.1146/annurev.ps.39.020188.000245
- Gibson, E. J. (1991). *An odyssey in learning and perception*. MIT.
- Gibson, E. J. (1997). Discovering the affordances of surfaces of support. *Monographs of the Society for Research in Child Development*, 62(3, Serial No. 251), 159–162.
- Gibson, E. J. (2002). *Perceiving the affordances: A portrait of two psychologists*. Erlbaum.
- Gibson, E. J., & Bergman, R. (1954). The effect of training on absolute estimation of distance over ground. *Journal of Experimental Psychology*, 48, 473–482.
- Gibson, E. J., & Pick, A. D. (2000). *An ecological approach to perceptual learning and development*. Oxford University Press.
- Gibson, E. J., Riccio, G., Schmuckler, M. A., Stoffregen, T. A., Rosenberg, D., & Taormina, J. (1987). Detection of the traversability of surfaces by crawling and walking infants. *Journal of Experimental Psychology: Human Perception and Performance*, 13, 533–544.
- Gibson, E. J., & Walk, R. D. (1960). The "visual cliff". *Scientific American*, 202, 64–71.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. Wiley.
- Held, R., Birch, E., & Gwiazda, J. (1980). Stereoaquity of human infants. *Proceedings of the National Academy of Sciences of the United States of America*, 77, 5572–5574.
- Held, R., & Hein, A. (1963). Movement-produced stimulation in the development of visually guided behavior. *Journal of Comparative and Physiological Psychology*, 56, 872–876.
- Hunter, M. A., & Ames, E. W. (1988). A multifactor model of infant preferences for novel and familiar stimuli. In C. Rovee-Collier, & L. P. Lipsitt (Eds.), *Advances in infancy research* (Vol. 5, pp. 69–95). Ablex Publishing.
- Johnson, S. P., Slemmer, J. A., & Amso, D. (2004). Where infants look determines how they see: Eye movements and object perception performance in 3-month-olds. *Infancy*, 6, 185–201.
- Karasik, L. B., Tamis-LeMonda, C. S., & Adolph, K. E. (2016). Decisions at the brink: Locomotor experience affects infants' use of social information on an adjustable drop-off. *Frontiers in Psychology*, 7, 797.
- Kellman, P. J., & Spelke, E. S. (1983). Perception of partly occluded objects in infancy. *Cognitive Psychology*, 15, 483–524.
- Kretch, K. S., & Adolph, K. E. (2017). The organization of exploratory behaviors in infant locomotor planning. *Developmental Science*, 20, e12421.
- Scherer, K. R., Zentner, M. R., & Stern, D. (2004). Beyond surprise: The puzzle of infants' expressive reactions to expectancy violation. *Emotion*, 4, 389–402.
- Schmuckler, M. A., & Gibson, E. J. (1989). The effect of imposed optical flow on guided locomotion in young walkers. *British Journal of Developmental Psychology*, 7, 193–206.
- Schmuckler, M. A., & Li, N. S. (1998). Looming responses to obstacles and apertures: The role of accretion and deletion of background texture. *Psychological Science*, 9, 49–52.
- Spelke, E. S. (1976). Infants' intermodal perception of events. *Cognitive Psychology*, 8, 553–560.



Spelke, E. S. (1979). Perceiving bimodally specified events in infancy. *Developmental Psychology*, 15, 626–636.


Spelke, E. S. (2022). *What babies know: Core knowledge and composition* (Vol. 1). Oxford University Press.

Tamis-LeMonda, C. S., Adolph, K. E., Lobo, S. A., Karasik, L. B., Dimitropoulou, K. A., & Ishak, S. (2008). When infants take mothers' advice: 18-month-olds integrate perceptual and social information to guide motor action. *Developmental Psychology*, 44, 734–746.

Teller, D. Y. (1979). The forced-choice preferential looking procedure: A psychophysical technique for use with human infants. *Infant Behavior and Development*, 2, 135–153.

Walk, R. D., & Gibson, E. J. (1961). A comparative and analytical study of visual depth perception. *Psychological Monographs*, 75(15, Whole No. 519), 1–44.

## Core knowledge, visual illusions, and the discovery of the self

Marlene D. Berke\*  and Julian Jara-Ettinger

Department of Psychology, Yale University, New Haven, CT, USA  
[marlene.berke@yale.edu](mailto:marlene.berke@yale.edu); [julian.jara-ettinger@yale.edu](mailto:julian.jara-ettinger@yale.edu)  
[marleneberke.github.io](https://github.com/marleneberke); [compdevlab.yale.edu](https://compdevlab.yale.edu)

\*Corresponding author.

doi:10.1017/S0140525X23003205, e122

### Abstract

Why have core knowledge? Standard answers typically emphasize the difficulty of learning core knowledge from experience, or the benefits it confers for learning about the world. Here, we suggest a complementary reason: Core knowledge is critical for learning not just about the external world, but about the mind itself.

Spelke (2022) approaches core knowledge as an empirical question. And it is on these grounds that *What Babies Know* compellingly argues for innate representational systems of knowledge. But this approach leaves a central question unanswered: Why do we have core knowledge in the first place?

The answer is often taken as self-evident, and it is largely left implicit in Spelke's book. If humans evolved innate representational knowledge, it must be because these systems are (1) too difficult to learn from experience alone, and (2) critical for making

the problem of learning about the world tractable. As Spelke puts it, humans develop a commonsense understanding of the world with much less data than machines require, and core knowledge (plus language) accounts for this difference.

Yet, these views center learning as a problem exclusively of building models of our environment. But core knowledge might be critical not only for learning about the world, but also for learning about our own minds. Metacognition – the ability to represent and build models of one's own mind – is widely agreed to be a foundational component of human intelligence, guiding how we learn (Flavell, 1979), shaping how we update our beliefs (Rollwage, Dolan, & Fleming, 2018), and possibly even forming the basis of self-awareness (Proust, 2013) and consciousness (Lau & Rosenthal, 2011; Peters, 2022). In some theories, it is uniquely human, and part of what separates humans from the rest of the animal kingdom (Carruthers, 2008). And yet, questions about how metacognitive representations might be learned or developed remain open. Core knowledge, we argue, might provide part of the answer.

Building a model of our own mind first requires that we distinguish mental representations that capture the external world from artifacts of how our mind works. This distinction can be far from clear-cut: Raw sensory data are processed by our perceptual systems with the goal of creating veridical representations (Berke, Walter-Terrill, Jara-Ettinger, & Scholl, 2022), but these computations can introduce (or fail to remove) distortions, sometimes leading to inaccurate representations of the outside world. In cases like those in Figure 1, our visual system's attempts to produce an accurate replica of the world end up, ironically, creating a compelling but incorrect representation. This is a deep challenge. Given the stream of mental representations built from sensory experience – and how phenomenologically compelling they all appear to us – how can we tell which parts reflect the external world and which parts ought to be mistrusted?

Consider how we might realize that the percepts in Figure 1 are illusions. As you approach the water on the highway ahead (Fig. 1A), it recedes and then vanishes as a function of your distance from it. If you see the artwork *Cercle et suite d'éclats* (by Felice Varini; Fig. 1B) in person, taking a few steps to the left or right would fragment and deform the floating circles, revealing that our viewpoint affects the objects' cohesion. And as we rove our eyes over the grid in Figure 1C, the circles at the intersections flicker between white and black, as though our eyes are somehow inducing action at a distance. In each of these cases, interpreting



**Figure 1 (Berke and Jara-Ettinger).** Examples of cases where perception provides compelling but incorrect representations of the world. (A) Mirage on the road. This is an atmospheric phenomenon where light bends upon encountering different densities of air. Perception fails to account for this distortion, leading to an illusory percept of water (in contrast to other light distortions that perception does account for and correct, such as color distortions because of shading, as in Adelson's checker shadow illusion; Adelson, 2001). (B) Felice Varini's (2009) *Cercle et suite d'éclats*, an art installation where curvatures were painted over a collection of houses, such that they appear as floating circles when seen from a particular viewpoint. (C) Scintillating grid illusion.

what we see against the backdrop of core knowledge enables us to realize that what we are seeing is not real.

When we identify that we are experiencing an illusion, we learn more than just that the world is not how it appears. We also learn about our own mind. Most directly, we learn to mistrust certain percepts, as when we realize that water on a highway can be safely ignored on a hot, sunny day. But the ability to identify illusions by comparing what we see against our core knowledge might also help us build more abstract models of our mind. Grounding experience on principles of reality helps us explicitly realize that mental representations are less trustworthy as a function of distance or lighting. Objects do not blur together when we take our glasses off; pencils don't break when dipped in water; and funhouse mirrors don't change our body.

These revelations are not limited to passive discoveries that happen when, by coincidence, we notice violations of core knowledge. As adults, we exploit core knowledge as a tool for reality testing. When we encounter something surprising, we carry out intuitive experiments over our mind: We are tempted to move our eyes and head to test how our visual experience changes. By having stable principles of how the world works, we can detect discrepancies between how the world seems (as determined by our perception) and how we know it ought to be (as determined by core knowledge), enabling us to build models of our own minds.

Our focus so far has been on object knowledge. Do other core knowledge systems also support learning about our minds? We believe this is the case. Consider how, when we lose our sense of direction, we do not wonder whether space has suddenly warped, but we instead attribute the lost sense of direction to a failure of our mind. Although such a realization might seem trivial, there are cases where core knowledge might help us make deeper discoveries about ourselves. Consider, for instance, the epistemic humility that comes knowing that we cannot always infer other people's mental states accurately. Such a realization might emerge from experiences where someone's behavior appears illogical, but we still hold on to the conviction that their actions must have resulted from a rational, goal-directed pursuit. These types of metacognitive representations not only help us understand ourselves, but they also make us better at navigating the world, at recognizing the limits of what we know, and at deciding when and how to explore so as to push those limits.

Does our proposal imply that all creatures that have core knowledge also have metacognition? This is unlikely: The process we propose requires core knowledge, but core knowledge alone is insufficient. At a minimum, an organism also needs (1) the capacity to instantiate representations over internal computations rather than over the external world – that is, metarepresentations – and (2) a learning algorithm that can build and refine metarepresentations by comparing experience against core knowledge (related work on artificial intelligence has shown proof-of-concept for such algorithms; Berke, Azerbayev, Belledonne, Tavares, & Jara-Ettinger, 2023). The prevalence of metacognition is therefore likely to be more restricted across species than core knowledge is.

Even if core knowledge did not evolve for the purpose of learning about our own minds, this does not make its ramifications for metacognition any less important. The use of core knowledge to learn metacognition may still be a major achievement in cognitive development. If we are correct, then core knowledge does not only play a pivotal role in learning about the external world, but also in learning about the internal world – how we think of ourselves, our mental lives, and who we are.

**Acknowledgments.** The authors thank Amanda Royka and Max Siegel for helpful comments.

**Financial support.** This work was supported by NSF award IIS-2106690.

**Competing interest.** None.

## References

- Adelson, E. H. (June 2001). On seeing stuff: The perception of materials by humans and machines. In B. E. Rogowitz, & T. N. Pappas (Eds.), *Human vision and electronic imaging VI* (Vol. 4299, pp. 1–12). SPIE.
- Berke, M. D., Azerbayev, Z., Belledonne, M., Tavares, Z., & Jara-Ettinger, J. (2023). MetaCOG: Learning a metacognition to recover what objects are actually there. arXiv preprint arXiv:2110.03105.
- Berke, M. D., Walter-Terrill, R., Jara-Ettinger, J., & Scholl, B. J. (2022). Flexible goals require that inflexible perceptual systems produce veridical representations: Implications for realism as revealed by evolutionary simulations. *Cognitive Science*, 46(10), e13195.
- Carruthers, P. (2008). Meta-cognition in animals: A skeptical look. *Mind & Language*, 23(1), 58–89.
- Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist*, 34(10), 906.
- Lau, H., & Rosenthal, D. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences*, 15(8), 365–373.
- Peters, M. A. (2022). Towards characterizing the canonical computations generating phenomenal experience. *Neuroscience & Biobehavioral Reviews*, 142, 104903.
- Proust, J. (2013). *The philosophy of metacognition: Mental agency and self-awareness*. Oxford University Press.
- Rollwage, M., Dolan, R. J., & Fleming, S. M. (2018). Metacognitive failure as a feature of those holding radical beliefs. *Current Biology*, 28(24), 4014–4021.
- Spelke, E. S. (2022). *What babies know: Core knowledge and composition*. Oxford University Press.
- Varini, F. (2009). Cercle et suite d'éclats. [Metallic paint on buildings]. Vercorin, Switzerland.

## The role of language in transcending core knowledge

Susan Carey\* 

Psychology, Harvard University, MA, Cambridge, USA  
scarey@wjh.harvard.edu

\*Corresponding author.

doi:10.1017/S0140525X23003060, e123

### Abstract

*What Babies Know* (WBK) argues that core knowledge has a unique place in cognitive architecture, between fully perceptual and fully conceptual systems of representation. Here I argue that WBK's core knowledge is on the perception side of the perception/cognition divide. I discuss some implications of this conclusion for the roles language learning might play in transcending core knowledge.

Spelke's monumental *What Babies Know* (Spelke, 2022) provides evidence for six domains of core knowledge: Innate systems of abstract, structured, representations with long evolutionary histories. In *The Origin of Concepts* (TOOC: Carey, 2009), I drew on Spelke's work to provide evidence for core knowledge, and I welcome and endorse Spelke's extended and more nuanced characterization in WBK.

Spelke and I each distinguish core knowledge from both perception and cognition. This has two parts: First, showing how perception and cognition differ, and second, characterizing how core knowledge has properties of each and lacks some properties of each. Here I argue for a different way of characterizing the difference between perception and cognition from that in *TOOC* and *WBK* that places core knowledge firmly on the perception side of the border.

**Distinguishing perception from cognition**

*WBK* and *TOOC* argue that perception is closer to sensory information than is cognition. Spelke adds that perception is modality specific. On the conception side, both point to the abstractness of the representations in core knowledge, and Spelke adds that core knowledge often has perceptual representations as crucial input. The problem is that these properties don't really distinguish perception from cognition. Perceptual systems often involve a series of computations, each with perceptual representations as their input (consider the representations of faces in successive face patches; Hesse & Tsao, 2020). Clear cases of perception (e.g., of the immediate spatial layout, in the service of reaching for objects) are amodal and computed by integrating information from audition, vision, and proprioception. Perceptual processes arguably create abstract representations (e.g., *cause* as in Michotte causality, *cardinal value*, *happy expression*, which all show hallmark signatures of perception, such as retinotopically specific adaptation effects; see Block, 2022). Such considerations have led many to deny a joint in nature between perception and cognition, with the consequence that core knowledge could not lie between the two.

To the contrary, I believe there is a deep divide between perception and cognition. Its essence is a fundamental difference in the kinds of representations within each: Differences in formats, and the kinds of structured representations and computations supported. With respect to format, Block (2022) argues that perceptual representations have iconic format, whereas conceptual representations have discursive format. Beck (2019) argues that perception is analog. Of course, analog and iconic representations can play a role in cognition, but cognition also involves logically structured representations, formulated over representations of structured propositions, and these are what distinguish cognition from perception.

In *TOOC* I speculate that the format of representation of core knowledge is iconic; whereas Spelke explicitly says she is agnostic about format. If my speculation is correct, then Spelke's and my claims for the place of core knowledge in cognitive architecture as its own natural kind between cognition and perception are false. If the format of the structured representations within core knowledge is exclusively iconic or analog, then core knowledge is a perceptual system of representation.

**Structured representations: Iconic/analog versus propositional formats**

Iconic representations are map like: They represent relations among parts of what is being represented with symbols that themselves instantiate those relations. The representations *resemble* or *mirror* (in Block's terminology) what is being represented. Pictures, including line drawings, are iconic symbols, as are maps and movies.

Analog representations involve symbols that are linear or logarithmic function values along a dimension that varies


continuously from small to big. There are analog representations of many dimensions: Of object size and weight, of number and density of ensembles, of temporal duration and loudness of sounds, of intensity of pain, and so on. While not map-like, there is a clear sense in which analog symbols are themselves iconic. A set of four is contained in a set of five; an analog symbol of five contains an analog symbol of four. Spelke establishes that there is an analog system of core number knowledge, the Analog Number System (ANS).

In contrast, propositional formats contain atomic symbols that do not mirror their referents and that participate in very different types of structured representation from maps and analog magnitude systems. The units of a proposition are typed syntactically, and these types determine the rules of combination into complex phrases and whole propositions. Clearly, language has a propositional format with words as units, composed into phrases, which in turn are still parts of sentences, which are truth-evaluable propositions.

Importantly, iconic/analog structures have no *explicit* symbols for logical relations. When I am looking at my table, my representation does not include a dog on it. But perception has no **symbol** *not*, nor symbols for many other abstract relations such as *all* or *same*.

**Iconic representation of structure in parallel individuation working memory models**

In addition to the ANS, Parallel Individuation (PI) is a second system of representation with numerical content. Spelke denies that PI is a system of number representation, for unlike the ANS, it contains no summary symbols for number. Indeed, the explicit atomic symbols in PI are representations of individuals (objects, events, sounds). But PI models iconically represent relations among those individuals. Numerical content is carried in the computations that guarantee that the model of three objects contains a symbol for each object. Numerical content is also carried by the computations these models support, such as planning the right number of reaches to retrieve those objects when hidden (Feigenson & Carey, 2003). *Figure 1* illustrates a working memory representation for sets of one, two, and three boxes. Note, there are no symbols for number, or for the relation *different* which holds between any two objects. Number and many relations, such a relative size and left to right spatial order, are represented iconically, being instantiated in the model.



**Figure 1 (Carey).** PI models.



Hochmann (2022) shows that in the context of Marcus' rule learning paradigm, the relations *same* and *different* are instantiated in PI models. Habituated to stimuli like "pi la pi," "du, no, du," "re bau re," infants (and adults) dishabituate to novel sequences "ta ku ku" or "ta ta ku" while generalizing habituation to "ta ku ta." Hochmann shows that the infants' working memory representation of sequence mirrors the relations among the syllables: ABA. Note there are no symbols for same in this representation, the content *same* is implicit in the match computations that allow all acts of recognition. There are also no explicit symbols for number. There are no explicit representations that there are three syllables, that two of them are the same, that not all of the syllables are the same. All those relations are implicit in the ABA grammar, though.

Hochmann (2022) provides two pieces of evidence for PI representations of syllable sequences in this paradigm. First, as with all PI models, there is a strict upper bound on number of syllables that can be represented (4, under these conditions). More relevant to us here, infants spectacularly fail to learn the generalization "all syllables are the same." They can represent what is in common between "pi pi"; "la la," treating "du mo" as an oddball, ditto for "pi pi pi"; "la, la la," treating "du du mo" as an outlier, and similarly for strings of four identical syllables. But familiarized to novel sequences of five or six identical syllables, they do not treat "du du du du mo" or "du du du du du mo" as outliers. They failed to represent the generalization that all the syllables in each sequence are the same. "All" and "same" have not been isolated within the iconic schema A A A.

Convergent evidence for failure to combine implicit representations of same with logical connectives (*not*) and (*all*) derives from Array Match to Sample experiments with large arrays (Fig. 2). Animals and young children can learn to match A (all-same) with novel A arrays and B (all-different) with novel B arrays. But they do so on the basis of analog magnitude representations of entropy, the degree of variability among the stimuli (Hochmann et al., 2017; Wasserman, Castro, & Fagot, 2017). In contrast, human beings over age 4 induce one of two propositional combinatorial rules: "Match all same with all same; not all same with not all same," or alternatively, "match all same with all same and all different with all different."

### How do propositionally structured representations arise?

Chapter 10 of *WBK* provides a spellbinding review of language learning in infancy. Although she doesn't call it so, language learning is supported by seventh system of core knowledge, differing from the others she reviews in two only two ways. First, it's unique to humans. Second, it is the only system whose proprietary internal structure is propositional. The capacity for




Figure 2 (Carey). Array match to sample stimuli.

propositionally structure representations is innate, encapsulated within core knowledge of language.

### Three roles for language transcending core knowledge: The relations same and different

Spelke and I agree that language has many roles to play in the creation of propositional representations with content that is embedded in other systems of core knowledge. She suggests that learning words for that content is important. I agree; learning the word "same" requires abstracting the relation same from the iconic schema A A A or ■ ■ ■, and may well be the impetus for doing so. But the above analysis predicts that words for aspects of the iconic models of core knowledge that are explicitly represented (e.g., the individuals in the PI models, including their properties) will be easier to learn than words for relations whose content is implicit, merely instantiated, in those models. This prediction is right. Words for objects, events like jump, and sounds are fast mapped. In contrast, Hochmann, Zhu, Zhu, and Carey (under review-b) show that the words for "same" and "different" are *hard* to learn. These words are common in speech to even 2-year-olds, and many 2-year-olds produce them. But virtually no 2-year-olds, less than 1/2 of 3-year-olds and only 80% of 4-year-olds have mastered their relational meanings.

Two kinds of bootstrapping, Gleitman's (1990) syntactic bootstrapping and Quinian bootstrapping (Carey, 2009), play crucial roles in this process. Both leverage propositional structures within language. In syntactic bootstrapping, the child uses the natural language syntactic and semantic representations they have already learned to constrain the meanings of newly encountered words. In Quinian bootstrapping, the child uses their knowledge of the propositional structure of language to create placeholder structures involving whole suites of interrelated concepts, none of which is yet mapped to any currently manifested meaning. *TOOC* shows how both of these processes are involved in creating the first explicit, non-analog, representations of number, and reviews evidence from the history of science and from science education for Quinian bootstrapping.

Hochmann et al. (under review-b) speculate how syntactic bootstrapping may play a role in learning the meanings of the words "same" and "different." One consequence of this analysis is that learning the meanings of these words should immediately support conceptual combination involving all of the linguistic propositional syntactic/semantic structure the child currently has. Hochmann, Zhu, and Carey (under review-a) review evidence that knowing the words "same" and "different" influences animal's and young children's performance on non-linguistic tasks drawing on these relations. Suggestively, it also finds that the proportion of 3- and 4-year-olds who can follow the rule "match the card where all of the pictures are the same to the card where all the pictures are the same; match the card where all of the pictures are not the same to the card where all the pictures are not the same" matches the proportion of children of these ages who know the words "same" and "different" in Hochmann et al. (under review-b). In comprehending this complex language, both children and adults make a small number of scope errors, to the same degree. Learning a discursive linguistically expressed symbol such as "same" provides immediate access to the logical functions expressed in language.

### Conclusions

Here I have argued that perception and core knowledge shared with other animals is not propositionally structured. This claim

is currently hotly debated, as attested by the BBS treatment of Quilty-Dunn, Porot, and Mandelbaum's (2023) target article on the current status of the language of thought hypothesis. *WBK* gives us further reason as a field to try to bring data to bear on the fundamental issues concerning formats of representation.

**Financial support.** The research reported here was funded by a James S. McDonnell Foundation Network grant "The Ontogenetic Origins of Abstract Combinatorial Thought". PI: Susan Carey; 16 co-PIs.

## References

- Beck, J. (2019). Perception is analog: The argument from Weber's Law. *The Journal of Philosophy*, 116, 319–349.
- Block, N. (2022). *The border between thinking and seeing*. Oxford University Press.
- Carey, S. (2009). *The origin of concepts*. Oxford University Press.
- Feigenson, L., & Carey, S. (2003). Tracking individuals via object files: Evidence from infants' manual search. *Developmental Science*, 6, 568–584.
- Gleitman L. (1990). The structural sources of verb meaning. *Language Acquisition*, 1, 3–55.
- Hesse, J., & Tsao, D. (2020). The macaque face patch system: A turtle's underbelly for the brain. *Nature Review Neuroscience*, 12, 695–716.
- Hochmann, J. R. (2022). Representations of abstract relations in infancy. *Open Mind*, 6, 291–310.
- Hochmann, J. R., Tuerk, A. S., Sanborn, S., Zhu, R., Long, R., Dempster, M., & Carey, S. (2017). Children's representations of abstract relations in relational/array match to sample tasks. *Cognitive Psychology*, 99, 17–43.
- Hochmann, J. R., Zhu, R., & Carey, S. (under review-a). Conceptual combination of mental representations of *same*, *not*, and *all* requires learning the words "same" and "different."
- Hochmann, J. R., Zhu, R., Zhu, L., & Carey, S. (under review-b). Learning the words "same" and "different" in the preschool years.
- Quilty-Dunn, J., Porot, N., & Mandelbaum, E. (2023). The best game in town: The reemergence of the language-of-thought hypothesis across the cognitive sciences. *Behavioral and Brain Sciences*, 46, 1–55.
- Spelke, E. S. (2022). *What babies know: Core knowledge and composition*. Oxford University Press.
- Wasserman, E., Castro, S., & Fagot, K (2017). Relational thinking in animals and humans: From percepts to concepts. In J. Call, G. M. Burgardt, I. M. Pepperberg, C. T. Snowden, & T. Zentall (Eds), *APA handbook of comparative psychology: Perception, learning & cognition* (pp. 359–384). American Psychological Association.

## Divisive language

Moira R. Dillon 

Department of Psychology, New York University, New York, NY, USA

### Corresponding author:

Moira R. Dillon; Email: [moira.dillon@nyu.edu](mailto:moira.dillon@nyu.edu)

doi:10.1017/S0140525X23003047, e124

### Abstract

What language devises, it might divide. By exploring the relations among the core geometries of the physical world, the abstract geometry of Euclid, and language, I give new insight into both the persistence of core knowledge into adulthood and our access to it through language. My extension of Spelke's language argument has implications for pedagogy, philosophy, and artificial intelligence.

As we wander the spaces of the physical world, our experience seems seamless, rich, and unitary, integrating the places we navigate with the visual forms in those places. Nevertheless, a rich series of studies in the psychological, cognitive, and neural

sciences – many of which were done by Elizabeth Spelke or her myriad mentees – suggest that different geometric representations underlie our experience of places and forms. In *What Babies Know*, Spelke (2022) argues that these different geometric representations are from different systems of "core knowledge," one system for *places* and another system for *forms*. Although core knowledge of places prioritizes distance and directional information for navigating paths through space, core knowledge of forms prioritizes hierarchically structured shape information for recognizing closed figures and objects. Spelke suggests, moreover, that human language allows the complementary geometries of the place and form systems to combine to support an intuitive abstract geometry that captures Euclidean geometry, a point she will expand on in her second volume, *How Children Learn*. Spelke's proposal is nevertheless committed to the persistence of the separate core systems of geometry throughout the human lifespan, remaining present and active even after older children learn Euclidean geometry, which is unitary in its integration of distance, direction, and shape.

In this commentary, I make two main points. First, I describe new evidence from a recent behavioral experiment in my lab that core knowledge about places and forms is indeed still present and active in educated human adults, consistent with Spelke's proposal (Lin & Dillon, 2023). My evidence complements evidence Spelke has put forward insofar as my tasks, unlike the tasks she reviews, relied only on simple and minimally contrastive linguistic descriptions – with no actual navigation or form analysis – to elicit core geometry of places and forms. Following this point, I then suggest that Spelke's "combined geometries," which are combined in language, can be later *re-isolated* through language, which is neither explicitly predicted by nor outlined in Spelke's proposal. I see my second point as a consistent – but not necessary – extension of Spelke's language argument that has implications for how we think about the relations among core knowledge systems and language more generally.

One pillar of core knowledge is its persistence throughout the lifespan, present – with all its original properties and limits – in human adults long after adults have developed the rich concepts that combine core knowledge. Spelke provides examples of this persistence in her review of each core system. For example, she describes how studies using brain-imaging techniques with adults navigating virtual environments reveal the signature limits of place geometry present in children and nonhuman animals (Doeller & Burgess, 2008; Doeller, King, & Burgess, 2008) and how studies using a two-alternative-forced-choice matching task show that adults, like infants, judge shapes as more similar when those shapes share the same skeletal structure versus the same 3D parts (Ayzenberg & Lourenco, 2019, 2022).

My lab's recent work, led by postdoctoral associate Yi Lin, provides new evidence for the persistence of both place and form geometry in adults, and it does so in a way that is complementary to the examples Spelke provides. In particular, we were able to elicit core geometry in adults for places and forms using simple, minimally contrastive linguistic descriptions and without adults' engaging in any actual navigation or form analysis. In our study, adults watched short videos of two points and two line segments forming an open figure on an otherwise blank screen. These simple figures were described with language that created different spatial contexts. After watching each video, adults were asked to provide a click response. In the *navigation* condition, they were told that they were seeing paths and stops that an agent traveled on a land. They were then asked to click on the

next step. In the *object* condition, they were told that they were seeing edges and corners of one side of an object. They were then asked to click on the next corner.

We wondered what geometry participants would preserve and perpetuate in their clicking responses given the language of their assigned condition. Could this minimal manipulation in language evoke core geometry? In particular, would adults in the navigation condition perpetuate the distance and direction of the figures' initial trajectories? Would those in the object condition instead preserve the initial figures' global shape?

Strikingly, adults produced responses reflecting different sets of geometric representations depending on the condition. In the navigation condition, adults perpetuated the figures' distance and directional information, producing open zig-zag paths. In the object condition, in contrast, they preserved the global shape of the initial figures and produced the third sides of what would be closed parallelograms. The clear and consistent reflections of the different geometries grabbed our attention because the procedure was open-ended and subjective and because the adult participants had been educated in formal geometry. These adults could have imagined a figure with *any* geometry. This task, inspired by other tasks' use of a simple and open-ended tapping procedure (e.g., Firestone & Scholl, 2014), was able to evoke effortlessly the particular geometric representations inherent to places and forms given minimally contrastive descriptions of the spatial context. Our results give new insight into both the persistence of core knowledge into adulthood and our access to it.

The power of language in this paradigm leads to my second point: What language joins it may unjoin. Spelke explains that when 9- and 10-month-old infants, like those in the studies of Xu and Carey (1996) and Xu (2002), see two different objects with two different shapes, like a cup and a shoe, emerge in alternation from either side of an occluder, they fail to predict that there are indeed two objects at play. The physical properties of the display trigger infants' object system and imply the presence of one moving body. This system outcompetes the infants' form system, which, from the spatial properties of the display, signals the presence of two different forms. When each object receives a different noun label upon emerging from its side of the occluder, however, infants can then use the objects' different forms to predict the presence of two objects. Spelke suggests that content words in language, like noun labels, allow for an efficient packaging of the activated core representations in a combined concept: In language, the cup and shoe are each simultaneously a moving body (*object* system) and a distinctive form (*form* system). Depending on the context, infants may choose between these core representations in a way that is relevant and efficient to the task at hand. For example, after hearing two different noun labels, infants may infer that the speaker intends to share with them two different experiences. Infants can then call upon their representation of two forms, from which follows the presence of two objects of different kinds. Older infants, children, and adults already have these combined concepts of cups and shoes as bodies with forms, concepts that were acquired through language, and so unlike younger infants, they do not need this initial labeling step to individuate the objects by their forms.

Content words thereby combine core concepts, as described above. But Spelke sees no evidence that content words express core concepts directly. Moreover, Spelke suggests that short and frequent function words, like *in* and *on*, may capture core knowledge "more directly" because such words express the mechanical relations between objects captured by the core

object system (Hespos & Spelke, 2004; Strickland, 2017). She states: "There is no word, in ordinary language, that refers to the objects, places, numerical magnitudes, forms, agents, or social beings revealed by the research on infants; that is why core knowledge is hard to write about." I suggest, however, that some content words are not too far off. After all, Spelke succeeds in writing elegantly about core knowledge! She was right, I think, to talk about *objects* instead of *schmobjects*. Despite the dangers of using ordinary language in scientific theory (Chomsky, 2000), and given that not conflating the ordinary-language combined concept *object* and the core-knowledge concept *object* is a challenge, the ordinary-language word *object* was nevertheless successfully used in Yi's and my experiment to evoke selectively a core inspiration. For example, describing the points, lines, and figures in our stimuli videos as "corners," "edges," and "objects" was enough to evoke the core geometry of the form system when adults were simply asked to click where the next "corner" would be. That this evocation was the geometry of *forms* that comes along with our everyday adult concept of *object*, however, proves both the power and limits of the ordinary language that names such concepts. My study with Yi makes the novel suggestion that at least some content words (or a brief collection of such content words) may also express core knowledge "more directly" – though only ever more or less directly. The re-isolation of core concepts within the medium of their combination reveals the scope and limits that define the combinatorial power of language: Language merges but does not meld (Chomsky, 1995; Chomsky et al., 2023).

So what seem like unitary concepts in ordinary language are never quite so, always already open to different evocations depending on context. If so, this may explain philosophical confusions (Reilly, 2019) and encourage pedagogical innovation. My example here again is geometry. My study with Yi also included an *abstract* condition, in which participants were told that they were seeing "points" and "lines" on an abstract "surface." They were then asked to click on the next point.

Participants in this *abstract* condition produced responses that were strikingly similar to participants' responses in the *navigation* condition. First, if abstract geometry is a combination of place and form geometries, then our findings suggest that this combination is not some *tertium quid*: The core knowledge in merged abstract geometry can be re-isolated through language. Spelke suggests that we lack a deep understanding of "the processes by which infants combine core representations with one another or with language," and I agree. Nevertheless, I suggest that in these combined representations, core concepts remain both intact and evocable through language. Second, abstract concepts in geometry maintain the competition between persistent core systems. In the case of our manipulation, place geometry wins. We suggest that under most conditions, in fact, the geometric representations humans call upon for reasoning about abstract points, lines, and figures may lie in representations we and other animals use for navigation: We wander the abstract world of Euclidean geometry like we wander the physical world of everyday life. Nevertheless, such competition raises the possibility that other manipulations probing abstract geometry, for example, those with different visuals or different language, could instead isolate form geometry over place geometry. These conclusions should inform philosophical debates over the origins of geometry (Husserl, 1970/1954; Kant, 1998/1781), interpretation of past empirical findings (e.g., Izard, Pica, Spelke, & Dehaene, 2011), the development of geometry pedagogies (e.g., Dillon, Kannan, Dean, Spelke, & Duflo,



2017), and the engineering of intelligent machines that aim to think mathematically like we do (McClelland, 2022; Sablé-Meyer, Ellis, Tenenbaum, & Dehaene, 2022).

Is the effect of content words in our study the norm or an exception? After all, Spelke describes how writing her book was a hard-won achievement. And, again, the term we used, *object*, evoked representations from the form system *not* from the object system! Nevertheless, I suggest that our findings call for explorations of whether and how language might isolate other core knowledge in adults' merged experiences, as, for example, our seemingly unified experience of the social world, which may, as Spelke suggests, instead rely on merged core knowledge of agents and social beings (see also Gray, Gray, & Wegner, 2007; Knobe & Prinz, 2008). For example, although our commonsense concept (or word) *person* might seem unified, similarly simple and minimal descriptions in language of people as either agents or social beings might re-isolate these core concepts underlying *person*. As for geometry, so for ethics. Spelke's core knowledge and language hypotheses, in combination, promise to be generative indeed in education and economics, philosophy and psychology, allowing us to probe the core of these domains.

**Acknowledgments.** The author thanks Elizabeth Spelke for her mentorship, support, and the opportunity to probe further just *What Babies Know*. The author thanks also Brian J. Reilly for his comments on the article.

**Financial support.** This work was supported by a National Science Foundation CAREER Award (DRL-1845924; to M. R. D.).

**Competing interests.** None.

## References

Ayzenberg, V., & Lourenco, S. (2022). Perception of an object's global shape is best described by a model of skeletal structure inhuman infants. *eLife*, *11*, e74943. doi:10.7554/eLife.74943

Ayzenberg, V., & Lourenco, S. F. (2019). Skeletal descriptions of shape provide unique perceptual information for object recognition. *Scientific Reports*, *9*, 1–12. doi: doi.org/10.1038/s41598-019-45268-y

Chomsky, N. (1995). *The minimalist program*. MIT Press.

Chomsky, N. (2000). *New horizons in the study of language and mind*. Cambridge University Press.

Chomsky, N., Seely, T. D., Berwick, R. C., Fong, S., Huybregts, M. A. C., Kitahara, H., ... Sugimoto, Y. (2023). *Merge and the strong minimalist thesis*. Cambridge University Press.

Dillon, M. R., Kannan, H., Dean, J. T., Spelke, E. S., & Duflo, E. (2017). Cognitive science in the field: A preschool intervention durably enhances intuitive but not formal mathematics. *Science (New York, N.Y.)*, *357*, 47–55. doi: doi.org/10.1126/science.aal4724

Doeller, C. F., & Burgess, N. (2008). Distinct error-correcting and incidental learning of location relative to landmarks and boundaries. *Proceedings of the National Academy of Sciences of the United States of America*, *105*, 5909–5914. doi: doi.org/10.1073/pnas.0711433105

Doeller, C. F., King, J. A., & Burgess, N. (2008). Parallel striatal and hippocampal systems for landmarks and boundaries in spatial memory. *Proceedings of the National Academy of Sciences of the United States of America*, *105*, 5915–5920. doi: doi.org/10.1073/pnas.0801489105

Firestone, C., & Scholl, B. J. (2014). "Please tap the shape, anywhere you like" shape skeletons in human vision revealed by an exceedingly simple measure. *Psychological Science*, *25*, 377–386. doi: doi.org/10.1177/095679613507

Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science (New York, N.Y.)*, *315*, 619. doi: doi.org/10.1126/science.1134475

Hespos, S. J., & Spelke, E. S. (2004). Conceptual precursors to language. *Nature*, *430*, 453–456. doi: doi.org/10.1038/nature02634

Husserl, E. (1970). *The crisis of European sciences and transcendental phenomenology* (D. Carr, Trans.). Northwestern University Press. (Original work published 1954).

Izard, V., Pica, P., Spelke, E. S., & Dehaene, S. (2011). Flexible intuitions of Euclidean geometry in an Amazonian indigene group. *Proceedings of the National Academy of Sciences of the United States of America*, *108*, 9782–9787. doi: doi.org/10.1073/pnas.1016686108

Kant, I. (1998). *Critique of pure reason* (P. Guyer & A. W. Wood, Trans. and Eds.). Cambridge University Press. (Original work published 1781).

Knobe, J., & Prinz, J. (2008). Intuitions about consciousness: Experimental studies. *Phenomenology and the Cognitive Sciences*, *7*(1), 67–83. doi: doi.org/10.1007/s11097-007-9066-y

Lin, Y., & Dillon, M. R. (2023). We are wanderers: Abstract geometry reflects spatial navigation. *Journal of Experimental Psychology: General*. doi:10.1037/xge0001504

McClelland, J. L. (2022). Capturing advanced human cognitive abilities with deep neural networks. *Trends in Cognitive Sciences*, *26*, 1047–1050. doi: doi.org/10.1016/j.tics.2022.09.018

Reilly, B. J. (2019). *Getting the blues: Vision and cognition in the middle ages*. Peter Lang.

Sablé-Meyer, M., Ellis, K., Tenenbaum, J., & Dehaene, S. (2022). A language of thought for the mental representation of geometric shapes. *Cognitive Psychology*, *139*, 101527. doi: doi.org/10.1016/j.cogpsych.2022.101527

Spelke, E. S. (2022). *What babies know: Core knowledge and composition*. Oxford University Press.

Strickland, B. (2017). Language reflects "core" cognition: A new theory about the origin of cross-linguistic regularities. *Cognitive Science*, *41*, 70–101. doi: doi.org/10.1111/cogs.12332

Xu, F. (2002). The role of language in acquiring object kind concepts in infancy. *Cognition*, *85*, 223–250. doi: doi.org/10.1016/S0010-0277(02)00109-9

Xu, F., & Carey, S. (1996). Infants' metaphysics: The case of numerical identity. *Cognitive Psychology*, *30*, 111–153. doi: doi.org/10.1006/cogp.1996.0005

## Is there only one innate modular system for spatial navigation?

Alexandre Duval\* 

School of Philosophy, Australian National University, Canberra, ACT, Australia  
alexandre.duval@anu.edu.au

\*Corresponding author.

doi:10.1017/S0140525X23003114, e125

### Abstract

Spelke convincingly argues that we should posit six innate modular systems beyond the periphery (i.e., beyond low-level perception and motor control). I focus on the case of spatial navigation (Ch. 3) to claim that there remain powerful considerations in favor of positing additional innate, nonperipheral modules. This opens the door to stronger forms of nativism and nonperipheral modularism than Spelke's.

A central thesis of *What Babies Know* (Spelke, 2022) is that there are (at least) six innate modular cognitive systems beyond the periphery of the mind, one for each of the following domains: objects, places, numbers, forms, agents, and social beings. Moreover, it seems clear from previous works (e.g., Spelke & Kinzler, 2007) and various discussions in the book that Spelke thinks that there are only a handful of systems that will turn out to be innate and/or nonperipheral modules – either exactly six or only slightly above six – and that research on core knowledge systems will therefore support *moderate* forms of both nativism and nonperipheral modularism.

My view on the book is that it does an excellent job of arguing for a *lower bound* on the number of such systems, but that it doesn't give strong reasons why we should stop at six and thus eschew stronger forms of nativism and nonperipheral modularism. It helps to distinguish two questions here: Are there additional innate modules operating *within* the six domains discussed in the book? Are there additional innate modules operating in *other* domains? I will make my case by focusing on the first question, and I will do so by taking spatial navigation (Ch. 3) as a case study. (Terminological note: In what follows, I

count the properties of *domain-specificity* and *encapsulation* as jointly sufficient for modularity.)

Chapter 3 defends an influential idea in navigation research commonly known as *the geometric-module hypothesis*. On a standard construal, it says that humans and many nonhuman species (including all mammals) possess an innate, domain-specific, encapsulated cognitive system that guides search behavior following sudden disorientation. Moreover, the system is encapsulated by virtue of operating on geometric representations of the three-dimensional surface layout of environments, and nothing else. The chapter doesn't explicitly argue for the view that this is the *only* innate module involved in spatial navigation. However, it rejects two challenges to that view, which I discuss in turn.

The first challenge relates to the ability to do *path integration*, which is well-documented in humans and many other species (Etienne & Jeffery, 2004). It is the process by which a subject keeps track of the distance and direction traveled from a certain origin point by relying on self-motion or *idiothetic* cues (i.e., proprioception, motor efference copy, vestibular signal related to head movements, and optic flow), perhaps along with other cues. Moreover, many researchers (e.g., Gallistel & King, 2010) believe that path integration is underpinned by an innate, domain-specific, encapsulated, nonperipheral cognitive system on something like the following grounds:

**Innateness:** Various species can perform path integration early in their life, with very little experience of the world (Bjerknes, Dagslott, Moser, & Moser, 2018; Newcombe, Huttenlocher, Drummey, & Wiley, 1998).

**Domain-specificity:** The system must use linear and angular velocity signals obtained from idiothetic cues to estimate the distance and direction traveled in recent bouts of spatial movements. To do so, it must perform the integration of velocity with respect to time, as well as other very specific mathematical operations suited to the task (Gallistel & King, 2010).

**Encapsulation:** Given the complexity and specificity of the mathematical operations involved, the system can only make use of input representations that have a very specific format. This in turn suggests that it will only rely on the inputs from a handful of systems, those that have evolved to cooperate with it – such as systems for dealing with idiothetic cues, as well as (possibly) systems encoding geometric or featural information about the environment (see below).

**Nonperipherality:** The system deals with abstract properties (location and heading of the subject), and it operates on information pertaining to multiple sense modalities (e.g., vestibular signal and optic flow). In addition, though it guides behavior in a variety of contexts (Etienne & Jeffery, 2004), it is not a low-level motor system either.

Finally, given that this system is triggered under different conditions (oriented navigation) than the geometric module (disoriented navigation), it is often thought that it is distinct from the geometric module.

Spelke's response to this challenge (p. 123) is to deny the last step. She holds that the core place system, a.k.a. the geometric module, is what deals with path integration. On this view, the geometric module is at work in the context of both oriented and disoriented navigation. In support of this claim, she argues that a number of navigation-related neurons in the mammalian hippocampus that underpin path integration display similar signature limits as the geometric module.

This response strikes me as problematic because of various findings about one category of navigation-related neurons: *place cells*. (Place cells are neurons that become active when an animal represents itself as being in a specific location in an environment.) In particular, I believe that there are good reasons to adopt the two following claims: (1) The implementation of the process of path integration in mammals directly involves place cells; and (2) place cells are sensitive to *featural cues* (e.g., odors, colors, textures, two-dimensional patterns on three-dimensional surfaces) in contexts where animals are performing path integration. Because it is a central commitment of the geometric-module hypothesis that the geometric module is completely insensitive to featural cues, (1) and (2) together entail that the geometric module can't be the system that implements path integration.

Why we should believe (1) and (2)? I will start by citing two strands of evidence in favor of (1). First, multiple studies suggest that lesions to the hippocampus proper, where place cells are located, undermines rodents' ability to go back to their home base when they are in the dark and deprived of olfactory cues (e.g., Maaswinkel, Jarrard, & Whishaw, 1999; Wallace & Whishaw, 2003). Second, Robinson et al. (2020) provide strong evidence that we can interfere with subjects' ability to perform path integration by intervening specifically on place cells. Robinson et al. began by training mice to move on a virtual-reality linear track and to perform licking behavior in a specific zone of the track, near the end, in order to receive a reward. Then, in one of the experimental conditions, when subjects reached a predetermined location around the midway point on the track, they underwent optogenetic activation of place cells that typically fired near the *beginning* of the track. In this context, mice started overshooting the reward zone and running straight through to the end of linear track significantly more often than before. This strongly suggests that the optogenetic activation of those cells around the midway point often caused the resetting of path integration to a previous position on the track.

Moving on to (2). Because this claim seems widely accepted among neuroscientists working on place cells, I will focus on only one paper: Fischler-Ruiz et al. (2021) showed that adding odors at specific points on a virtual-reality linear track significantly increases the number of hippocampal cells that qualify as place cells (according to standard methods for identifying such cells based on imaging data) as well as significantly improving the ability of mice to reach a reward zone at the end of the track in the dark. This supports the view that odors, which count as featural cues, can affect place-cell activity in path-integration contexts.

In sum, these findings suggest that proponents of the geometric-module hypothesis must accept that there is an additional innate, nonperipheral module that implements path integration.

The second challenge pertains to a theoretical paper (Duval, 2019) that argues, among other things, that extant versions of the geometric-module hypothesis are incomplete because they do not explain how subjects can reliably select the geometric representation of the current environment from memory following a sudden disorientation event.

Drawing on a variety of experiments that involve multiple enclosures, Duval further suggests that geometric-module theorists should posit a domain-specific and encapsulated cognitive mechanism that performs something like environment recognition by virtue of selecting a geometric representation of the current environment in memory. It operates according to the following principle: *Select the geometric representation in memory whose content best matches the current environment. If multiple*

representations match it about equally well, pick the one whose associated featural information best matches the featural cues in the current environment. Assuming that the selection mechanism exists as characterized here, it has to be distinct from the geometric module because it is a central commitment of the geometric-module hypothesis that the latter is insensitive to featural cues. Furthermore, there are number of reasons to think that it would be innate and nonperipheral:

**Innateness:** By hypothesis, the selection mechanism feeds geometric representations to the geometric module that the latter needs to perform its behavior-guiding functions. So, if the latter is innate and operating early in life (as Spelke argues on pp. 134–135), the former would likely be as well.

**Nonperipherality:** The selection mechanism deals with abstract properties (geometry of the three-dimensional surface layout of environments). Moreover, though it guides behavior indirectly through the information it feeds to the geometric module, it is far from a low-level motor system.

Spelke's response to the challenge raised by Duval consists in holding that there would *not* have been strong evolutionary pressures for a specialized mechanism in charge of environment recognition following a disorientation event. She writes: "Sudden, unknown, passive displacements to entirely new environments [...] happened close to never in the lives of animals or people in preindustrial times. [...] Although hurricanes or tidal waves may produce this situation, it is unlikely that we or other animals evolved specialized mechanisms for dealing with such rare events" (p. 93). She also points out that animals that actively navigate the world almost always change positions in a continuous fashion: For example, "one step at a time" (p. 93) in the case of animals who stay on the ground. Thus, the process of path integration can help them maintain a sense of where they are in cases when they are *not* undergoing unexpected, passive displacements (which are very rare).

I want to push back on this analysis. I believe that, contrary to what Spelke claims, there are specific, recurrent situations in the wild where animals would benefit from a specialized mechanism for environment recognition. These are precisely situations where path integration is unreliable. One example comes from exploratory looping behavior. Many species perform looping paths in uncharted territories for purposes of exploration (Eilam, 2014). Animals in this situation would benefit from a system in charge of environment recognition to determine whether they have come back to the environment where they started their exploration and have thus completed their loop. There is no way they can systematically rely on pure *idiothetic* path integration alone to determine whether have done so, as much work shows that *idiothetic* path integration quickly accumulates noise (see, e.g., Cheung, Ball, Milford, Wyeth, & Wiles, 2012; Thrun 2002). Another case pertains to animals that follow a familiar route in low-visibility conditions – because of fog, smoke, or the lack of sunlight at night – toward a known environment some distance away. For similar reasons about the unreliability of *idiothetic* path integration, such animals would benefit from a process of environment recognition to determine where they are on their route when there are sudden increases in visibility (e.g., a temporary clear-up in the fog, a better angle of the moon).

Hence, it seems that Spelke's response leaves intact the case, inspired by Duval (2019), for an evolved, innate, modular, and nonperipheral system in charge of environment recognition through geometric-representation selection. More generally, the

foregoing discussion supports the view that there (are least) *two* innate, nonperipheral modules for spatial navigation in human and nonhuman mammals besides the geometric module brilliantly championed by Spelke.

Let me conclude by emphasizing that Spelke has done an enormous service to the cognitive science community with this book by providing a careful, detailed, and extremely important analysis of a very wide range of experimental findings in support of moderate forms of nativism and nonperipheral modularism. Although I don't think that Spelke has given strong reasons to stop at the six innate modular systems that she identifies, the value of *What Babies Know* cannot be overstated.

**Financial support.** This work was supported by the Templeton World Charity Foundation (grant TWCF0539).

**Competing interest.** None.

## References

- Bjerknes, T. L., Dagslott, N. C., Moser, E. I., & Moser, M.-B. (2018). Path integration in place cells of developing rats. *Proceedings of the National Academy of Sciences of the United States of America*, 115(7), E1637–E1646.
- Cheung, A., Ball, D., Milford, M., Wyeth, G., & Wiles, J. (2012). Maintaining a cognitive map in darkness: The need to fuse boundary knowledge with path integration. *PLoS Computational Biology*, 8(8), e1002651.
- Duval, A. (2019). The representation selection problem: Why we should favor the geometric-module framework of spatial reorientation over the view-matching framework. *Cognition*, 192, 103985.
- Eilam, D. (2014). Of mice and men: Building blocks in cognitive mapping. *Neuroscience and Biobehavioral Reviews*, 47, 393–409.
- Etienne, A. S., & Jeffery, K. J. (2004). Path integration in mammals. *Hippocampus*, 14(2), 180–192.
- Fischler-Ruiz, W., Clark, D. G., Joshi, N. R., Devi-Chou, V., Kitch, L., Schnitzer, M., ... Axel, R. (2021). Olfactory landmarks and path integration converge to form a cognitive spatial map. *Neuron*, 109(24), 4036–4049.e5.
- Gallistel, C. R., & King, A. P. (2010). *Memory and the computational brain: Why cognitive science will transform neuroscience*. Wiley-Blackwell.
- Maaswinkel, H., Jarrard, L. E., & Whishaw, I. Q. (1999). Hippocampectomized rats are impaired in homing by path integration. *Hippocampus*, 9(5), 553–561.
- Newcombe, N. S., Huttenlocher, J., Drummey, A. B., & Wiley, J. G. (1998). The development of spatial location coding: Place learning and dead reckoning in the second and third years. *Cognitive Development*, 13(2), 185–200.
- Robinson, N. T. M., Descamps, L. A. L., Russell, L. E., Buchholz, M. O., Bicknell, B. A., Antonov, G. K., ... Häusser, M. (2020). Targeted activation of hippocampal place cells drives memory-guided spatial behavior. *Cell*, 183(6), 1586–1599.e10.
- Spelke, E. S. (2022). *What babies know: Core knowledge and composition* (Vol. 1, 1st ed.). Oxford University Press.
- Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental Science*, 10(1), 89–96.
- Thrun, S. (2002). Robotic mapping: A survey. In G. Lakemeyer & B. Nebel (Eds.), *Exploring artificial intelligence in the New Millennium* (pp. 1–35). Morgan Kaufmann.
- Wallace, D. G., & Whishaw, I. Q. (2003). NMDA lesions of Ammon's horn and the dentate gyrus disrupt the direct and temporally paced homing displayed by rats exploring a novel environment: Evidence for a role of the hippocampus in dead reckoning. *European Journal of Neuroscience*, 18(3), 513–523.

## Investigating infant knowledge with representational similarity analysis

Cameron T. Ellis\* 

Department of Psychology, Stanford University, Palo Alto, CA, USA  
cte@stanford.edu

\*Corresponding author.

doi:10.1017/S0140525X23003187, e126



## Abstract

Decades of research have pushed us closer to understanding what babies know. However, a powerful approach – representational similarity analysis (RSA) – is underused in developmental research. I discuss the strengths of this approach and what it can tell us about infant conceptual knowledge. As a case study, I focus on numerosity as a domain where RSA can make unique progress.

“What Babies Know” (Spelke, 2022) is a state-of-the-art, inclusive, and thoughtful survey of what babies know about a few key domains, making the book a must-read for junior scientists in the field. Perhaps most striking is the breadth of research methods and model systems used in the experiments covered: From various behavioral approaches, to neuroimaging, to computational models, to animal neurophysiology. Even so, this commentary will highlight a method mostly absent from the text (and field) that could offer unique insight into infants’ knowledge: Representational similarity analysis (RSA; Kriegeskorte, Mur, & Bandettini, 2008). I will discuss this method and explain its strengths, then focus on numerosity as a case study. Specifically, I will consider ways to address questions that the target book highlights as unresolved.

RSA is a method for evaluating the degree of representational similarity between exemplars in a set. RSA, along with the more straightforward “pattern similarity” method (Carlson, Schrater, & He, 2003), has been used frequently in vision science to compare the representations of exemplars, like faces, scenes, and objects. RSA is most often performed with functional magnetic resonance imaging (fMRI) (i.e., comparing the patterns of activity across voxels that exemplars evoke) but can be applied to any data where metric differences can be measured between exemplars, including electroencephalography (EEG; Xie et al., 2022) and behavior (Spriet, Abassi, Hochmann, & Papeo, 2022). One example of a question that RSA can ask is whether the representation of exemplars cluster according to their perceptual similarity (i.e., exemplars that are perceptually different have low similarity), or conceptual similarity (i.e., exemplars that are conceptually different have low similarity). Analyses of this type show that different parts of the brain can divide exemplars either perceptually or conceptually (Bankson, Hebart, Groen, & Baker, 2018).

A key benefit of RSA, and what makes it distinct from many popular methods in developmental science, is that it can test both the degree of difference between representations and the degree of *similarity*. By way of contrast, a habituation paradigm can test when infants think one exemplar (i.e., the habituation stimulus) is different from another (i.e., the test stimulus). However, suppose infants do not dishabituate to the test stimulus: That does not necessitate infants think the habituation and test exemplars are similar, but instead may result from poor study design (e.g., exposure is too short) or low statistical power. By contrast, high similarity in RSA typically (although not always, e.g., Spriet et al., 2022) entails positive evidence of similarity. For instance, in RSA with neuroimaging, high similarity means that neural patterns have a high correlation, which should not be expected by chance. An additional benefit of RSA arises when it uses neuroimaging data: It can disentangle simultaneous representations (Bankson et al., 2018), as illustrated in the example above where conceptual and perceptual information were localized in different brain regions.

Despite this potential value, RSA remains an uncommon tool in infant research. A few exceptions exist: Spriet et al. (2022) recently used RSA to understand the clustering of visual stimulus representations based on gaze behavior. Moreover, Xie et al. (2022) used EEG to measure neural responses to categories and evaluated their clustering. Studies using fMRI are surprisingly absent from this list, but this is likely to change soon given the recent success of awake infant fMRI (Ellis et al., 2020, 2021; Kosakowski et al., 2022; Yates et al., 2022). The aforementioned articles use RSA to address questions about visual perception, but RSA is flexible enough to tackle broad questions about what infants know. In what follows, I consider how RSA can resolve lingering questions in infant’s knowledge of numerosity.

Over the last 30 years, research has led to incredible progress in understanding what infants know about number. Nonetheless, there are at least two questions that RSA can advance. The first question is whether numerosity is represented conceptually, rather than reflecting mere magnitude of perceptual content. Both historically (Clearfield & Mix, 1999) and recently (Leibovich, Katzin, Harel, & Henik, 2017), researchers have argued that attempts to de-confound perceptual magnitude from conceptual number have been insufficient. For instance, many experimental controls for perceptual confounds have prioritized the wrong stimulus properties (Yousif & Keil, 2021). A second question RSA can tackle is the extent to which a signature of the numerosity system – namely, that differences in number depend on a ratio (AKA ratio dependence) – is consistent across the number line. Ratio dependence means that arrays containing 4 and 8 dots are perceived as just as different as arrays containing 8 and 16 dots. There is compelling evidence that infants represent numbers according to a ratio for values greater than 3 (i.e., above the subitizing range), but it remains unclear whether quantities lower than four are processed with ratio dependence (Hyde & Spelke, 2011; McCrink & Wynn, 2004; Starr, Libertus, & Brannon, 2013; Uller, Carey, Huntley-Fenner, & Klatt, 1999; see the target book for an extended discussion on this topic). These two questions can be addressed in the following RSA study.

Infants would see arrays of dots while undergoing fMRI (EEG, MEG or functional near-infrared spectroscopy [fNIRS] would be a viable substitute for some of these analyses). Across trials, arrays will differ in the number of dots from 1, 2, 3, 4, 8, and 16. These numbers include the subitizing range (1–3) and beyond (4–16). The dots will be presented in one of the two sizes: (1) the dots are all the same size, regardless of the quantity of dots in the array, and (2) the dots are scaled so that they span the same total size (where size is based on additive area; Yousif & Keil, 2021). The pattern of activity in different brain regions that are evoked by each array, averaged across repetitions, will be compared to all other array types to complete the RSA.

With this relatively simple design, four distinct analyses can assess infant knowledge of number:

- (1) Nonnumeric magnitude: are arrays with the same area (i.e., nonnumeric magnitude) more similar to each other than arrays with different areas, even when the number of dots differ? Brain regions sensitive to this nonnumeric magnitude will likely include both sensory systems and regions that support number processing in adults (e.g., the parietal cortex), as shown previously in adults (Sokolowski, Fias, Ononye, & Ansari, 2017).
- (2) Numeric magnitude: are arrays with the same number of dots more similar to each other than arrays with different

numbers, even when the size of the dots differs? Numeric magnitude has been shown to recruit the parietal cortex in infants (Hyde, Boas, Blair, & Carey, 2010). In adults, neural regions that code for numeric and nonnumeric magnitudes overlap, but only partially (Sokolowski et al., 2017). Using fMRI with awake infants, it's possible to test the degree to which these computations are supported by different systems: a viable hypothesis is that they start out similar and diverge during development.

- (3) Ordinal number representations: akin to Lyons, Ansari, and Beilock (2015), are numbers represented ordinally (e.g., 2 is more similar to 3 than it is to 4)? This is particularly interesting for quantities in the subitizing range, where the processing of numeric magnitude may be served by object tracking (Uller et al., 1999); thus, representations may not be ordinal in this range.
- (4) Ratio dependence: does representational similarity correspond more to absolute numerical differences (e.g., 2 is equally similar to 1 and 3) or a ratio difference (e.g., 2 is more similar to 3 than it is to 1)? Ratio dependence is a key indicator of a magnitude-estimation system (Meck & Church, 1983) and has been found for values in the approximate number range in infants (Xu, 2003). Whether ratio dependence exists for quantities in the subitizing range remains unclear (Hyde & Spelke, 2011; although see Starr et al., 2013).

These analyses highlight the flexibility of RSA and neuroimaging: It can test questions about the neural substrate of infant cognition (e.g., is the neural implementation of numerosity continuous across development) and also the nature of cognitive representations (e.g., what is the relationship between the representations of quantities). In this case study, neuroimaging gives both confirmatory and unique answers to questions regarding infant numerical perception. Even more exciting is that this is just a taste of what RSA and neuroimaging can do to help us understand the infant mind.

**Competing interest.** None.

## References

Bankson, B. B., Hebart, M. N., Groen, I. I., & Baker, C. I. (2018). The temporal evolution of conceptual object representations revealed through models of behavior, semantics and deep neural networks. *NeuroImage*, *178*, 172–182.

Carlson, T. A., Schrater, P., & He, S. (2003). Patterns of activity in the categorical representations of objects. *Journal of Cognitive Neuroscience*, *15*(5), 704–717.

Clearfield, M. W., & Mix, K. S. (1999). Number versus contour length in infants' discrimination of small visual sets. *Psychological Science*, *10*(5), 408–411.

Ellis, C. T., Skalaban, L. J., Yates, T. S., Bejjanki, V. R., Córdova, N. I., & Turk-Browne, N. B. (2020). Re-imagining fMRI for awake behaving infants. *Nature Communications*, *11*, 4523.

Ellis, C. T., Yates, T. S., Skalaban, L. J., Bejjanki, V. R., Arcaro, M. J., & Turk-Browne, N. B. (2021). Retinotopic organization of visual cortex in human infants. *Neuron*, *109*, 1–11.

Hyde, D. C., Boas, D. A., Blair, C., & Carey, S. (2010). Near-infrared spectroscopy shows right parietal specialization for number in pre-verbal infants. *NeuroImage*, *53*(2), 647–652.

Hyde, D. C., & Spelke, E. S. (2011). Neural signatures of number processing in human infants: Evidence for two core systems underlying numerical cognition. *Developmental Science*, *14*(2), 360–371.

Kosakowski, H. L., Cohen, M. A., Takahashi, A., Keil, B., Kanwisher, N., & Saxe, R. (2022). Selective responses to faces, scenes, and bodies in the ventral visual pathway of infants. *Current Biology*, *32*(2), 265–274.

Kriegeskorte, N., Mur, M., & Bandettini, P. A. (2008). Representational similarity analysis-connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, *4*, 1–28.

Leibovich, T., Katzin, N., Harel, M., & Henik, A. (2017). From “sense of number” to “sense of magnitude”: The role of continuous magnitudes in numerical cognition. *Behavioral and Brain Sciences*, *40*, e164.

Lyons, I. M., Ansari, D., & Beilock, S. L. (2015). Qualitatively different coding of symbolic and nonsymbolic numbers in the human brain. *Human Brain Mapping*, *36*(2), 475–488.

McCrink, K., & Wynn, K. (2004). Large-number addition and subtraction by 9-month-old infants. *Psychological Science*, *15*(11), 776–781.

Meck, W. H., & Church, R. M. (1983). A mode control model of counting and timing processes. *Journal of Experimental Psychology: Animal Behavior Processes*, *9*(3), 320.

Sokolowski, H. M., Fias, W., Ononye, C. B., & Ansari, D. (2017). Are numbers grounded in a general magnitude processing system? A functional neuroimaging meta-analysis. *Neuropsychologia*, *105*, 50–69.

Spelke, E. S. (2022). *What babies know: Core knowledge and composition*. Oxford University Press.

Spriet, C., Abassi, E., Hochmann, J. R., & Papeo, L. (2022). Visual object categorization in infancy. *Proceedings of the National Academy of Sciences of the United States of America*, *119*(8), e2105866119.

Starr, A., Libertus, M. E., & Brannon, E. M. (2013). Infants show ratio-dependent number discrimination regardless of set size. *Infancy*, *18*(6), 927–941.

Uller, C., Carey, S., Huntley-Fenner, G., & Klatt, L. (1999). What representations might underlie infant numerical knowledge? *Cognitive Development*, *14*(1), 1–36.

Xie, S., Hoehl, S., Moeskops, M., Kayhan, E., Kliesch, C., Turtleton, B., ... Cichy, R. M. (2022). Visual category representations in the infant brain. *Current Biology*, *32*(24), 5422–5432.

Xu, F. (2003). Numerosity discrimination in infants: Evidence for two systems of representations. *Cognition*, *89*(1), B15–B25.

Yates, T. S., Skalaban, L. J., Ellis, C. T., Bracher, A. J., Baldassano, C., & Turk-Browne, N. B. (2022). Neural event segmentation of continuous experience in human infants. *Proceedings of the National Academy of Sciences of the United States of America*, *119*(43), e2200257119.

Yousif, S. R., & Keil, F. C. (2021). How we see area and why it matters. *Trends in Cognitive Sciences*, *25*(7), 554–557.

## How important is it to *learn* language rather than *create* it?

Susan Goldin-Meadow 

Department of Psychology, The University of Chicago, Chicago, IL, USA  
<http://goldin-meadow-lab.uchicago.edu/>

### Corresponding author:

Susan Goldin-Meadow;  
 Email: [sgm@uchicago.edu](mailto:sgm@uchicago.edu)

doi:10.1017/S0140525X23003254, e127

### Abstract

I focus here on concepts that are *not* part of core knowledge – the ability to treat people as social agents with shareable mental states. Spelke proposes that *learning language from another* might account for the development of these concepts. I suggest that homesigners, who *create* language rather than learn it, may be a potential counterexample to this hypothesis.

Spelke has done a masterful job of describing six early emerging domain-specific systems of core knowledge that appear in infancy not only in children, but also in other animals. I focus here on two aspects of knowledge that Spelke argues are *not* present at birth and that humans develop over the first year of life – the ability to treat others as social agents, which develops at 10 months, and the ability to treat themselves and other social agents as having shareable experiences of objects and events, which develops at 12 months. Around their first birthday, children begin to construct a new system of knowledge of themselves and others as actors, collaborators, and sentient beings. How do these new concepts, which go beyond core knowledge, develop?

Spelke begins by considering whether infants' concepts of social agents and their mental states are, in fact, not a newly

developed skill at all, but rather present from the start. Under this view, the concepts are late-emerging because their appearance depends on other late-emerging capacities. She entertains, and rejects, three abilities that could potentially account for the development of these concepts. (1) The ability to share experiences of objects with others – shared intentionality – which onsets around the end of the first year but comes so naturally to human children that it might be considered innate (Tomasello, 2018). Spelke applies the criteria she uses to argue for core systems to shared intentionality. But she finds that it does not have the characteristics of a core system, nor is there an obvious explanation for its appearance at 12 months. (2) The ability to learn from others by interpreting the information directed to them as generic (Csibra & Gergely, 2009). Children do gain different information when an adult actively draws their attention to objects (thus creating a pedagogical atmosphere), compared to when an adult acts on objects without looking, speaking, or gesturing to the child. But a bias to interpret information from an adult as generic is just as likely to get in the way of learning about a social partner's mental states as to foster discovery of those states (cf. Powell & Spelke, 2018). (3) The ability to interact with objects symbolically; for example, to recognize an object from its picture. Interacting with people over pictorial symbols might lead infants to view those people as social agents who share their experiences of objects. Spelke rejects this explanation because infant's new conceptions of people, their actions, and their mental states at 12 months precede, rather than follow, children's understanding of pictures as symbols. The emergence of concepts of social agents and their mental states is therefore not likely to depend on any of these three abilities.

The fourth hypothesis is, for me, the most interesting. Spelke argues that learning a language brings children a new understanding of people and their mental states. Her proposal is that this understanding arises when children “learn enough of their language to interpret people's speech as simultaneously social and object-directed (at about 10 months) and to view their acts of speaking as invitations to share their experiences of the objects and events that they speak about (at about 12 months)” (Spelke, 2022, p. 403). Mastery of a specific natural language is thus thought to underlie the onset of social agents and their mental states.

Language may indeed play an essential role in the development of these late-emerging skills. But *learning* language from another cannot be the whole story. The evidence comes from homesigners – deaf children whose profound hearing losses prevent them from learning spoken language, and whose hearing parents have not exposed them to sign language. These children construct gestures, called *homesigns*, which have the properties of language that Spelke considers essential to the acquisition of social agents and their mental states – an open-ended lexicon, a productively combinatorial grammar generating abstract structured representations, and a compositional semantics (see Goldin-Meadow, 2020). Even though homesigners do not have a model for a conventional language, the gestures that their hearing parents produce when they talk to them might provide a model for combinatoriality or compositionality. But they don't (Goldin-Meadow, Mylander, & Butcher, 1995, 2007), confirming that homesigners do not learn these linguistic properties from others.

Given these facts about homesign, a good test of Spelke's hypothesis is to ask whether homesigners are able to treat people as social agents and to treat themselves and others as having mental states about shareable experiences of objects. They might, of

course, develop these skills late, particularly because homesigners are delayed in the development of some linguistic skills (e.g., communicating about the non-here-and-now, Morford & Goldin-Meadow, 1997). But the crucial question is *do they have them at all?* It's hard to imagine that adult homesigners, who can do relatively sophisticated things with their home-made languages (e.g., Coppola & Newport, 2005; Goldin-Meadow, Brentari, Coppola, Horton, & Senghas, 2015), do *not* have these concepts. But Pyers and Senghas (2009) have found that the homesigners who initially created Nicaraguan Sign Language have difficulty attributing false beliefs to others and thus are not proficient at understanding other peoples' minds. Moreover, as Spelke points out, nonhuman animals can *appear* to have an understanding of social agents in predictable contexts but show no understanding in novel contexts. In other words, homesigners might look more socially adept than they actually are. So this is not merely a thought-experiment – the test needs to be carried out.

If it turns out that adult homesigners are *not* able to treat people as social agents and treat themselves and others as having mental states about shareable experiences, this would provide strong support for Spelke's hypothesis – that learning language from another plays an important role in the development of these late-emerging concepts.

But if adult homesigners *do* display an understanding of these concepts, then *learning* language from another is not necessary to develop these abilities – *creating* language works too. Note, however, that learning some aspects of language might be necessary for children to develop other concepts; for example, homesigners have difficulty developing concepts of large exact numbers (Spaepen, Coppola, Spelke, Carey, & Goldin-Meadow, 2011) and certain spatial relations (Gentner, Özyurek, Gurcanli, & Goldin-Meadow, 2013), presumably because they lack a model from others for how to express these notions.

One final, clarifying point is worth making. Whether homesigners are truly a good test of Spelke's theory depends on what is central to the theory. Is it crucial that children learn from others *only* that they can share ideas about objects and people? If so, then homesigners might *not* be a good test case for the theory simply because their hearing parents gesture when they talk and use those gestures to point out objects and share attention to the objects with their children. Alternatively, is it crucial to the theory that children learn a compositional language from others? If so, homesigners do present a potential problem for the theory because their hearing parents do not provide them with a usable model for compositionality – the children *create* rather than *learn* these structures.

**Financial support.** This study was supported by NIDCD R01 DC00491; NSF BNS 8497941; and BCS-1654154.

**Competing interests.** None.

## References

- Coppola, M., & Newport, E. L. (2005). Grammatical subjects in home sign: Abstract linguistic structure in adult primary gesture systems without linguistic input. *PNAS*, 102(52), 19249–19253. <https://doi.org/10.1073/pnas.0509306102>
- Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, 13(4), 148–153. doi:10.1016/j.tics.2009.01.005
- Gentner, D., Özyurek, A., Gurcanli, O., & Goldin-Meadow, S. (2013). Spatial language facilitates spatial cognition: Evidence from children who lack language input. *Cognition*, 12(3), 318–330.
- Goldin-Meadow, S. (2020). Discovering the biases children bring to language learning. *Perspectives in Child Development*, 14, 195–201. doi: [doi.org/10.1111/cdep.12379](https://doi.org/10.1111/cdep.12379)



- Goldin-Meadow, S., Brentari, D., Coppola, M., Horton, L., & Senghas, A. (2015). Watching language grow in the manual modality: Nominals, predicates, and hand-shapes. *Cognition*, *135*, 381–395.
- Goldin-Meadow, S., Mylander, C., & Butcher, C. (1995). The resilience of combinatorial structure at the word level: Morphology in self-styled gesture systems. *Cognition*, *56*, 195–262.
- Goldin-Meadow, S., Mylander, C., & Franklin, A. (2007). How children make language out of gesture: Morphological structure in gesture systems developed by American and Chinese deaf children. *Cognitive Psychology*, *55*, 87–135.
- Morford, J. P., & Goldin-Meadow, S. (1997). From here to there and now to then: The development of displaced reference in homesign and English. *Child Development*, *68*, 420–435.
- Powell, L. J., & Spelke, E. S. (2018). Human infants' understanding of social imitation: Inferences of affiliation from third party observations. *Cognition*, *170*, 31–48. doi:10.1016/j.cognition.2017.09.007
- Pyers, J., & Senghas, A. (2009). Language promotes false-belief understanding: Evidence from Nicaraguan Sign Language. *Psychological Science*, *20*, 805–812.
- Spaepen, E., Coppola, M., Spelke, E., Carey, S., & Goldin-Meadow, S. (2011). Number without a language model. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(8), 3163–3168.
- Spelke, E. S. (2022). *What babies know: Core knowledge and composition* (Vol. 1). Oxford Press.
- Tomasello, M. (2018). How children come to understand false beliefs: A shared intentionality account. *PNAS*, *115*, 8491–8498. doi:10.1073/Pnas.1804761115

## The brain origins of early social cognition

Tobias Grossmann 

Department of Psychology, University of Virginia, Charlottesville, VA, USA  
[www.tobiasgrossmann.com](http://www.tobiasgrossmann.com)

### Corresponding author:

Tobias Grossmann;  
 Email: [tg3ny@virginia.edu](mailto:tg3ny@virginia.edu)

doi:10.1017/S0140525X23003151, e128

### Abstract

This commentary challenges Spelke's view on the early development of social cognition from a neuroscience perspective by presenting an overlooked body of evidence from neuroimaging research on joint attention with human infants. Indeed, evidence demonstrating adult-like, neural sensitivity to joint attention in young infants, supports alternative theoretical views concerning the origins of uniquely human forms of social cognition.

Spelke (2022) presents an intriguing account of the ontogenetic origins of human knowledge by reviewing research on human infants using diverse methods. Although the evidence put forward to support core knowledge systems in the physical domain, including places, objects, number, and geometry, is compelling, research examining the neuroscience of social cognition and especially the neural bases of joint attention in infancy has been overlooked. Specifically, there now exists a mounting body of evidence, showing that young infants, like adults, recruit medial prefrontal cortical regions supporting sophisticated social-cognitive functions, including joint (triadic) engagement (see Grossmann, 2013, 2015, for reviews). In adults, medial prefrontal cortex (mPFC) has been shown to play a key role in attributing and interpreting mental states (Amodio & Frith, 2006). There also is behavioral evidence suggesting that already

newborns show a basic sensitivity to eye gaze cueing of object locations foundational to joint attention (Farroni, Pividori, Simion, Massaccesi, & Johnson, 2004) and that, at least by 3 months of age, infants discriminate between dyadic and triadic social interactions (Striano & Stahl, 2005). First insights into the neural underpinnings of joint attention came from a series of event-related brain potential (ERP) studies. Specifically, Striano, Reid, and Hoehl (2006) examined the ERP correlates of joint engagement in 9-month-old infants in a paradigm in which an adult interacted live with the infant in two contexts. In the joint-attention context, the adult looked at the infant and then at the computer screen displaying a novel object. In the non-joint-attention context, the adult only looked at the chest of the infant and then at the novel object presented on the screen. Objects presented in the joint-attention context, compared to objects in the non-joint-attention context, were found to elicit a greater negative component (Nc) over frontal and central electrodes, an ERP component known to be generated within the mPFC (Reynolds & Richards, 2005). Based on these neural-level findings it was concluded that infants are sensitive to joint-attention interactions. Critically, this ERP paradigm has also been used to examine joint attention in younger infants (Parise, Reid, Stets, & Striano, 2008). This study reported that, already by the age of 5 months, infants show a selectively enhanced Nc during the joint-attention condition.

Relatedly, Grossmann and Johnson (2010) examined localized brain responses in 5-month-old infants during triadic social interactions using functional near-infrared spectroscopy (fNIRS). In this study, infants were presented with interactive scenarios in which a social partner (virtual agent presented on a screen): (a) Engaged in joint attention by gaze cueing the infant's attention to an object after establishing eye contact [joint-attention condition], (b) gaze cued the infant's attention to an empty location [no referent condition], or (c) looked at an object without prior eye contact with the infant [no eye contact condition]. Only in the joint-attention condition, infants selectively recruited a brain region within the mPFC, demonstrating that 5-month-old infants are sensitive to triadic interactions. Moreover, this study showed that 5-month-old infants employed a similar region of mPFC as seen during joint attention in human adults (Schilbach et al., 2010). Furthermore, there is neuroscience evidence to show that young infants are sensitive to when a social partner follows their gaze. Grossmann, Lloyd-Fox, and Johnson (2013) examined 5-month-olds' sensitivity to when a social partner follows their gaze by measuring infant brain responses using fNIRS during scenarios in which a social partner either followed the infants' gaze to an object that they had previously looked at (congruent condition) or a social partner shifted attention to look at a different object (incongruent condition). The fNIRS results of this study demonstrated that a selective region in the mPFC displayed an enhanced response to the congruent condition, suggesting that infants are sensitive to when someone follows their gaze. This finding provides early developmental evidence for theories, positing that brain processes are flexibly engaged by self- and other-initiated social interactions, including during joint-attentional engagement (Schilbach et al., 2013).

Taken together, the findings summarized above suggest the early developmental emergence of the brain system involved in joint attention by at least 5 months of age (Grossmann et al., 2013; Grossmann & Johnson, 2010; Parise et al., 2008; Striano et al., 2006). This is also in agreement with behavioral evidence

showing that young infants, by around 5 months of age, reliably follow pointing gestures, suggesting that the sensitivity to joint (triadic) engagement extends beyond eye cues to human-unique, gestural means of triadic communication (Bertenthal, Boyer, & Harding, 2014; Rohlfing, Longo, & Bertenthal, 2012). These findings, demonstrating the early emergence of the brain system involved in joint attention as a sophisticated form of human social cognition, challenge Spelke's account arguing for the relatively late emergence of such social-cognitive skills around 10 or 12 months of age. These neural-level findings are difficult to reconcile with, or integrate into, Spelke's current account, unless a new core system is added, or the core systems (agents and social partners) implicated in social cognition are revised or extended to adequately include existing evidence from the overlooked line of neuroimaging research on joint attention with infants. Moreover, this line of evidence from neuroimaging research with young infants also challenges Spelke's claims that the purported social-cognitive changes at the end of the first postnatal year of life are underpinned by "the mastery of language." Indeed, the neuroscience work overlooked by Spelke is much better aligned with alternative theoretical accounts, assigning developmental primacy to nonverbal, social-cognitive capacities, namely joint or shared attention/intentionality, being of foundational importance for the acquisition of spoken language and other communicative, cooperative, and cultural feats (Grossmann, 2017; Kuhl, 2007; Tomasello, 2019). In sum, this commentary challenges Spelke's view on the early development of social cognition from a neuroscience perspective, presenting an alternative of what may be considered uniquely human forms of social cognition in the service of cooperation, communication, and culture.

**Competing interests.** None.

## References

- Amodio, D. M., & Frith, C. D. (2006). Meeting of minds: The medial frontal cortex and social cognition. *Nature Reviews Neuroscience*, 7(4), 268–277. <http://dx.doi.org/10.1038/nrn1884>, <http://www.nature.com/nrn/journal/v7/n4/pdf/nrn1884.pdf>
- Bertenthal, B. I., Boyer, T. W., & Harding, S. (2014). When do infants begin to follow a point? *Developmental Psychology*, 50(8), 2036–2048. <https://doi.org/10.1037/a0037152>
- Farroni, T., Pividori, D., Simion, F., Massaccesi, S., & Johnson, M. H. (2004). Eye gaze cueing of attention in newborns. *Infancy*, 5, 39–60. [http://onlinelibrary.wiley.com/store/10.1207/s15327078in0501\\_2/asset/s15327078in0501\\_2.pdf?v=1&t=h1xhryvg&s=465228139544c235ede070d9b66702fd8983cc11](http://onlinelibrary.wiley.com/store/10.1207/s15327078in0501_2/asset/s15327078in0501_2.pdf?v=1&t=h1xhryvg&s=465228139544c235ede070d9b66702fd8983cc11)
- Grossmann, T. (2013). The role of medial prefrontal cortex in early social cognition. *Frontiers in Human Neuroscience*, 7. <https://doi.org/10.3389/fnhum.2013.00340>
- Grossmann, T. (2015). The early development of social brain functions in infancy. *Psychological Bulletin*, 141, 1266–1297.
- Grossmann, T. (2017). The eyes as windows into other minds: An integrative perspective. *Perspectives on Psychological Science*, 12, 107–121.
- Grossmann, T., & Johnson, M. H. (2010). Selective prefrontal cortex responses to joint attention in early infancy. *Biology Letters*, 6(4), 540–543. <https://doi.org/10.1098/rsbl.2009.1069>
- Grossmann, T., Lloyd-Fox, S., & Johnson, M. H. (2013). Brain responses reveal young infants are sensitive to when a social partner follows their gaze. *Developmental Cognitive Neuroscience*, 6, 155–161.
- Kuhl, P. K. (2007). Is speech learning gated by the social brain? *Developmental Science*, 10, 110–120. <https://doi.org/10.1111/j.1467-7687.2007.00572.x>
- Parise, E., Reid, V. M., Stets, M., & Striano, T. (2008). Direct eye contact influences the neural processing of objects in 5-month-old infants. *Social Neuroscience*, 3, 141–150.
- Reynolds, G. D., & Richards, J. E. (2005). Familiarization, attention, and recognition memory in infancy: An ERP and cortical source analysis study. *Developmental Psychology*, 41, 598–615.
- Rohlfing, K. J., Longo, M. R., & Bertenthal, B. I. (2012). Dynamic pointing triggers shifts of visual attention in young infants. *Developmental Science*, 15(3), 426–435. <https://doi.org/10.1111/j.1467-7687.2012.01139.x>
- Schilbach, L., Timmermans, B., Reddy, V., Costall, A., Bente, G., Schlicht, T., & Vogeley, K. (2013). Toward a second-person neuroscience. *Behavioral and Brain Sciences*, 36(04), 393–414. <https://doi.org/doi:10.1017/S0140525X12000660>
- Schilbach, L., Wilms, M., Eickhoff, S. B., Romanzetti, S., Tepest, R., Bente, G., ... Vogeley, K. (2010). Minds made for sharing: Initiating joint attention recruits reward-related neurocircuitry. *Journal of Cognitive Neuroscience*, 22(12), 2702–2715.
- Spelke, E. S. (2022). *What babies know: Core knowledge and composition*. Oxford University Press.
- Striano, T., & Stahl, D. (2005). Sensitivity to triadic attention in early infancy. *Developmental Science*, 4, 333–343. <http://onlinelibrary.wiley.com/store/10.1111/j.1467-7687.2005.00421.x/asset/j.1467-7687.2005.00421.x.pdf?v=1&t=h3mzje7b&s=1a86b6e63a35e34d175968386bcd9d2086bc3e6b>
- Striano, T., Reid, V. M., & Hoehl, S. (2006). Neural mechanisms of joint attention in infancy. *European Journal of Neuroscience*, 23, 2819–2823. <http://onlinelibrary.wiley.com/store/10.1111/j.1460-9568.2006.04822.x/asset/j.1460-9568.2006.04822.x.pdf?v=1&t=h3mzjvcm&s=c8937dbe937519cbf88d91ef1adde16defb57e5>
- Tomasello, M. (2019). *Becoming human: A theory of ontogeny*. Harvard University Press.

## Where is the baby in core knowledge?

Hyowon Gweon\*  and Peter Zhu

Department of Psychology, Stanford University, Stanford, CA, USA  
[gweon@stanford.edu](mailto:gweon@stanford.edu); [pgzhu@stanford.edu](mailto:pgzhu@stanford.edu)  
<https://sll.stanford.edu>

\*Corresponding author.

doi:10.1017/S0140525X2300314X, e129

### Abstract

What we know about what babies know – as represented by the core knowledge proposal – is perhaps missing a place for the baby itself. By studying the baby as an actor rather than an observer, we can better understand the origins of human intelligence as an interface between perception and action, and how humans think and learn about themselves in a complex world.

*What Babies Know* (Spelke, 2022) is a masterful synthesis of what we have discovered about the origins of human intelligence. The use of looking time methods has opened a window into the minds of incredibly young infants, enabling science to weigh in on questions that philosophers have debated for centuries: The content of the human mind and how it develops in the earliest phase of life. In this book, Spelke proposes core knowledge as a way to explain both the remarkable systematicity and striking diversity in human behavior; humans are endowed with a system of knowledge in key domains (systematicity) that provides the foundational concepts that support learning from experience (systematicity and diversity).

The comprehensive review of the empirical evidence for this proposal, however, ironically highlights a glaring gap: Where is the baby (i.e., the self)? As much as it tells a rich, compelling story about what infants know about the external world, it also reveals how the field of infant cognition has focused primarily on infants' representations of their surroundings as onlookers. By characterizing the baby as a learner building a model of the world – and by constraining the scope of "the world" to what is external to the learner – researchers have often overlooked the baby as a key constituent of that world; infants occupy a coordinate in space, are subject to physical principles, and can exert causal power on objects and agents (including themselves). How do infants represent themselves in the physical and the

social world, not just as observers, but as actors who navigate in space, manipulate objects, and interact with others? How do infants build a model of the self, both as a physical constituent of the external world and as a mental agent capable of representing their internal world (e.g., affective and mental states)? We cannot answer these questions yet because we still lack a scientific theory of how a baby might represent, model, and learn about the self.

It may be tempting to offload these questions onto the realm of learning from experience, treating them as a topic for *How Children Learn* rather than *What Babies Know*. However, if the core knowledge proposal is meant to be a scientific account of the origins of human intelligence – an evolved system of knowledge that is adaptive for survival and learning – it is almost impossible to imagine how this knowledge could be so rich, extensive, and coherent in content yet so passive in nature, detached from the organism itself.

The idea that humans know and learn about the self is far from new. Early philosophers have theorized about the role of introspection (Descartes), first-person sensory experience (Hume), and attribution of self-concept (Locke). William James' distinction between the self as "Me" (the empirical self constructed through interactions with the world) versus "I" (the subjective self capable of representing "Me"; James, 1890; see also Gergely, 2002; Neisser, 1988) and Piaget's emphasis on the self and agentive experience (Piaget, 1952, 1954) have provided grounds for the initial empirical findings that hint at a nascent sense of self-awareness in infancy, especially regarding their own bodies and actions. By 3 months of age, infants not only recognize their own movements (Rochat & Striano, 2000) but also grasp the contingency between their own actions and changes in the physical world (e.g., kicking and the movement of a crib mobile; Rovee & Rovee, 1969) and in their interactions with others (Bigelow & Rochat, 2006; Gergely & Watson, 1999; Nadel, Carchon, Kervella, Marcelli, & Reserbat-Plantey, 1999). These studies suggest that infants at this age can already represent themselves as both physical and social entities that exert influence on the external world. Although questions remain about the richness of such understanding, these findings nonetheless highlight the need to characterize both the nature and the content of the knowledge that could give rise to these behaviors, and the relationship between a baby's representation of the self and their active action planning.

The key functions of the human mind – perception and cognition – are often conceptualized in the context of action. The limited behavioral repertoire of young infants may perhaps explain why infants have often been characterized as observers looking out the window. While limited, however, infants' behaviors are systematic; researchers can use infants' looks to tap into their minds precisely because these behaviors reflect what infants perceive and think. Within a few months of life, infants also begin to reach, vocalize, and make facial expressions in ways that suggest a mapping between themselves and their physical and social surroundings. It is possible that the absence of these behaviors in very young infants reflects little more than their inability to execute motor actions; prereaching infants may *want* to approach or attain an object but are limited in their capacity to do so. If core knowledge is meant to be adaptive for survival and learning, then it seems reasonable to expect this system of knowledge to serve as the interface between perception and action. Even if this knowledge may not yet manifest as explicit, interpretable actions, it should nonetheless allow infants to use their own experience to

learn about themselves and the external world. This critical link between perception and action, which must incorporate at least a crude representation of the self, is curiously absent from the core knowledge proposal.

Spelke's book, in fact, offers many examples of studies where the baby as a participant is more than just an onlooker. Infants visually recognize objects based on the motion patterns they themselves produced via haptic exploration (Streri & Spelke, 1988) and preferentially explore objects that violated their expectations to test for the specific violations they had observed (Stahl & Feigenson, 2015); toddlers use their (shockingly limited) representation of space to orient themselves (Hermer & Spelke, 1994); infants direct their gaze based on the contingency between their own behaviors and another agent (Johnson, Slaughter, & Carey, 1998), and the list goes on. Although some of these findings come from relatively older infants and toddlers, the systematic relationship between what they observe and what they do as a consequence nonetheless raises important questions about the developmental origins of the interface between perception, knowledge, and action.

One particularly striking example comes from studies using the "sticky mittens" training. The "Woodward effect" (Woodward, 1998), indicating an abstract understanding of others' object-directed goals, is observed only when infants view actions that they themselves can perform. Intriguingly, 3-month-olds (who cannot yet reach for objects) show this effect after a short training with a velcro-lined mitten that allows them to entrain objects without performing targeted reach and grasp (Sommerville, Woodward, & Needham, 2005; see also Skerry, Carey, & Spelke, 2013, for the effect of this training on expectation of action efficiency). Why does this training work? Spelke argues against the interpretation that the training allows prereaching infants to *learn* about goals; they may represent goals but do not yet understand what constitutes goal-directed actions, and the training highlights physical contact as a key feature of causally meaningful actions (Liu, Brooks, & Spelke, 2019). This compelling argument, however, still presupposes something that remains unaddressed by core knowledge; the baby understands the equivalence between its own hand movement and someone else's hand movement (i.e., self-other mapping). Although Chapter 8 (Core Social Cognition) begins to hint at a core system that incorporates the baby itself (Sect. 8.2: Infants' Sensitivity to Social Engagement), perhaps because of the difficulty of directly studying infants' internal, phenomenal experiences, many questions remain open about the nature of knowledge that supports infants' ability to formulate their own goals.

Stepping back, the discrepancy between what babies *know* and what babies *do* gives us an opportunity to reflect upon how the field has progressed in the last few decades. Moving on from behaviorism and away from the initial Piagetian emphasis on agentive experience, infant cognition research has been a remarkably fruitful enterprise filled with seminal discoveries on infants' ability to think and represent their surroundings. Although many of these studies have used methods that place infants in a relatively passive position as an observer, others have leveraged self-initiated behaviors such as manual search (e.g., Feigenson & Carey, 2003), crawling (e.g., Denison & Xu, 2010), object-based exploration (e.g., Gweon, Tenenbaum, & Schulz, 2010), and even socially oriented actions such as help-seeking (e.g., Goupil, Romand-Monnier, & Kouider, 2016; Gweon & Schulz, 2011). Although these behaviors primarily serve as "dependent measures" that are meant to inform researchers about the representations and inferential processes that reside in infants' minds, they are not to be taken for granted; they emerge



only when babies know enough to understand what is going on, and are motivated enough to act.

More broadly, infants' understanding of the external world has remained a rather separate topic of scientific inquiry from infants' understanding of their own bodies and locomotor ability (e.g., Adolph & Hoch, 2019; Rochat & Striano, 2000) or the motivational drive that underlies their own goal-directed actions (e.g., Dweck, 2017). Yet, infants are clearly motivated to perform actions that exert a systematic influence on the external world (Rovee & Rovee, 1969), are puzzled when their actions fail and driven to figure out why (Gweon & Schulz, 2011), and actively seek information about their own motor abilities (Adolph & Hoch, 2019). Beyond the self as a physical agent, as noted in the book, infants begin to use their own social relationships to reason about potential social partners (Thomas, Saxe, & Spelke, 2022). By preschool years, children expect their parent, but not the experimenter's parent, to know private information about themselves (Chuey, Jara-Ettinger, & Gweon, 2023); they also deliberately seek information about what others think of them (e.g., Zhu, Dweck, & Gweon, 2023) and even try to manage these representations by communicating about themselves (e.g., Asaba & Gweon, 2022; Heyman, Compton, Amemiya, Ahn, & Shao, 2021). These findings suggest that an abstract, socially constructed self is already present by the preschool years, and likely develops earlier.

By imagining how core knowledge could incorporate aspects of the self, perhaps we could inch closer toward a unified theory that explains how humans come to build a model of the world that incorporates themselves. One possibility is that core knowledge in each of the key domains have a placeholder for "the self"; for instance, infants' representations of space might incorporate information about one's own position, orientation, and movement, as suggested by the discovery of place cells, grid cells, speed cells, and head-orientation cells that together may give rise to a sense of the physical self in space and time (Moser et al., 2014). The recent discovery of social place-cells (Danjo, Toyozumi, & Fujisawa, 2018; Omer, Maimon, Las, & Ulanovsky, 2018) that represent the spatial locations of conspecifics with respect to the self also supports the idea that the core system for spatial navigation has carved out a "place" for the self (note: it also raises a question about a strong version of modularity). Another possibility is that there is a separate domain of core knowledge about the self, which might support the use of proprioceptive senses for haptic feedback and motor coordination. Additionally, core knowledge of agents (Ch. 7) and core social cognition (Ch. 8) might also provide particularly useful grounds for building a model of the self that represents the self as an agent in relation to other agents, and learns about its properties through self-guided exploration of objects and interactions with others. These possibilities are not mutually exclusive, and can offer ways in which core knowledge can make contact with infants' understanding of the self.

By synthesizing decades of work, Spelke's book already offers a rich foundation for new methods and approaches for exploring these possibilities. If core knowledge serves to make learning more tractable, understanding how babies represent themselves may provide key clues that can help explain what makes humans such powerful learners: Humans compete against themselves to improve in ways that are far more efficient than the most powerful machines today, deploy their own intelligence to problem-solve both for themselves and for others, and perhaps most importantly, appreciate themselves as entities that can bond, interact, and communicate with other beings.

**Acknowledgments.** The authors are grateful to Aaron Chuey, Michael C. Frank, and Julian Jara-Ettinger for helpful discussions.

**Financial support.** This commentary was supported by a James S. McDonnell Foundation Scholars Award for Understanding Human Cognition, a Jacobs Foundation Early Career Research Fellowship, and an NSF Award BCS-2019567.

**Competing interest.** None.

## References

- Adolph, K. E., & Hoch, J. E. (2019). Motor development: Embodied, embedded, enculturated, and enabling. *Annual Review of Psychology*, 70(1), 141–164. <https://doi.org/10.1146/annurev-psych-010418-102836>
- Asaba, M., & Gweon, H. (2022). Young children infer and manage what others think about them. *Proceedings of the National Academy of Sciences of the United States of America*, 119(32), e2105642119. <https://doi.org/10.1073/pnas.2105642119>
- Bigelow, A. E., & Rochat, P. (2006). Two-month-old infants' sensitivity to social contingency in mother–infant and stranger–infant interaction. *Infancy*, 9(3), 313–325. [https://doi.org/10.1207/s15327078in0903\\_3](https://doi.org/10.1207/s15327078in0903_3)
- Chuey, A., Jara-Ettinger, J., & Gweon, H. (2023). Violation of epistemic expectations: Children monitor what others know and recognize unexpected sources of knowledge. In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society* (pp. 2751–2758). Austin, TX: Cognitive Science Society.
- Danjo, T., Toyozumi, T., & Fujisawa, S. (2018). Spatial representations of self and other in the hippocampus. *Science*, 359(6372), 213–218. <https://doi.org/10.1126/science.aa03898>
- Denison, S., & Xu, F. (2010). Twelve- to 14-month-old infants can predict single-event probability with large set sizes: Twelve- to 14-month-olds can predict single-event probability. *Developmental Science*, 13(5), 798–803. <https://doi.org/10.1111/j.1467-7687.2009.00943.x>
- Dweck, C. S. (2017). From needs to goals and representations: Foundations for a unified theory of motivation, personality, and development. *Psychological Review*, 124(6), 689–719. <https://doi.org/10.1037/rev0000082>
- Feigenson, L., & Carey, S. (2003). Tracking individuals via object-files: Evidence from infants' manual search. *Developmental Science*, 6(5), 568–584. <https://doi.org/10.1111/1467-7687.00313>
- Gergely, G. (2002). The development of understanding self and agency. In U. Goswami (Ed.), *Blackwell handbook of childhood cognitive development* (pp. 26–46). Blackwell Publishers.
- Gergely, G., & Watson, J. S. (1999). Early socio-emotional development: Contingency perception and the social-biofeedback model. In P. Rochat (Ed.), *Early social cognition: Understanding others in the first months of life* (Vol. 60, pp. 101–136). Taylor & Francis.
- Goupil, L., Romand-Monnier, M., & Kouider, S. (2016). Infants ask for help when they know they don't know. *Proceedings of the National Academy of Sciences of the United States of America*, 113(13), 3492–3496. <https://doi.org/10.1073/pnas.1515129113>
- Gweon, H., & Schulz, L. E. (2011). 16-month-olds rationally infer causes of failed actions. *Science*, 332(6037), 1524. <https://doi.org/10.1126/science.1204493>
- Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2010). Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences of the United States of America*, 107(20), 9066–9071. <https://doi.org/10.1073/pnas.1003095107>
- Hermer, L., & Spelke, E. S. (1994). A geometric process for spatial reorientation in young children. *Nature*, 370(6484), 57–59. <https://doi.org/10.1038/370057a0>
- Heyman, G. D., Compton, A. M., Amemiya, J., Ahn, S., & Shao, S. (2021). Children's reputation management: Learning to identify what is socially valued and acting upon it. *Current Directions in Psychological Science*, 30(4), 315–320. <https://doi.org/10.1177/09637214211009516>
- James, W. (1890). The consciousness of self. In W. James (Ed.), *The principles of psychology* (Vol. 1, pp. 291–401). Henry Holt and Co.
- Johnson, S., Slaughter, V., & Carey, S. (1998). Whose gaze will infants follow? Features that elicit gaze-following in 12-month-olds. *Developmental Science*, 1, 233–238. [https://doi.org/10.1016/S0163-6383\(96\)90581-0](https://doi.org/10.1016/S0163-6383(96)90581-0)
- Liu, S., Brooks, N. B., & Spelke, E. S. (2019). Origins of the concepts cause, cost, and goal in prereaching infants. *Proceedings of the National Academy of Sciences of the United States of America*, 2, 201904410–201904416. <https://doi.org/10.1073/pnas.1904410116>
- Moser, E. I., Roudi, Y., Witter, M. P., Kentros, C., Bonhoeffer, T., & Moser, M.-B. (2014). Grid cells and cortical representation. *Nature Reviews Neuroscience*, 15(7), 466–481. <https://doi.org/10.1038/nrn3766>
- Nadel, J., Carchon, I., Kervella, C., Marcelli, D., & Reserbat-Plantey, D. (1999). Expectancies for social contingency in 2-month-olds. *Developmental Science*, 2(2), 164–173. <https://doi.org/10.1111/1467-7687.00065>
- Neisser, U. (1988). Five kinds of self-knowledge. *Philosophical Psychology*, 1(1), 35–59. <https://doi.org/10.1080/09515088808572924>

Omer, D. B., Maimon, S. R., Las, L., & Ulanovsky, N. (2018). Social place-cells in the bat hippocampus. *Science*, 359(6372), 218–224. <https://doi.org/10.1126/science.aaa3474>

Piaget, J. (1952). *The origins of intelligence in children*. International Universities Press.

Piaget, J. (1954). *The construction of reality in the child*. Routledge.

Rochat, P., & Striano, T. (2000). Perceived self in infancy. *Infant Behavior and Development*, 23(3–4), 513–530. [https://doi.org/10.1016/S0163-6383\(01\)00055-8](https://doi.org/10.1016/S0163-6383(01)00055-8)

Rovee, C. K., & Rovee, D. T. (1969). Conjugate reinforcement of infant exploratory behavior. *Journal of Experimental Child Psychology*, 8(1), 33–39. [https://doi.org/10.1016/0022-0965\(69\)90025-3](https://doi.org/10.1016/0022-0965(69)90025-3)

Skerry, A. E., Carey, S. E., & Spelke, E. S. (2013). First-person action experience reveals sensitivity to action efficiency in prereaching infants. *Proceedings of the National Academy of Sciences of the United States of America*, 110(46), 18728–18733. <https://doi.org/10.1073/pnas.1312322110>

Sommerville, J. A., Woodward, A. L., & Needham, A. (2005). Action experience alters 3-month-old infants' perception of others' actions. *Cognition*, 96(1), B1–B11. <https://doi.org/10.1016/j.cognition.2004.07.004>

Spelke, E. S. (2022). *What babies know: Core knowledge and composition*. Oxford University Press.

Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, 348(6230), 91–94. <https://doi.org/10.1126/science.aaa3799>

Streri, A., & Spelke, E. S. (1988). Haptic perception of objects in infancy. *Cognitive Psychology*, 20(1), 1–23. [https://doi.org/10.1016/0010-0285\(88\)90022-9](https://doi.org/10.1016/0010-0285(88)90022-9)

Thomas, A. J., Saxe, R., & Spelke, E. S. (2022). Infants infer potential social partners by observing the interactions of their parent with unknown others. *Proceedings of the National Academy of Sciences of the United States of America*, 119(32), e2121390119. <https://doi.org/10.1073/pnas.2121390119>

Woodward, A. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, 69(1), 1–34. [https://doi.org/10.1016/S0010-0277\(98\)00058-4](https://doi.org/10.1016/S0010-0277(98)00058-4)

Zhu, P., Dweck, C., & Gweon, H. (2023). Young children's curiosity about what others think about the self. In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society* (pp. 835–842). Austin, TX: Cognitive Science Society

## Evidence for core social goal understanding (and, perhaps, core morality) in preverbal infants

J. Kiley Hamlin 

Department of Psychology, University of British Columbia, Vancouver, BC, Canada  
[cic.psych.ubc.ca](http://cic.psych.ubc.ca)

### Corresponding author:

J. Kiley Hamlin; Email: [Kiley.hamlin@psych.ubc.ca](mailto:Kiley.hamlin@psych.ubc.ca)

doi:10.1017/S0140525X23003059, e130

### Abstract

Spelke's *What Babies Know* masterfully describes infants' impressive repertoire of core cognitive concepts, from which the suite of human knowledge is eventually built. The current commentary argues for the existence of a core concept that Spelke claims preverbal infants lack: social goal. Core social goal concepts, operative extremely early in human development, underlie infants' basic abilities to interpret and evaluate entities within the moral world; such abilities support claims for a core moral domain.

### 1. What babies know

In *What Babies Know* (2022), legendary cognitive scientist Elizabeth Spelke reviews decades of evidence supporting the claim that humans possess “core knowledge”: Innate systems of

abstract knowledge that are evolutionarily ancient, early-emerging, and invariant over development; within these systems knowledge operates automatically, unconsciously, and independently of belief. Spelke provides a detailed account of the evidence for and against six systems, or domains, of core knowledge: Objects, number, places, forms, agents, and social beings. Each system contains skeletal versions of particular concepts (e.g., how objects behave), allowing inexperienced human and nonhuman animals to reason productively about the entities within it.

Crucial to Spelke's analysis of which concepts are and are not part of core knowledge is the claim that each domain competes with the others for attention, resulting in an initial failure of domains to communicate with each other. This failure precludes thinking about any concept that involves entities from more than one domain, thereby limiting the range of concepts that can be core to the human mind. Spelke argues that these limitations are overcome with the emergence of language around 10 months; language facilitates communication between domains and allows children to begin constructing myriad novel, non-core concepts.

### 2. What babies don't know: Social goals

After reviewing concepts she believes young infants possess, Spelke devotes significant discussion to one she believes they lack: social goal. Specifically, although two of Spelke's six core domains are for reasoning about the social world (the agent and social being systems), Spelke holds that neither is capable of social goal understanding by itself. On the one hand, the core agent system solely considers agents' physical and instrumental goals; it does not consider goals underlying agent-directed actions (e.g., social looking, communication), nor instrumental actions undertaken for social reasons (e.g., cooperating, helping, hindering, other prosocial and antisocial acts). On the other hand, the core social being system solely considers engagements between social entities and between those entities and infants themselves; it does not consider the mental states driving those engagements. Given that understanding social goals requires thinking about social beings as agents whose mental states refer to other agents, it requires communication between the agent and social being systems. Because this communication is impossible prior to language, the concept social goal cannot be part of core knowledge.

Although the observation that core knowledge is limited and the provision of an explanation for how its limitations are eventually overcome are crucial to a proper account of the origins of human cognition, it is curious why Spelke spends so much time arguing that core knowledge lacks social goal understanding. Presumably, this is because she wishes to argue against a potential additional domain of core knowledge that has recently attracted attention: core morality. Indeed, a growing number of cognitive scientists have recently argued that human moral systems are supported by evolved, domain-specific mechanisms for thinking about the moral world (e.g., Baumard, André, & Sperber, 2013; Buyukozer Dawkins, Sloane, & Baillargeon, 2019; Hamlin, 2013a; Krebs, 2008; Macnamara, 1991; Mikhail, 2011; Premack, 2007; Woo, Tan, & Hamlin, 2022; Wynn & Bloom, 2014). Among other things, these mechanisms might allow inexperienced humans to identify morally relevant (inter)actions, evaluate those actions and agents who engage in them as positive or negative, and generate expectations for further actions the agents might perform. All such capacities would, at minimum, require an understanding of social goals.

Although not exclusively, much recent argumentation for core morality cites evidence that preverbal infants appear to positively and negatively evaluate agents based on their prosocial and antisocial acts (for review, see, Margoni & Surian, 2018; Woo et al., 2022), and/or possess expectations for the prosocial versus antisocial acts that individuals are likely to perform in distinct contexts (see Buyukozer Dawkins et al., 2019). Spelke acknowledges that this evidence could suggest that young infants understand social goals, but argues that it actually does not, for two reasons. First, much of the evidence involves infants older than 10 months. At 10 months, the core agent and social being systems can communicate, meaning a concept of social goals can be constructed rather than “core.” Second, evidence with infants <10 months has a viable alternative explanation that does not implicate social goal understanding. Specifically, young infants’ preferences for helpers over hinderers (e.g., Hamlin & Wynn, 2011; Hamlin, Wynn, & Bloom, 2007) need not reflect their understanding of (pro- and/or anti-) social goals, as those preferences may stem from mere sensitivity to social beings in states of engagement (or not) with other social beings: Social engagement is handled by the social being system alone.

Specifically, Spelke points out that prototypical helpful/unhelpful acts shown to infants can also be described as *imitative/not imitative*: In order to help, the helper generally reproduces the actions of a needy protagonist, whereas the hinderer produces opposing actions (see, e.g., Powell & Spelke, 2018). Under Spelke’s core social being system, imitation is a powerful cue that one social being is engaged with another, but certainly lacks moral content (see also Powell, 2022). Thus, Spelke holds that much of the evidence used to argue for core moral capacities has no moral content after all.

In what follows, I review evidence that although infants appear sensitive to cues to social engagement like imitation, they can and do reason about social goals prior to 10 months. Indeed, consistent with claims for core morality, preverbal infants may be *particularly* sensitive to social goals with moral content, including helping/hindering, protection/harm, and fairness/unfairness. Because of space constraints I can only touch on the relevant evidence below; interested readers can find more detailed discussion elsewhere (Hamlin, 2023; Woo, Tan, Yuen, & Hamlin, 2023).

### 3. Evidence infants’ preference for helpers cannot be explained by imitation

Spelke argues that young infants’ preference for helpers reflects sensitivity to imitation rather than to prosocial/antisocial goals. Indeed, Spelke and her former student Lindsay Powell have demonstrated that infants prefer imitators over non-imitators within scenarios purported to demonstrate infants’ preferences for helpers over hinderers (Powell & Spelke, 2018). Although Powell (2022) argues that infants’ imitator preference itself reflects social goal understanding (e.g., that one agent has adopted another’s “utility,” or goal, as its own), Spelke argues that infants could instead prefer imitators without representing utility adoption/goals at all, by inferring that imitators are engaged with their targets.

Inconsistent with Spelke’s analysis, several studies now suggest that infants’ preferences focus on helping rather than imitating. For instance, infants’ preferences rely on their understanding of a needy protagonist’s goals. Hamlin (2015) manipulated whether or not 6–10-month-olds could recognize the goal of trying but failing to climb a hill, by showing some infants the protagonist’s

eyes pointing toward the hilltop (suggesting a goal to reach the top) and others the protagonist’s eyes pointing away from the hilltop (rendering its goal ambiguous). Critically, in both conditions one character imitated the protagonist (pushed it up), and another character did not imitate the protagonist (pushed it down). Critically, only those infants who saw the protagonist looking toward the hilltop, demonstrating a clear unfulfilled goal, preferred the pusher-upper (here, a helper) to the pusher-downer (a hinderer). Similarly, Tan and Hamlin (2022) showed that infants’ own looking toward the hilltop during the protagonist’s failed attempts, arguably indicating goal inference (see Elsner & Adam, 2020), predicted their individual preference for the agent who pushed it to the top: Only those infants who ever looked to the top of the hill preferred the “helper.”

Other studies more directly compare helpers and imitators. For instance, in Hamlin et al. (2007) and Chae and Song (2018), 6- and 10-month-olds were asked to choose between an agent who pushed a needy protagonist uphill, moving like the protagonist *and* causing it to achieve its goal, and an agent who moved up the hill in exactly the same way, but independently from the protagonist. Here, infants chose between an imitative helper and a mere imitator, and consistently selected the helper. In a study from Spelke’s own laboratory led by her former student Brandon Woo (Woo & Spelke, 2023), 8-month-olds were led to infer that an agent’s goal was one of the two possible options, either to open a specific box or to obtain a specific toy. Subsequently, one agent facilitated the goal they inferred the agent to have, whereas the other agent facilitated the other goal; critically for the present purposes, in one condition the helpful agent was less imitative. Here again, infants’ choices suggested they consistently preferred helpers, but not imitators.

In each of the above studies, one character’s actions always matched the protagonist’s more closely than the other’s. However, infants only chose those imitators whose actions were also helpful: They preferred helpful imitators over mere imitators, and failed to distinguish differentially imitative characters who were not differentially helpful. Thus, Spelke’s claim that young infants’ helper preferences can be reduced to preferences for imitators seems unlikely.

### 4. Evidence young infants can represent (pro- and anti-) social goals

Although the above work suggests that infants are more sensitive to helping than to imitating, it need not indicate that infants represent that helpers possess the prosocial goal to facilitate another’s goal or that hinderers possess the antisocial goal to prevent a goal. Indeed, perhaps young infants merely represent whether or not one agent *causes* another to achieve its goal (which imitators do not do). Of course, adults’ moral concepts privilege others’ prosocial and antisocial *intentions* (e.g., Cushman, 2008; Malle, 1999): Do young infants also consider social intentions?

Multiple studies now suggest that infants represent and evaluate prosocial and antisocial intentions before 10 months (see also Hamlin, Ullman, Tenenbaum, Goodman, & Baker, 2013; Kanakogi et al., 2017; Strid & Meristo, 2020; Woo, Steckler, Le, & Hamlin, 2017). First, Hamlin (2013b) demonstrated that 8-month-olds privileged intentions over outcomes in their preferences for pro- and antisocial others; for instance, preferring an agent who tried but failed to help a protagonist achieve its unfulfilled goal over an agent who tried but failed to hinder the protagonist. Second, Woo and Spelke (2022) showed that, remarkably,



8-month-olds preferred an agent who *believed* it was helping, even though it was not, over an agent who believed it was not helping, even though it was. Finally, Geraci and Surian (2023) and Geraci, Simion, and Surian (2022) demonstrated that 4- and 9-month-olds preferred an agent who tried but failed to distribute resources equally between two recipients over one who tried but failed to distribute resources unequally. Because no resources were ever actually given out, infants' choices must have been based on intent; further control conditions suggest that it was not that infants simply like agents appearing to have more social partners (for related controls with acts of protection, see Kanakogi et al., 2017). These papers suggest that even young infants understand prosocial and/or antisocial intentions within two morally relevant domains, at times at or before they appear to be able to recognize intention in nonmoral contexts (for discussion, see Woo et al., 2023).

## 5. Conclusion

In sum, in contrast to Spelke's claims, young infants appear to possess the concept social goal. Although this may seem like a rather insignificant topic of debate, an inability to understand social goals early in development would, in turn, render recent arguments that humans possess core *moral* capacities (e.g., Baumard et al., 2013; Buyukozer Dawkins et al., 2019; Hamlin, 2013a; Krebs, 2008; Macnamara, 1991; Mikhail, 2011; Premack, 2007; Woo et al., 2022; Wynn & Bloom, 2014) moot, given that many if not most moral concepts are fundamentally rooted in notions of social goals. Of course, there is more to moral concepts than social goals, and any effective claim to core morality will include much more data than reviewed here, including from animals and diverse adult humans (see, e.g., Anderson et al., 2017; Anderson, Kuroshima, Takimoto, & Fujita, 2013a; Anderson, Takimoto, Kuroshima, & Fujita, 2013b; Brosnan, 2023; Cosmides, Guzmán, & Tooby, 2018; Cosmides & Tooby, 1992; Curry, Mullins, & Whitehouse, 2019; Darden, James, Cave, Brask, & Croft, 2020; Isik, Koldewyn, Beeler, & Kanwisher, 2017; Mikhail, 2011). Spelke's *What Babies Know* (2022) provides an enviable model of what effective arguments for core moral knowledge – and indeed any account of the origins of knowledge more generally – must look like.

**Financial support.** This work was supported by the National Sciences and Engineering Research Council of Canada, RGPIN-2022-05037.

**Competing interests.** None.

## References

- Anderson, J. R., Bucher, B., Chijiwa, H., Kuroshima, H., Takimoto, A., & Fujita, K. (2017). Third-party social evaluations of humans by monkeys and dogs. *Neuroscience & Biobehavioral Reviews*, 82, 95–109. <https://doi.org/10.1016/j.neubiorev.2017.01.003>
- Anderson, J. R., Kuroshima, H., Takimoto, A., & Fujita, K. (2013a). Third-party social evaluation of humans by monkeys. *Nature Communications*, 4(1), 1561. <https://doi.org/10.1038/ncomms2495>
- Anderson, J. R., Takimoto, A., Kuroshima, H., & Fujita, K. (2013b). Capuchin monkeys judge third-party reciprocity. *Cognition*, 127(1), 140–146. <https://doi.org/10.1016/j.cognition.2012.12.007>
- Baumard, N., André, J.-B., & Sperber, D. (2013). A mutualistic approach to morality: The evolution of fairness by partner choice. *Behavioral and Brain Sciences*, 36(1), 59–78. <https://doi.org/10.1017/S0140525X11002202>
- Brosnan, S. F. (2023). A comparative perspective on the human sense of justice. *Evolution and Human Behavior*, 44(3), 242–249. <https://doi.org/10.1016/j.evolhumbehav.2022.12.002>
- Buyukozer Dawkins, M., Sloane, S., & Baillargeon, R. (2019). Do infants in the first year of life expect equal resource allocations? *Frontiers in Psychology*, 10, 116. <https://doi.org/10.3389/fpsyg.2019.00116>
- Chae, J. K., & Song, H. (2018). Negativity bias in infants' expectations about agents' dispositions. *British Journal of Developmental Psychology*, 36(4), 620–633. <https://doi.org/10.1111/bjdp.12246>
- Cosmides, L., Guzmán, R. A., & Tooby, J. (2018). The evolution of moral cognition. In A. Zimmerman, K. Jones, & M. Timmons (Eds.), *Routledge handbook on moral epistemology* (pp. 174–228). Routledge.
- Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. In *The adapted mind: Evolutionary psychology and the generation of culture* (pp. 163–228). Oxford University Press.
- Curry, O. S., Mullins, D. A., & Whitehouse, H. (2019). Is it good to cooperate? Testing the theory of morality-as-cooperation in 60 societies. *Current Anthropology*, 60(1), 47–69. <https://doi.org/10.1086/701478>
- Cushman, F. (2008). Crime and punishment: Distinguishing the roles of causal and intentional analyses in moral judgment. *Cognition*, 108(2), 353–380. <https://doi.org/10.1016/j.cognition.2008.03.006>
- Darden, S. K., James, R., Cave, J. M., Brask, J. B., & Croft, D. P. (2020). Trinidadian guppies use a social heuristic that can support cooperation among non-kin. *Proceedings of the Royal Society B: Biological Sciences*, 287(1934), 20200487. <https://doi.org/10.1098/rspb.2020.0487>
- Elsner, B., & Adam, M. (2020). Infants' goal prediction for simple action events: The role of experience and agency cues. *Topics in Cognitive Science*, 13(1), 45–62. <https://doi.org/10.1111/tops.12494>
- Geraci, A., Simion, F., & Surian, L. (2022). Infants' intention-based evaluations of distributive actions. *Journal of Experimental Child Psychology*, 220, 105429. <https://doi.org/10.1016/j.jecp.2022.105429>
- Geraci, A., & Surian, L. (2023). Intention-based evaluations of distributive actions by 4-month-olds. *Infant Behavior and Development*, 70, 101797. <https://doi.org/10.1016/j.infbeh.2022.101797>
- Hamlin, J. K. (2013a). Moral judgment and action in preverbal infants and toddlers: Evidence for an innate moral core. *Current Directions in Psychological Science*, 22(3), 186–193. <https://doi.org/10.1177/0963721412470687>
- Hamlin, J. K. (2013b). Failed attempts to help and harm: Intention versus outcome in preverbal infants' social evaluations. *Cognition*, 128(3), 451–474. <https://doi.org/10.1016/j.cognition.2013.04.004>
- Hamlin, J. K. (2015). The case for social evaluation in preverbal infants: Gazing toward one's goal drives infants' preferences for helpers over hinderers in the hill paradigm. *Frontiers in Psychology*, 5, 1563. <https://doi.org/10.3389/fpsyg.2014.01563>
- Hamlin, J. K. (2023). Core morality? Or merely core agents and social beings? A response to Spelke's *What babies know*. *Mind and Language*, 38, 1323–1335.
- Hamlin, J. K., Ullman, T., Tenenbaum, J., Goodman, N., & Baker, C. (2013). The mentalistic basis of core social cognition: Experiments in preverbal infants and a computational model. *Developmental Science*, 16(2), 209–226. <https://doi.org/10.1111/desc.12017>
- Hamlin, J. K., & Wynn, K. (2011). Young infants prefer prosocial to antisocial others. *Cognitive Development*, 26(1), 30–39. <https://doi.org/10.1016/j.cogdev.2010.09.001>
- Hamlin, J. K., Wynn, K., & Bloom, P. (2007). Social evaluation by preverbal infants. *Nature*, 450(7169), 557–559. <http://dx.doi.org/10.1038/nature06288>
- Isik, L., Koldewyn, K., Beeler, D., & Kanwisher, N. (2017). Perceiving social interactions in the posterior superior temporal sulcus. *Proceedings of the National Academy of Sciences of the United States of America*, 114(43), E9145–E9152. <https://doi.org/10.1073/pnas.1714471114>
- Kanakogi, Y., Inoue, Y., Matsuda, G., Butler, D., Hiraki, K., & Myowa-Yamakoshi, M. (2017). Preverbal infants affirm third-party interventions that protect victims from aggressors. *Nature Human Behaviour*, 1(2), 1–7. <https://doi.org/10.1038/s41562-016-0037>
- Krebs, D. L. (2008). Morality: An evolutionary account. *Perspectives on Psychological Science*, 3(3), 149–172. <https://doi.org/10.1111/j.1745-6924.2008.00072.x>
- Macnamara, J. (1991). The development of moral reasoning and the foundations of geometry. *Journal for the Theory of Social Behaviour*, 21(2), 125–150.
- Malle, B. F. (1999). How people explain behavior: A new theoretical framework. *Personality and Social Psychology Review*, 3(1), 23–48. [https://doi.org/10.1207/s15327957pspr0301\\_2](https://doi.org/10.1207/s15327957pspr0301_2)
- Margoni, F., & Surian, L. (2018). Infants' evaluation of prosocial and antisocial agents: A meta-analysis. *Developmental Psychology*, 54(8), 1445–1455. <https://doi.org/10.1037/dev0000538>
- Mikhail, J. (2011). *Elements of moral cognition: Rawls' linguistic analogy and the cognitive science of moral and legal judgment*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511780578>
- Powell, L. J. (2022). Adopted utility calculus: Origins of a concept of social affiliation. *Perspectives on Psychological Science*, 17(5), 1215–1233. <https://doi.org/10.1177/17456916211048487>
- Powell, L. J., & Spelke, E. S. (2018). Third-party preferences for imitators in preverbal infants. *Open Mind*, 2(2), 61–71. [https://doi.org/10.1162/opmi\\_a\\_00018](https://doi.org/10.1162/opmi_a_00018)
- Premack, D. (2007). Foundations of morality in the infant. In O. Vilarroya, & F. Forn i Argimon (Eds.), *Social brain matters* (pp. 161–167). Brill. [https://doi.org/10.1163/9789401204491\\_015](https://doi.org/10.1163/9789401204491_015)
- Spelke, E. S. (2022). *What babies know: Core knowledge and composition* (Vol. 1). Oxford University Press.

- Strid, K., & Meristo, M. (2020). Infants consider the distributor's intentions in resource allocation. *Frontiers in Psychology, 11*, 596213. <https://doi.org/10.3389/fpsyg.2020.596213>
- Tan, E., & Hamlin, J. K. (2022). Mechanisms of social evaluation in infancy: A preregistered exploration of infants' eye-movement and pupillary responses to prosocial and antisocial events. *Infancy, 27*(2), 255–276. <https://doi.org/10.1111/inf.12447>
- Woo, B. M., & Spelke, E. (2022). Eight-month-old infants' social evaluations of agents who act on false beliefs. *Proceedings of the Annual Meeting of the Cognitive Science Society, 44*, 1184–1189. <https://escholarship.org/uc/item/8k02x1mx>
- Woo, B. M., & Spelke, E. S. (2023). Infants and toddlers leverage their understanding of action goals to evaluate agents who help others. *Child Development, 94*(3), 734–751. doi:10.1111/cdev.13895
- Woo, B. M., Steckler, C. M., Le, D. T., & Hamlin, J. K. (2017). Social evaluation of intentional, truly accidental, and negligently accidental helpers and harmers by 10-month-old infants. *Cognition, 168*, 154–163. <https://doi.org/10.1016/j.cognition.2017.06.029>
- Woo, B. M., Tan, E., & Hamlin, J. K. (2022). Human morality is based on an early-emerging moral core. *Annual Review of Developmental Psychology, 4*(1), 41–61. <https://doi.org/10.1146/annurev-devpsych-121020-023312>
- Woo, B. M., Tan, E., Yuen, F. L., & Hamlin, J. K. (2023). Socially evaluative contexts facilitate mentalizing. *Trends in Cognitive Sciences, 27*(1), 17–29. <https://doi.org/10.1016/j.tics.2022.10.003>
- Wynn, K., & Bloom, P. (2014). The moral baby. In M. Killen & J. G. Smetana (Eds.), *Handbook of moral development* (pp. 435–453). Psychology Press.

## Substances as a core domain

Susan J. Hespos<sup>a,b\*</sup>  and Lance J. Rips<sup>a</sup> 

<sup>a</sup>Psychology Department, Northwestern University, Evanston, IL, USA and

<sup>b</sup>MARCS Institute for Brain, Behaviour, and Cognition, Western Sydney University, Penrith, NSW, Australia

[S.Hespos@westernsydney.edu.au](mailto:S.Hespos@westernsydney.edu.au); [Rips@northwestern.edu](mailto:Rips@northwestern.edu)

[https://www.westernsydney.edu.au/babylab/people/researchers/professor\\_susan\\_hespos/](https://www.westernsydney.edu.au/babylab/people/researchers/professor_susan_hespos/);  
<https://sites.northwestern.edu/ripslab/>

\*Corresponding author.

doi:10.1017/S0140525X23003163, e131

### Abstract

Central to *What Babies Know* (Spelke, 2022) is the thesis that infants' understanding is divided into independent modules of core knowledge. As a test case, we consider adding a new domain: core knowledge of substances. Experiments show that infants' understanding of substances meets some criteria of core knowledge, and they raise questions about the relations that hold between core domains.

*What Babies Know* (Spelke, 2022) summarizes and systematizes several decades of ingenious, influential research on infants' cognition. Its central thesis is that “infants' learning rests on a set of cognitive systems that we share with animals and that evolved over hundreds of millions of years. At least six distinct systems serve to represent highly abstract properties of the unchanging navigable environment, of movable objects, of number, and of the living, animate, and social beings who populate our world” (Spelke, 2022, p. xvii). In this commentary, we ask whether there may be room for one more – a core domain of substances for nonsolid things like liquid or sand – and we examine some consequences of adding it to Spelke's “at least six.”

The key tenet of Spelke's core-knowledge approach is that underneath all the things that vary across humans, there exists a

set of conceptual capacities common to everyone. The research motivated by core knowledge strives to characterize these abilities and their development. In the case of object knowledge, children never receive explicit instruction about how objects behave and interact, yet they draw universally similar expectations about them. For example, they universally expect that hidden objects do not cease to exist when they are hidden from view (Baillargeon, Spelke, & Wasserman, 1985). Such expectations appear to be universal across age groups as well as across individuals of the same age. Expectations about objects are evident in other species as well. Rhesus macaques expect an object to stop when it comes in contact with a wall rather than pass through it (Santos & Hauser, 2002). Humans and chickens have similar expectations about partially occluded objects (Chiandetti & Vallortigara, 2011; Kellman & Spelke, 1983).

However, our world includes more than physical objects. Spelke has highlighted five other core-knowledge domains, arguing that core-knowledge systems are evolutionarily important abilities, each solving a limited set of problems. These systems are encapsulated, that is, they operate independently of other cognitive systems. An advantage of encapsulation is that core abilities are universal and are effortlessly acquired with little experience. A limitation of encapsulation is that these abilities lack flexibility and precision (e.g., Samuels, 2012). Evidence supporting the independence of core systems comes from neurological research identifying specialized brain areas for processing this information, as well as clinical studies showing that these domain-specific abilities can be lost while other cognitive abilities remain intact.

Spelke often cites nonobjects like sand as a contrast case to items in the core domain of objects. Our focus in this commentary is on arguing that instead of being merely a fringe contrast example, substances may be a core domain of their own. As adults, we automatically react differently when we upset a glass of beer than when we upset a bowl of pretzels. These reactions to spills result from our understanding that objects and liquids have different physical properties and so behave differently. These reactions may seem obvious, but when do we develop the notion that liquids deform to fill space whereas solids don't? This ontological categorical distinction has captivated linguists who trace differences in the meaning of quantitative expressions, such as “many pretzels” (vs. \*“much pretzels”) and “much beer” (vs. \*“many beer”) (e.g., Rothstein, 2017; Wellwood, 2019). In the philosophical domain of metaphysics, there are distinctions between entities that come in atomic units (like pretzels) and those that have no clear units (like beer; e.g., Frege, 1980; Koslicki, 1997). In the field of psychology, we look at the origins and development of knowledge about substances and how it compares to knowledge about objects (Hespos & VanMarle, 2012; Rips & Hespos, 2015, 2019). For example, unlike objects, liquids deform to fit a container and a solid object can pass through them. Yet, like objects, liquids are omnipresent, and knowledge of how substances behave is probably universal across cultures and species.

Early evidence suggested that infants had principled expectations about objects, but not about substances (Cherries, Mitroff, Wynn, & Scholl, 2008; Chiang & Wynn, 2000; Huntley-Fenner, Carey, & Solimando, 2002). In one study, Huntley-Fenner et al. (2002) showed infants a pile of sand poured on a stage; then the pile was concealed by a screen, and a second pile of sand was poured behind a nearby but separate screen. The test trials alternated between an expected (by adults) outcome and an unexpected outcome, and looking time was the dependent measure. The expected outcome was to reveal two piles of sand, one behind

each screen. The unexpected outcome was to reveal only a single pile of sand behind one of the screens and nothing behind the other. Infants looked equally at the expected and unexpected outcomes, providing evidence that they did not detect the violation when one sand pile disappeared. In contrast, when the sand was replaced with solid objects that were shaped like sand piles, the infants looked significantly longer at the unexpected test trials. Infants' difficulties in tracking sand extended to collections of objects, like a disassembled pile of Legos (Chiang & Wynn, 2000). Together, these findings were interpreted as evidence that infants have principled expectations for objects but not for substances (Spelke & Kinzler, 2007).

However, evidence for infants' knowledge of substances began to appear in later studies. Bourgeois, Khawar, Neal, and Lockman (2005) introduced a different approach to ask if infants had concepts of distinct materials. They presented infants with entities that varied in whether they were rigid (particle board) versus flexible (sponge) versus liquid (water) versus discontinuous (netting) and found that 6- to 10-month-old infants adjusted their actions toward the entities based on their material-specific qualities. This finding was important because it demonstrated that infants applied different behaviors to objects and liquids.

Our first experiment on substance knowledge asked whether infants have material-specific ideas about liquids. Using a looking paradigm, we habituated infants to either a glass that contained a liquid or a glass that had solid contents but was otherwise perceptually similar. The glass was tipped back and forth, and the motion cues revealed whether the contents were liquid or solid. Next, in the test trials, all infants saw a straw lowered into the glass. On half the trials, the contents of the glass were liquid, and the straw penetrated the surface of the liquid, coming to rest at the bottom of the glass. On the other half of the trials, the contents of the glass were solid, and the straw stopped when it met the surface of the solid. Infants dishabituated (i.e., showed a significant increase in looking time compared to their last habituation trials) when there was a state change from liquid to solid or from solid to liquid. These studies show that infants have distinct ideas about how objects and liquids behave (Hespos, Ferry, & Rips, 2009).

We started with a water-like substance because it is the most prevalent example of its kind. However, core principles go beyond information in the immediate environment. Our initial findings raised questions about how far infants' ideas about substances extend. Do infants develop expectations about liquids because of their experience in drinking and bathing or would they generalize such expectations to unfamiliar events with similar physical attributes? More specifically, would an infant who has never been to a beach know that the sand in a cup should pour out and not tumble? Our next study provided a positive answer (Hespos, Ferry, Anderson, Hollenbeck, & Rips, 2016). The events were like those of the liquid experiment except the liquid was replaced with sand. We again found that the infants distinguished sand from a solid object. Together these findings provide evidence that knowledge of substances emerges early, based on little or no experience.

The results we have presented suggest that infants can grasp simple physical properties that apply to nonsolid substances, and Spelke now acknowledges these findings (Spelke, 2022, pp. 62–63, footnote 8). However, the previous research showing success with objects and failure with substances in otherwise identical paradigms poses questions about the extent of infants' knowledge (Chiang & Wynn, 2000; Huntley-Fenner et al., 2002; Rosenberg & Carey, 2009). Infants seem unable to predict the

number of piles that result from pouring sand behind adjacent screens. But could this be because of the working memory demands of the pouring event rather than to lack of substance knowledge? We tested infants' expectations about a simplified pouring event. Although nonsolid substances can sometimes spread to fit the space allotted, constraints particular to sand limit its ability to do so. If infants see two cups of sand poured at opposite ends of a tray behind a screen, would it violate their expectations to reveal a single pile? What if just one cup was poured behind a screen? Would it be surprising if a single pour resulted in two separate piles? Our findings provide a yes answer to both questions (Anderson, Hespos, & Rips, 2018).

The emerging picture is that infants have ideas about how substances behave, divide, and accumulate. These ideas are distinct from those governing how objects behave and interact, but they have at least some of the standard core properties (Spelke, 2022, Ch. 5): They appear as early as 4 months of age. They seem to operate automatically and unconsciously. They are abstract, in that motion cues cause infants to make inferences about later division and accumulation. They go beyond information in the immediate perceptual array.

Of course, many of the criteria for core domains remain open to investigation in the case of substances, but it is worth considering the implications of a possible substance domain for the general core-knowledge framework. The experiments described earlier show that infants understand some of the relations that hold between objects and substances. They know that solid objects can pass through nonsolid substances but not through other solid objects. What's unclear is how infants can know about these relations if objects and substances belong to distinct domains. Encapsulation of domains is a hallmark of core knowledge, as Spelke (2022) emphasizes. So, if objects and substances belong to different core domains (that by definition can't talk to each other), then to capture the relations, the object domain must contain information about substances (e.g., that objects can pass through substances), or the substance domain must contain information about objects (ditto) or both. We're unsure how this overlap could be consistent with encapsulation. There are similar overlaps between object and number domains, and future research could work out these details.

Similarly, all objects consist of substances, and the causal properties of the objects are inherited in part from those substances. A toy's behavior during a collision will differ for a toy made of rubber than for an otherwise similar toy made of metal. Assuming that infants are sensitive to such differences in behavior, then if it's the object domain that predicts the nature of collisions (Spelke, 2022, sects. 2.3 and 2.5), it must have access to knowledge about the substances that compose the objects. But if objects and substances belong to separate encapsulated core domains, this kind of coordination is difficult to explain. Perhaps the correct domains are not objects and substances, but solid objects and nonsolid substances. However, that would leave wood, metal, and other solid substances and socks, raindrops, and other nonsolid objects in limbo until the knowledge is elaborated or refined later in development.

Issues like these reveal a tension between the core-knowledge program's expanding scope and its architectural restrictions. The experimental program that Spelke launched has successfully revealed an increasing range of knowledge that infants deploy. But this increase comes with problems of accommodating the information within core knowledge's separate fiefdoms.

**Acknowledgement.** The authors thank Grace Coram, Kexin Que, and Zixin Zeng for comments on an earlier version of this commentary.



**Financial support.** This research received no specific grant from any funding agency, commercial, or not-for-profit sectors.

**Competing interest.** None.

## References

- Anderson, E. M., Hespos, S. J., & Rips, L. J. (2018). Five-month-old infants have expectations for the accumulation of nonsolid substances. *Cognition*, *175*, 1–10.
- Baillargeon, R., Spelke, E. S., & Wasserman, S. (1985). Object permanence in five-month-old infants. *Cognition*, *20*, 191–208.
- Bourgeois, K. S., Khawar, A. W., Neal, S. A., & Lockman, J. J. (2005). Infant manual exploration of objects, surfaces, and their interrelations. *Infancy*, *8*, 233–252. <https://doi.org/10.1207/s15327078in0803>
- Cherries, E. W., Mitroff, S. R., Wynn, K., & Scholl, B. J. (2008). Cohesion as a constraint on object persistence in infancy. *Developmental Science*, *11*, 427–432.
- Chiandetti, C., & Vallortigara, G. (2011). Intuitive physical reasoning about occluded objects by inexperienced chicks. *Proceedings of the Royal Society B: Biological Sciences*, *278*, 2621–2627.
- Chiang, W.-C., & Wynn, K. (2000). Infants' tracking of objects and collections. *Cognition*, *77*, 169–195.
- Frege, G. (1980). *The foundations of arithmetic: A logico-mathematical enquiry into the concept of number* (J. L. Austin, Trans., 2nd ed.). Northwestern University Press. (Original work published 1884).
- Hespos, S. J., Ferry, A. L., Anderson, E. M., Hollenbeck, E. N., & Rips, L. J. (2016). Five-month-old infants have general knowledge of how nonsolid substances behave and interact. *Psychological Science*, *27*, 244–256.
- Hespos, S. J., Ferry, A. L., & Rips, L. J. (2009). Five-month-old infants have different expectations for solids and liquids. *Psychological Science*, *20*, 603–611.
- Hespos, S. J., & VanMarle, K. (2012). Physics for infants: Characterizing the origins of knowledge about objects, substances, and number. *Wiley Interdisciplinary Reviews: Cognitive Science*, *3*, 19–27.
- Huntley-Fenner, G., Carey, S., & Solimando, A. (2002). Objects are individuals but stuff doesn't count: Perceived rigidity and cohesiveness influence infants' representations of small groups of discrete entities. *Cognition*, *85*, 203–221.
- Kellman, P. J., & Spelke, E. S. (1983). Perceptions of partly occluded objects in infancy. *Cognitive Psychology*, *15*(4), 483–524.
- Koslicki, K. (1997). Isolation and non-arbitrary division: Frege's two criteria for counting. *Synthese*, *112*, 403–430.
- Rips, L. J., & Hespos, S. J. (2015). Divisions of the physical world: Concepts of objects and substances. *Psychological Bulletin*, *141*, 786.
- Rips, L. J., & Hespos, S. J. (2019). Concepts of objects and substances in language. *Psychonomic Bulletin & Review*, *26*, 1238–1256.
- Rosenberg, R. D., & Carey, S. (2009). Infants' representations of material entities. In B. M. Hood & L. R. Santos (Eds.), *The origins of object knowledge* (pp. 165–188). Oxford University Press.
- Rothstein, S. (2017). *Semantics for counting and measuring*. Cambridge University Press.
- Samuels, R. (2012). Massive modularity. In E. Margolis, R. Samuels, & S. Stich (Eds.), *Oxford handbook of philosophy of cognitive science* (pp. 60–91). Oxford University Press.
- Santos, L. R., & Hauser, M. D. (2002). A non-human primate's understanding of solidity: Dissociations between seeing and acting. *Developmental Science*, *5*, F1–F7.
- Spelke, E. S. (2022). *What babies know: Core knowledge and composition* (Vol. 1). Oxford University Press.
- Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental Science*, *10*, 89–96.
- Wellwood, A. (2019). *The meaning of "more"*. Oxford University Press.

## Learning in the social being system

Zoe Jenkin<sup>a\*</sup> and Lori Markson<sup>b</sup>

<sup>a</sup>Department of Philosophy and Center for Cognitive Science, Rutgers University, New Brunswick, NJ, USA and <sup>b</sup>Department of Psychological and Brain Sciences, Washington University in St. Louis, St Louis, MO, USA  
[zoe.jenkin@rutgers.edu](mailto:zoe.jenkin@rutgers.edu)  
[markson@wustl.edu](mailto:markson@wustl.edu)

\*Corresponding author.

doi:10.1017/S0140525X23003138, e132

## Abstract

We argue that the core social being system is unlike other core systems in that it participates in frequent, widespread learning. As a result, the social being system is less constant throughout the lifespan and less informationally encapsulated than other core systems. This learning supports the development of the precursors of bias, but also provides avenues for preempting it.

Among Spelke's important and novel contributions in *What Babies Know* is the idea of a core cognitive system for representing social beings. We share Spelke's view that such a system exists, and that it functions to produce abstract conceptual representations of social beings, social engagement, shareable experience, and social value. However, we argue that the social being system differs from other core cognitive systems with respect to several of the characteristic properties of core knowledge. The social being system is (1) more affected by learning, (2) less constant throughout the lifespan, and (3) less informationally encapsulated than other core systems, especially with respect to its representations of social value. Unlike other core systems, the social being system's computations are environmentally dependent, allowing the system to amplify and morph in ways other core systems do not.

Spelke presents an array of powerful data in favor of the core social being system. In our view, whereas this data on the whole supports the existence of such a system, it also supports certain interesting divergences from the paradigm of core cognition. For example, Spelke reviews evidence that infants are especially sensitive to perceptual cues that carry important social information, such as faces (e.g., Field, Cohen, Garcia, & Greenberg, 1984; Meltzoff & Moore, 1994) and speech (e.g., Werker, 1989). However, as Spelke notes, the developmental patterns of infants' face and speech processing between birth and 12 months are highly dependent on input from the infant's social environment. Newborn infants are equally good at recognizing and differentiating faces of all races, and even discriminate between human and chimpanzee faces. Sometime around 6 months of age infants begin to display the "Other Race Effect," a diminished ability to recognize and differentiate human faces of other races (Kelly *et al.*, 2005, 2007; Pascalis, Scott, & Nelson, 2005; Sangrigoli & De Schonen, 2004). However, perceptual training (i.e., exposure to faces of other races) can diminish or even eliminate the Other Race Effect in 6- to 9-month-old infants (Heron-Delaney *et al.*, 2011). Similar patterns appear in infants' differentiation of non-human primate faces (the "Other Species Effect," Pascalis *et al.*, 2005) and non-native speech sounds (Kuhl, Tsao, & Liu, 2003; Werker, 1989): Between birth and 12 months, wide initial sensitivity to a type of perceptual social cue (faces or speech sounds) either narrows or remains wide, depending on the breadth and diversity of the infants' experiences (see Jenkin 2023a, 2023b for further discussion of infant perceptual learning).

These developmental patterns indicate that whereas the social being system emerges early, it is remarkably malleable. Compared to the core systems for object, place, number, and agent, which are relatively insensitive to environmental factors, the social being system learns from the input it receives. For example, in the case of face processing, during the first year of life the system learns which kinds of faces are regularly present and thus worth recognizing and differentiating. Similarly, in the case of speech

processing, the system learns which speech sounds, or phonemes, are meaningful in one's language. These capacities reflect the faces and speech sounds with which an infant is familiar. Whereas such capacity limits may simply reflect what is familiar, not bias (Wang, Laming, & Andrews, 2022), they may be precursors to learned in-group/out-group divisions that have the potential to lead to the development of bias (Hughes *et al.*, 2019; Vingilis-Jaremko, Kawakami, & Friesen, 2020).

The social being system's susceptibility to learning implies that it is not constant throughout the lifespan. Whereas the system is present from infancy to adulthood, as is characteristic of core cognitive systems, it operates differently at different periods, depending on when and how it has learned. This is especially evident with respect to social preferences and values. For example, Singh, Phneah, Wijayarathne, Lee, and Quinn (2022) studied infants living in Singapore, which is a multiracial society (with a predominantly Chinese population, but with also large Indian and Malay populations), who are raised by caregivers of other races. At 3, 6, and 9 months these infants showed an increasing visual preference for faces of their caregiver's race. This other-race preference was predicted by the extent of contact with members of the other race. This data indicates that experience shapes the social being system over time, such that an infant's social being system at 3 months may look very different from the same infant's social being system at 9 months or older. Another study in Singapore found that extensive experience with caregivers of other races had mitigated the development of preschool age children's explicit racial bias (Setoh, Sudo, Quinn, & Lee, 2023).

Such environmentally driven effects on the social being system continue into childhood. For example, 3- to 7-year-old children's neighborhood and school demographics can affect their racial preferences, as can the race of white children's preschool teachers (Hwang & Markson, 2023). This data thus reflects both micro- and macro-level influences on preschoolers' social preferences. Imitation behaviors are also socially influenced from infancy to early childhood. For example, 14-month-old infants are more likely to imitate a native- over foreign-language speaker, suggesting differential learning and affiliation based on social elements of the input (Buttelmann, Zmyj, Daum, & Carpenter, 2013). Four- and 5-year-old children overimitate actions only of adults who demonstrate social affiliations (Nielsen & Blank, 2011), and 4- and 6-year-old children overimitate actions that are described as normative (Clay, Over, & Tennie, 2018). As children begin to develop richer cognitive concepts of race, gender, class, status, and norms, it becomes more difficult to disentangle effects on the core social being system from effects on other forms of social cognition. Nonetheless, these are examples of effects on representations of social value and imitation behavior, which are among the outputs of Spelke's core social being system.

These examples also indicate that the social being system may be less informationally encapsulated than other core systems. As Spelke argues, the core object, place, number, form, and agent systems are informationally encapsulated, both from central cognition and from each other (Spelke, 2022, p. 194). However, the examples discussed above indicate that the core social being system can be influenced by children's beliefs about social norms and affiliations, as well as by environmental factors such as neighborhood and school demographics. Such influences may be automatic and unconscious, but they do reflect that the social being system is sensitive to a large variety of inputs, and that its boundaries are more porous.

Other core cognitive systems also participate in learning, but they typically do so in different or more minimal ways. For

example, the core number system facilitates young children's learning of natural number concepts (Spelke, 2022, pp. 171–185), but this learning does not alter the core number system itself. Rather, it is an example of a core system supporting the construction of a separate cognitive system. Minimal learning does occur within the core object system, when between 5- and 7-months infants learn that objects will fall when unsupported (Kim & Spelke, 1992, see also Jenkin, 2020 for discussion). But this is a relatively minor addition to the constraints of the system. In contrast, the effects of learning on the social being system are greater and more pervasive. It can be influenced at various developmental periods, and by various factors, such as the members of our family, the people who live in our neighborhoods and attend our schools, children's friends and their parents' social networks (Eason, Kaiser, & Somerville, 2018; Hwang & Markson, 2023; Markson & Luo, 2020; Roberts, Williams, & Gelman, 2017). As children's social interactions and circles expand, or their neighborhoods and schools become more diverse, so might their racial attitudes, biases, and preferences for members of different social groups. This extensive susceptibility to learning sets the social being system apart from other core cognitive systems.

In conclusion, we share Spelke's view that the early emergence of sensitivity to faces and speech sounds support positing a core social being system that is at least in part innate. We depart from Spelke in emphasizing that the system also allows for differential learning trajectories and takes in a wider range of inputs than other core systems do. The malleability of the social being system enables the development of the precursors of bias, but it also enables the prevention of the development of bias. Because the system is relatively plastic, environmental interventions are effective in ways they are not for other core systems.

The environmental dependency of the social being system also raises questions about the relationship between the characteristic properties of core knowledge, and more broadly about what it means for a cognitive system to be "core." Spelke outlines 12 characteristic properties of core knowledge, and argues that if a system has some of them, it is likely to have them all (Spelke, 2022, pp. 190–200). However, if the above arguments are correct, the core social being system has some, but not all, of these characteristic properties (or at least varies in the way these properties manifest). In particular, it is neither developmentally invariant nor entirely encapsulated, and whereas it is innate, it can also be changed by learning. What then should we make of developmental invariance, encapsulation, and innateness as characteristics of core cognition? One option is to say that only some of the characteristic properties of core cognition are necessary to classify a system as part of core knowledge, whereas others are common among core systems but irrelevant to their classification. A second option is to say that the characteristic properties of core cognition function like a cluster concept, such that no individual property from the set is necessary for a system to count as core knowledge, but the presence of a sufficient number of the properties is jointly necessary. In either case, the resulting picture of core cognition is one on which some systems, such as the social being system, are less paradigmatically "core" than others.

**Acknowledgements.** We thank the students at Washington University in St. Louis who participated in each of our seminars on *What Babies Know* for their valuable discussion.


**Financial support.** We have no funding to disclose.

**Competing interest.** None.

## References

- Büttelmann, D., Zmyj, N., Daum, M. M., & Carpenter, M. (2013). Selective imitation of in-group over out-group members in 14-month-old infants. *Child Development, 84*, 422–428. <https://doi.org/10.1111/j.1467-8624.2012.01860.x>
- Clay, Z., Over, H., & Tennie, C. (2018). What drives young children to over-imitate? Investigating the effects of age, context, action type, and transitivity. *Journal of Experimental Child Psychology, 166*, 520–534. <https://doi.org/10.1016/j.jecp.2017.09.008>
- Eason, A. E., Kaiser, C. R., & Somerville, J. A. (2018). Underrepresentation and the perception of others' racial attitudes. *Social Psychological and Personality Science, 10*(6), 757–767. <https://doi.org/10.1177/1948550618788855>
- Field, T. M., Cohen, D., Garcia, R., & Greenberg, R. (1984). Mother-stranger face discrimination by the newborn. *Infant Behavior and Development, 7*(1), 19–25. [https://doi.org/10.1016/S0163-6383\(84\)80019-3](https://doi.org/10.1016/S0163-6383(84)80019-3)
- Heron-Delaney, M., Anzures, G., Herbert, J. S., Quinn, P. C., Slater, A. M., Tanaka J. W., ... Pascalis, O. (2011). Perceptual training prevents the emergence of the other race effect during infancy. *PLoS ONE, 6*(5), e19858. <https://doi.org/10.1371/journal.pone.0019858>
- Hughes, B. L., Camp, N. P., Gomez, J., Natu, V. S., Grill-Spector, K., & Eberhardt, J.L. (2019). Neural adaptation to faces reveals racial outgroup homogeneity effects in early perception. *Proceedings of the National Academy of Sciences, USA, 116*(29), 14532–14537. <https://doi.org/10.1073/pnas.1822084116>
- Hwang, H. G., & Markson, L. (2023). Black and white children's race-based information-endorsement and teacher preference: Effects of school and neighborhood demographics. *Developmental Psychology, 59*(5), 893–907. <https://doi.org/10.1037/dev0001507>
- Jenkin, Z. (2020). The epistemic role of core cognition. *The Philosophical Review, 129*(2), 251–298. <https://doi.org/10.1215/00318108-8012850>
- Jenkin, Z. (2023a). Perceptual learning. *Philosophy Compass, 18*(6), e12932. <https://doi.org/10.1111/phc3.12932>
- Jenkin, Z. (2023b). The function of perceptual learning. *Philosophical Perspectives, 31*(1), 172–186. <https://doi.org/10.1111/phpe.12186>
- Kelly, D. J., Liu, S., Ge, L., Quinn, P. C., Slater, A. M., Lee, K., ... Pascalis, O. (2007). Cross-race preferences for same-race faces extend beyond the African versus Caucasian contrast in 3-month-old infants. *Infancy, 11*, 87–95. <https://doi.org/10.1080/15250000709336871>
- Kelly, D. J., Quinn, P. C., Slater, A. M., Lee, K., Gibson, A., Smith, M., ... Pascalis, O. (2005). Three-month-olds, but not newborns, prefer own-race faces. *Developmental Science, 8*, F31–F36. <https://doi.org/10.1111/j.1467-7687.2005.04344.x>
- Kim, I. K., & Spelke, E. S. (1992). Infants' sensitivity to effects of gravity on visible object motion. *Journal of Experimental Psychology: Human Perception and Performance, 18*(2), 385–393. <https://doi.org/10.1037/0096-1523.18.2.385>
- Kuhl, P. K., Tsao, F. M., & Liu, H. M. (2003). Foreign-language experience in infancy: Effects of short-term exposure and social interaction on phonetic learning. *Proceedings of the National Academy of Sciences, USA, 100*(15), 9096–9101. <https://doi.org/10.1073/pnas.1532872100>
- Markson, L., & Luo, Y. (2020). Trust in early childhood. *Advances in Child Development and Behavior, 58*, 137–162. <https://doi.org/10.1016/bs.acdb.2020.01.005>
- Meltzoff, A. N., & Moore, M. K. (1994). Imitation, memory, and the representation of persons. *Infant Behavior and Development, 25*(1), 39–61. [https://doi.org/10.1016/0163-6383\(94\)90024-8](https://doi.org/10.1016/0163-6383(94)90024-8)
- Nielsen, M., & Blank, C. (2011). Imitation in young children: Who gets copied is more important than what gets copied. *Developmental Psychology, 47*(4), 1050–1053. <https://doi.org/10.1037/a0023866>
- Pascalis, O., Scott, L. S., & Nelson, C. A. (2005). The plasticity of face processing in infancy. *Biological Sciences, 102*(14), 5297–5300. <https://doi.org/10.1073/pnas.0406627102>
- Roberts, S. O., Williams, A. D., & Gelman, S. A. (2017). Children's and adults' predictions of black, white, and multiracial friendship patterns. *Journal of Cognition and Development, 18*(2), 189–208. <https://doi.org/10.1080/15248372.2016.1262374>
- Sangrigoli, S., & De Schonen, S. (2004). Recognition of own-race and other-race faces by three-month-old infants. *Journal of Child Psychology and Psychiatry, 45*(7), 1219–1227. <https://doi.org/10.1111/j.1469-7610.2004.00319.x>. PMID: 15335342.
- Setoh, P., Sudo, M., Quinn, P. C., & Lee, K. (2023). Does extended experience with other-race nannies predict racial bias in the preschool years? *Journal of Experimental Child Psychology, 235*, 105729. <https://doi.org/10.1016/j.jecp.2023.105729>
- Singh, L., Phneah, K. T., Wijayaratne, D. C., Lee, K., & Quinn, P. C. (2022). Effects of interracial experience on the race preferences of infants. *Journal of Experimental Child Psychology, 216*, 105352. <https://doi.org/10.1016/j.jecp.2021.105352>
- Spelke, E. (2022). *What babies know: Core knowledge and composition*, (Vol. 1). Oxford University Press.
- Vingilis-Jaremko, L., Kawakami, K., & Friesen, J. P. (2020). Other-groups bias effects: Recognizing majority and minority outgroup faces. *Social Psychological and Personality Science, 11*(7), 908–916. <https://doi.org/10.1177/1948550620919562>
- Wang, A. O., Laming, C., & Andrews, T. J. (2022). Covariation in the recognition of own-race and other-race faces argues against the role of group bias in the other race effect. *Scientific Reports, 12*, 13088. <https://doi.org/10.1038/s41598-022-17330-9>
- Werker, J. F. (1989). Becoming a native listener. *American Scientist, 77*(1), 54–59.

## Is core knowledge a natural subdivision of infant cognition?

Caroline M. Kaicher\* , Julia J. Conti,  
Abhishek M. Dedhe, Lauren S. Aulet and  
Jessica F. Cantlon

Department of Psychology, Carnegie Mellon University, Pittsburgh, PA, USA  
ckaicher@andrew.cmu.edu; jjconti@andrew.cmu.edu; adedhe@andrew.cmu.edu; laulet@andrew.cmu.edu; jcantlon@andrew.cmu.edu  
<https://www.cmu.edu/dietrich/psychology/kidneurolab/>

\*Corresponding author.

doi:10.1017/S0140525X23003229, e133

### Abstract

We examine Spelke's core knowledge taxonomy and test its boundaries. We ask whether Spelke's core knowledge is a distinct *type* of cognition in the sense that the cognitive processes it includes and excludes are biologically and mechanically coherent.

Spelke's central thesis classifies infant knowledge into the following distinct core knowledge systems: Objects, places, number, forms, agents, and social beings. These systems apply to specific domains of entities in the world and capture specific properties of those entities. They share many features – perhaps the most critical being that they are ancient, automatic, center on abstract concepts, and emerge early in life. Spelke characterizes these systems as innate, invariant over development, impervious to explicit beliefs, and dependent on attention, with a primary function of supporting learning and operating (in part) through mental stimulation. Thus, core knowledge systems are defined not only by their functional role in human cognition but also by characteristics of the mechanisms supporting them. Spelke's core knowledge taxonomy provides a framework for understanding the evolutionary and developmental origins of human knowledge, including the foundations of complex cognition.

However, we question the ability of Spelke's core knowledge taxonomy to “carve nature at its joints.” Testing the boundaries of Spelke's core knowledge framework is important because it helps refine the theory, enabling more precise predictions about the emergence and progression of infant cognition. Here, we ask whether Spelke's collection of core knowledge domains represent a meaningfully distinct *type* of cognition. Infants have several other cognitive functions that share characteristics with the core domains in that they are ancient, automatic, early-emerging, and abstract cognitive processes that are integral to infants' information processing, and equally essential for explaining what they know. We describe three examples – categorical perception, referential understanding, and algebraic rule learning – to demonstrate this point, raising the question: Is Spelke's “core knowledge” a natural subdivision of infant cognition?

Spelke uses the common characteristics shared by object, place, number, form, agent, and social systems to argue that core knowledge is a distinct cognitive *type*. Specifically, she claims that all the shared characteristics of the core domains “go together,” and that, “any ancient, abstract conceptual system that has some of these



properties is likely to have all of them” (Spelke, 2022, p. 198). Critically, Spelke states that core knowledge systems focus on “the problem of understanding what the sensed world consists of: what entities inhabit it, how those entities behave, and why they do what they do” (Spelke, 2022, p. 36). But, given the criteria, which mechanisms of infant cognition are *not* core knowledge – and why not?

Here we argue that there is no sharp boundary between Spelke’s core knowledge and the rest of infant cognition by showing how three fundamental cognitive processes – categorical perception, referential understanding, and algebraic rule learning – are not only automatic, unconscious, ancient, and abstract, but also support knowledge in infants and are essential to their understanding of the perceptual world.

**Categorical perception:** Infants display categorical perception, or the propensity to assign discrete boundaries among stimuli varying along a continuum. This process is demonstrated when perceptual discriminations are easier for items belonging to different categories, and more difficult for items belonging to the same category, even when their physical differences are objectively equal (Goldstone & Hendrickson, 2010). The most prominent example of this is speech perception, during which we perceive phonemes that are abstracted from the pure acoustic properties of the signal. This process is automatic (Kasai et al., 2003), and is demonstrated in infants as young as 1 month old (Eimas, Siqueland, Jusczyk, & Vigorito, 1971). This ability is not uniquely human – macaques exhibit the same phoneme boundary effect as 1-month-old infants (Kuhl & Padden, 1982), European starlings can learn vowel sound categories (Kluender, Lotto, Holt, & Bloedel, 1998), and chinchillas detect changes along phoneme boundaries in particular (Kuhl, 1981). Additionally, the neural basis of this ability is shared between infants and adults (Dehaene-Lambertz & Gliga, 2004), and also among humans and non-human primates (Ley et al., 2012). Beyond speech, 2-month-old infants categorically perceive some non-speech sounds (Jusczyk, Rosner, Cutting, Foard, & Smith, 1977), 7-month-old infants display categorical perception of facial expressions of emotion (Kotsoni, de Haan, & Johnson, 2001), and 4-month-olds categorically perceive color (Franklin et al., 2008). Categorical perception of color has also been shown in goldfish (Goldman, Lanson, & Brown, 1990) and zebra finches (Zipple et al., 2019).

Categorical perception makes items that are meaningfully different *more* distinct, and makes those that are meaningfully similar *more* similar. This helps infants eliminate unnecessary information and allows them to more efficiently represent the stimuli around them (Oakes & Madole, 2003). This capacity provides the “building blocks” for higher-order categories (Harnad, 1987), which is not only critical for learning language (Werker & Lalonde, 1988), but it may also be a basis for social categorization in infants (Lieberman, Woodward, & Kinzler, 2017). Thus, the propensity to create discrete category representations is a core aspect of infant cognition that is abstract, ancient, early-emerging, automatic, and supports learning.

**Referential understanding:** Referential understanding refers to the ability to understand that communicative signals such as words and pointing are linked to something concrete in the world, and to use such signals to imply intended referents (Wynne & Udell, 2013). Infants as young as 3 months old demonstrate this by using words to help them categorize objects (Ferry, Hespos, & Waxman, 2010). At 1 year old, infants understand the referential nature of deictic gestures (Gliga & Csibra,

2009) and begin to utilize pointing (Tomasello, Carpenter, & Liszkowski, 2007). Referential understanding is also automatic, as demonstrated every time we use language. The ability to match symbols or gestures to referents is also present in dogs (Kaminski, Call, & Fischer, 2004), dolphins (Herman, Richards, & Wolz, 1984), and apes (Savage-Rumbaugh, Shanker, & Taylor, 1998). Spelke argues that our propensity to use symbols is rooted in human-specific language abilities, but since this capacity is shared between species, it may be more primitive. In fact, linking labels to referents can be considered an associative process during which children use space/object and space/word associations to link words to objects (Samuelson, Smith, Perry, & Spencer, 2011). Associative processes such as this are abstract (Delamater, Desouza, Rivkin, & Derman, 2014) and are present in a variety of non-human animals (Rescorla & Holland, 1982).

Referential understanding is key for word learning in infants (Gentner & Boroditsky, 2001), and in their second year they begin to utilize the non-arbitrary referential actions of others (i.e., looking and pointing) to establish arbitrary referential relationships, such as mapping words onto objects (Baldwin, 1993). Additionally, 12-month-old infants rely on referential cues to connect others’ emotional messages with novel objects (Moses, Baldwin, Rosicky, & Tidball, 2001). Thus, referential understanding is not only abstract, early-emerging, ancient, and automatic, but it also plays an important role in infants’ learning about the world around them.

**Algebraic rule learning:** Algebraic rule learning requires one to detect relations between entities, and is characterized by an ability to generalize patterns to novel items (Dehaene, Meyniel, Wacogne, Wang, & Pallier, 2015). Infants demonstrate this through their remarkable ability to extract rules from visual and auditory input. For example, 3-month-olds can generalize same/different relations among arrays of toys (Anderson, Chang, Hespos, & Gentner, 2018) and 4-month-olds can do so with geometric shapes (Addyman & Mareschal, 2010). Additionally, newborns can discriminate spoken syllable patterns (Gervain, Macagno, Cogoi, Peña, & Mehler, 2008). Our detection of algebraic rules is also an automatic and unconscious process (Dehaene et al., 2015; Miller, 1967). Kanzi the chimpanzee demonstrated the ability to understand word order grammatical rules (Schoenemann, 2022), dolphins display key elements of syntax (Kako, 1999), and crows and monkeys can even generate recursive sequences (Ferrigno, Cheyette, Piantadosi, & Cantlon, 2020; Liao, Brecht, Johnston, & Nieder, 2022). In addition, macaques can learn context-free grammars based on embedded spatial sequences (Ferrigno, 2022; Jiang et al., 2018), demonstrating an evolutionarily conserved propensity for algebraic rule learning.

Infants’ rule learning abilities are essential to the development of complex capacities such as language. For instance, 4-month-olds can detect non-adjacent grammatical dependencies in a novel language after only one learning session (Friederici, Mueller, & Oberecker, 2011), and 17-month-olds can segment words in fluent speech based on non-adjacent dependencies using statistical learning (Frost et al., 2020). Infants’ rule-learning abilities also help them learn the “grammar” of music (McMullen & Saffran, 2004). Thus, abundant evidence demonstrates that algebraic rule learning is an abstract, early-emerging, automatic, unconscious, and ancient aspect of infant knowledge, supporting learning in multiple domains (Rabagliati, Ferguson, & Lew-Williams, 2019).

Perhaps what makes the core domains in Spelke’s theory distinct is that they “operate on a limited domain of entities” and

“capture only a limited subset of properties that our perceptual systems deliver” (Spelke, 2022, p. 190). However, we question whether Spelke’s core domains are more selective, rigid, or filtered than other systems. For instance, knowledge of number can be used with any discrete set of things or events, and adapts to new, evolutionarily recent information such as digits and verbal counting. Numerical information automatically interacts with perceptual and semantic information from disparate domains during development (e.g., Gebuis, Cohen Kadosh, De Haan, & Henik, 2009). Ferrigno, Jara-Ettinger, Piantadosi, and Cantlon (2017) showed that when both numerical and surface area information is available for approximate magnitude discrimination, numerical biases are uniquely enhanced in humans compared to non-human primates. Additionally, they found that within the Tsimane’, a non-industrialized group in Bolivia, adults who have learned to count display a greater number bias than those who have not. Spelke herself even discusses how Mundurucu children and adults who have been exposed to formal education have more precise numerical representations than those who have not (Piazza, Pica, Izard, Spelke, & Dehaene, 2013). Spelke uses this evidence to show that the core number system supports learning of the symbolic number system, but it also shows that the core number system can be penetrated by novel domains and inputs. Thus, the number system may not be as independent, rigid, or limited as it is made out to be. Similarly, the limitations of the “core” systems, such as the numerical system, are not greater than the biases and constraints on other informational systems such as categorical perception, referential understanding, and rule learning. All mechanisms have their own unique cognitive signatures and constraints for abstracting information across diverse entities while adapting to novel inputs and problems.

Mechanisms that are (perhaps erroneously) considered more “general purpose” than the core domains also exhibit biases and constraints on processing. This is even the case for reinforcement learning, in which avoidance responses to different reinforcers (induced nausea or shock) are more readily associated with certain cues (gustatory and audiovisual, respectively) than others in rats (Garcia & Koelling, 1966). This bias is present in humans, as shown through the privileged role of nausea in the acquisition of food dislikes (Pelchat & Rozin, 1982). Thus, deep information processing biases are present in this “general” mechanism and influence learning in humans. Our three purportedly general-purpose mechanisms also display innate biases and are subject to information constraints and filters. For instance, rule learning, like the number system, has capacity limits – just as larger numerical differences are easier to discriminate than smaller ones, shorter range dependencies are easier to learn than longer ones (Futrell, Mahowald, & Gibson, 2015). For referential understanding, children display specific biases, such as the whole-object, taxonomic, and mutual exclusivity assumptions, that constrain how they map words onto referents (Markman, 1991). Additionally, information processing through categorical perception is constrained so that objective similarities between stimuli are filtered based on useful category boundaries (Goldstone & Hendrickson, 2010). Category formation can also be constrained by the number of exemplars, their variability, and their similarity (Needham, Dueker, & Lockhead, 2005).

Thus, categorical perception, referential understanding, and algebraic rule learning are three examples of key components of infant cognition – things that babies “know” and that are integral to their understanding of the world. These processes exhibit innate biases and are subject to information constraints,

and filters similar to Spelke’s core knowledge domains. The range of infant abilities that are early-emerging, abstract, automatic, ancient, and *not* considered core knowledge indicates that infant knowledge emerges independently of the purported specificity of its domain. In this sense, the boundaries of core knowledge set by Spelke are not biologically and mechanically coherent, and are displaced from the evolutionary and developmental origins of infant cognition and the knowledge it generates. The disconnection between well-known evolved cognitive functions and Spelke’s lens limits the explanatory and predictive power of “core knowledge” as a taxonomy – if the boundaries of core knowledge arbitrarily exclude key forms of infant cognition, then the framework cannot anticipate what babies naturally know.

**Financial support.** This work was supported by the National Institute of Child Health and Human Development (JFC, grant number R01HD107715); and National Science Foundation (JFC, grant number 2148343).


**Competing interest.** None.

## References

- Addyman, C., & Mareschal, D. (2010). The perceptual origins of the abstract same/different concept in human infants. *Animal Cognition*, 13, 817–833. <https://doi.org/10.1007/s10071-010-0330-0>
- Anderson, E. M., Chang, Y., Hespos, S., & Gentner, D. (2018). Comparison within pairs promotes analogical abstraction in three-month-olds. *Cognition*, 176, 74–86. <https://doi.org/10.1016/j.cognition.2018.03.008>
- Baldwin, D. A. (1993). Early referential understanding: Infants’ ability to recognize referential acts for what they are. *Developmental Psychology*, 29(5), 832. <https://doi.org/10.1037/0012-1649.29.5.832>
- Dehaene, S., Meyniel, F., Wacongne, C., Wang, L., & Pallier, C. (2015). The neural representation of sequences: From transition probabilities to algebraic patterns and linguistic trees. *Neuron*, 88(1), 2–19. <https://doi.org/10.1016/j.neuron.2015.09.019>
- Dehaene-Lambertz, G., & Gliga, T. (2004). Common neural basis for phoneme processing in infants and adults. *Journal of Cognitive Neuroscience*, 16(8), 1375–1387. <https://doi.org/10.1162/0898929042304714>
- Delamater, A. R., Desouza, A., Rivkin, Y., & Derman, R. (2014). Associative and temporal processes: A dual process approach. *Behavioural Processes*, 101, 38–48. <https://doi.org/10.1016/j.beproc.2013.09.004>
- Eimas, P. D., Siqueland, E. R., Jusczyk, P., & Vigorito, J. (1971). Speech perception in infants. *Science*, 171(3968), 303–306. <https://doi.org/10.1126/science.171.3968.303>
- Ferrigno, S. (2022). Sequencing, artificial grammar, and recursion in primates. In B. L. Shwartz, & M. J. Beran (Eds.), *Primate cognitive studies* (pp. 260–290). Cambridge University Press.
- Ferrigno, S., Cheyette, S. J., Piantadosi, S. T., & Cantlon, J. F. (2020). Recursive sequence generation in monkeys, children, US adults, and native Amazonians. *Science Advances*, 6(26), eaaz1002. <https://doi.org/10.1126/sciadv.aaz1002>
- Ferrigno, S., Jara-Ettinger, J., Piantadosi, S. T., & Cantlon, J. F. (2017). Universal and uniquely human factors in spontaneous number perception. *Nature Communications*, 8(1), 13968. <https://doi.org/10.1038/ncomms13968>
- Ferry, A. L., Hespos, S. J., & Waxman, S. R. (2010). Categorization in 3- and 4-month-old infants: An advantage of words over tones. *Child Development*, 81(2), 472–479. <https://doi.org/10.1111/j.1467-8624.2009.01408.x>
- Franklin, A., Drivonikou, G. V., Bevis, L., Davies, I. R., Kay, P., & Regier, T. (2008). Categorical perception of color is lateralized to the right hemisphere in infants, but to the left hemisphere in adults. *Proceedings of the National Academy of Sciences*, 105(9), 3221–3225. <https://doi.org/10.1073/pnas.0712286105>
- Friederici, A. D., Mueller, J. L., & Oberecker, R. (2011). Precursors to natural grammar learning: Preliminary evidence from 4-month-old infants. *PLoS ONE*, 6(3), e17920. <https://doi.org/10.1371/journal.pone.0017920>
- Frost, R. L., Jessop, A., Durrant, S., Peter, M. S., Bidgood, A., Pine, J. M., ... Monaghan, P. (2020). Non-adjacent dependency learning in infancy, and its link to language development. *Cognitive Psychology*, 120, 101291. <https://doi.org/10.1016/j.cogpsych.2020.101291>
- Futrell, R., Mahowald, K., & Gibson, E. (2015). Large-scale evidence of dependency length minimization in 37 languages. *Proceedings of the National Academy of Sciences*, 112(33), 10336–10341. <https://doi.org/10.1073/pnas.1502134112>
- Garcia, J., & Koelling, R. A. (1966). Relation of cue to consequence in avoidance learning. *Psychonomic Science*, 4, 123–124. <https://doi.org/10.3758/BF03342209>

- Gebuis, T., Cohen Kadosh, R., De Haan, E., & Henik, A. (2009). Automatic quantity processing in 5-year olds and adults. *Cognitive Processing*, 10, 133–142. <https://doi.org/10.1007/s10339-008-0219-x>
- Gentner, D., & Boroditsky, L. (2001). Individuation, relativity, and early word learning. *Language Acquisition and Conceptual Development*, 3, 215–256.
- Gervain, J., Macagno, F., Cogo, S., Peña, M., & Mehler, J. (2008). The neonate brain detects speech structure. *Proceedings of the National Academy of Sciences*, 105(37), 14222–14227. <https://doi.org/10.1073/pnas.0806530105>
- Gluga, T., & Csibra, G. (2009). One-year-old infants appreciate the referential nature of deictic gestures and words. *Psychological Science*, 20(3), 347–353. <https://doi.org/10.1111/j.1467-9280.2009.02295.x>
- Goldman, M., Lanson, R., & Brown, G. G. (1990). Wavelength categorization by goldfish (*Carassius auratus*). *International Journal of Comparative Psychology*, 4(3). <https://doi.org/10.46867/C4JS4X>
- Goldstone, R. L., & Hendrickson, A. T. (2010). Categorical perception. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(1), 69–78. <https://doi.org/10.1002/wcs.26>
- Harnad, S. (1987). Psychophysical and cognitive aspects of categorical perception: A critical overview. In S. Harnad (Ed.), *Categorical perception: The groundwork of cognition* (pp. 1–25). Cambridge University Press.
- Herman, L. M., Richards, D. G., & Wolz, J. P. (1984). Comprehension of sentences by bottlenosed dolphins. *Cognition*, 16(2), 129–219. [https://doi.org/10.1016/0010-0277\(84\)90003-9](https://doi.org/10.1016/0010-0277(84)90003-9)
- Jiang, X., Long, T., Cao, W., Li, J., Dehaene, S., & Wang, L. (2018). Production of supra-regular spatial sequences by macaque monkeys. *Current Biology*, 28(12), 1851–1859. <https://doi.org/10.1016/j.cub.2018.04.047>
- Jusczyk, P. W., Rosner, B. S., Cutting, J. E., Foard, C. F., & Smith, L. B. (1977). Categorical perception of nonspeech sounds by 2-month-old infants. *Perception & Psychophysics*, 21, 50–54. <https://doi.org/10.3758/BF03199467>
- Kako, E. (1999). Elements of syntax in the systems of three language-trained animals. *Animal Learning & Behavior*, 27, 1–14. <https://doi.org/10.3758/BF03199424>
- Kaminski, J., Call, J., & Fischer, J. (2004). Word learning in a domestic dog: Evidence for “fast mapping”. *Science*, 304(5677), 1682–1683. <https://doi.org/10.1126/science.1097859>
- Kasai, K., Yamada, H., Kamio, S., Nakagome, K., Iwanami, A., Fukuda, M., ... Kato, N. (2003). Neuromagnetic correlates of impaired automatic categorical perception of speech sounds in schizophrenia. *Schizophrenia Research*, 59(2–3), 159–172. [https://doi.org/10.1016/S0920-9964\(01\)00382-6](https://doi.org/10.1016/S0920-9964(01)00382-6)
- Kluender, K. R., Lotto, A. J., Holt, L. L., & Bloedel, S. L. (1998). Role of experience for language-specific functional mappings of vowel sounds. *The Journal of the Acoustical Society of America*, 104(6), 3568–3582. <https://doi.org/10.1121/1.423939>
- Kotsoni, E., de Haan, M., & Johnson, M. H. (2001). Categorical perception of facial expressions by 7-month-old infants. *Perception*, 30(9), 1115–1125. <https://doi.org/10.1068/p3155>
- Kuhl, P. K. (1981). Discrimination of speech by nonhuman animals: Basic auditory sensitivities conducive to the perception of speech-sound categories. *The Journal of the Acoustical Society of America*, 70(2), 340–349. <https://doi.org/10.1121/1.386782>
- Kuhl, P. K., & Padden, D. M. (1982). Enhanced discriminability at the phonetic boundaries for the voicing feature in macaques. *Perception & Psychophysics*, 32(6), 542–550. <https://doi.org/10.3758/BF03204208>
- Ley, A., Vroomen, J., Hausfeld, L., Valente, G., De Weerd, P., & Formisano, E. (2012). Learning of new sound categories shapes neural response patterns in human auditory cortex. *Journal of Neuroscience*, 32(38), 13273–13280. <https://doi.org/10.1523/JNEUROSCI.0584-12.2012>
- Liao, D. A., Brecht, K. F., Johnston, M., & Nieder, A. (2022). Recursive sequence generation in crows. *Science Advances*, 8(44), eabq3356. <https://doi.org/10.1126/sciadv.abq3356>
- Liberman, Z., Woodward, A. L., & Kinzler, K. D. (2017). The origins of social categorization. *Trends in Cognitive Sciences*, 21(7), 556–568. <https://doi.org/10.1016/j.tics.2017.04.004>
- Markman, E. M. (1991). The whole-object, taxonomic, and mutual exclusivity assumptions as initial constraints on word meanings. In S. A. Gelman, & J. P. Byrnes (Eds.), *Perspectives on language and thought: Interrelations in development* (pp. 72–106). Cambridge University Press. <https://doi.org/10.1017/CBO9780511983689.004>
- McMullen, E., & Saffran, J. R. (2004). Music and language: A developmental comparison. *Music Perception*, 21(3), 289–311. <https://doi.org/10.1525/mp.2004.21.3.289>
- Miller, G. A. (1967). The psychology of communication. *Human Resource Management*, 6(3), 43.
- Moses, L. J., Baldwin, D. A., Rosicky, J. G., & Tidball, G. (2001). Evidence for referential understanding in the emotions domain at twelve and eighteen months. *Child Development*, 72(3), 718–735. <https://doi.org/10.1111/1467-8624.00311>
- Needham, A., Dueker, G., & Lockhead, G. (2005). Infants’ formation and use of categories to segregate objects. *Cognition*, 94(3), 215–240. <https://doi.org/10.1016/j.cognition.2004.02.002>
- Oakes, L. M., & Madole, K. L. (2003). Principles of developmental change in infants’ category formation. In D. H. Rakison, & L. M. Oakes (Eds.), *Early category and concept development: Making sense of the blooming, buzzing confusion* (pp. 132–158). Oxford University Press.
- Pelchat, M. L., & Rozin, P. (1982). The special role of nausea in the acquisition of food dislikes by humans. *Appetite*, 3(4), 341–351. [https://doi.org/10.1016/S0195-6663\(82\)80052-4](https://doi.org/10.1016/S0195-6663(82)80052-4)
- Piazza, M., Pica, P., Izard, V., Spelke, E. S., & Dehaene, S. (2013). Education enhances the acuity of the nonverbal approximate number system. *Psychological Science*, 24(6), 1037–1043. <https://doi.org/10.1177/0956797612464057>
- Rabagliati, H., Ferguson, B., & Lew-Williams, C. (2019). The profile of abstract rule learning in infancy: Meta-analytic and experimental evidence. *Developmental Science*, 22(1), e12704. <https://doi.org/10.1111/desc.12704>
- Rescorla, R. A., & Holland, P. C. (1982). Behavioral studies of associative learning in animals. *Annual Review of Psychology*, 33(1), 265–308.
- Samuelson, L. K., Smith, L. B., Perry, L. K., & Spencer, J. P. (2011). Grounding word learning in space. *PLoS ONE*, 6(12), e28095. <https://doi.org/10.1371/journal.pone.0028095>
- Savage-Rumbaugh, E. S., Shanker, S., & Taylor, T. J. (1998). *Apes, language, and the human mind*. Oxford University Press.
- Schoenemann, P. T. (2022). Evidence of grammatical knowledge in apes: An analysis of Kanzi’s performance on reversible sentences. *Frontiers in Psychology*, 13, 885605. <https://doi.org/10.3389/fpsyg.2022.885605>
- Spelke, E. S. (2022). *What babies know: Core knowledge and composition volume 1* (Vol. 1). Oxford University Press.
- Tomasello, M., Carpenter, M., & Liszkowski, U. (2007). A new look at infant pointing. *Child Development*, 78(3), 705–722. <https://doi.org/10.1111/j.1467-8624.2007.01025.x>
- Werker, J. F., & Lalonde, C. E. (1988). Cross-language speech perception: Initial capabilities and developmental change. *Developmental Psychology*, 24(5), 672. <https://doi.org/10.1037/0012-1649.24.5.672>
- Wynne, C. D., & Udell, M. A. (2013). *Animal cognition: Evolution, behavior and cognition*. Bloomsbury Publishing.
- Zippel, M. N., Caves, E. M., Green, P. A., Peters, S., Johnsen, S., & Nowicki, S. (2019). Categorical colour perception occurs in both signalling and non-signalling colour ranges in a songbird. *Proceedings of the Royal Society B: Biological Sciences*, 286(1903), 20190524. <https://doi.org/10.1098/rspb.2019.0524>

## Wired for society? From ego-logy to eco-logy

Laurence Kaufmann<sup>a</sup> and Fabrice Clément<sup>b</sup> 

<sup>a</sup>Institute of Social Sciences, University of Lausanne, Lausanne, Switzerland and

<sup>b</sup>Cognitive Science Centre, University of Neuchâtel, Neuchâtel, Switzerland

[laurence.kaufmann@unil.ch](mailto:laurence.kaufmann@unil.ch),

<https://applicationspub.unil.ch/interpub/noauth/php/Un/UnPers.php>

[PerNum=871282&LanCode=37](https://www.unine.ch/iscc/home/equipe/fabrice-clement.html)

<https://www.unine.ch/iscc/home/equipe/fabrice-clement.html>

### Corresponding author:

Fabrice Clément;

Email: [fabrice.clement@unine.ch](mailto:fabrice.clement@unine.ch)

doi:10.1017/S0140525X23003126, e134

### Abstract

Somewhat questioning Elizabeth Spelke’s attempt to account for infants’ social knowledge, our commentary argues that social cognition might be divided into several specialized systems. In addition to the core system dedicated to the intersubjective dimension of close relationships, infants could be prewired to process social relationships, such as dominance, characterized by their impersonal, normative dimension.

As exciting as it is erudite, Spelke’s (2022) book provides a remarkable overview of her groundbreaking research into “innate systems of core knowledge.” To her, those domain-specific, innate cognitive capacities are essential for carving out “our cognitive territory into more manageable units.” As shown by Spelke’s successive attempts, identifying core knowledge systems is however a



complicated task. In her book, she argues for six core systems – objects, numbers, places, forms, agents, and social beings. The phylogenetic and ontogenetic functions of those core systems, interposed between percepts and beliefs, are to support the “first informative encounters with the entities they serve to represent” (p. xix). Although Spelke’s hypothesis is quite convincing, we would like to make two critical arguments. Our first argument concerns the ontological divergence between the six core systems and the “entities” they are supposed to deal with. The second argument focuses on the conceptual primitives at the heart of the “core system of social beings.”

From an ontological point of view, numbers and shapes are not individual entities in themselves, but universal properties of given entities. Without reactivating here the old metaphysical debate between the ontology of individuals and the ontology of universals, it is hard to ignore that properties such as quantity, form, or place have a higher degree of abstraction than objects or agents. Whereas first-order concrete entities are defined by their material existence and causal power in the physical world, second-order abstract properties are more complex ways of referring to and qualifying those entities. There is an ontological hierarchy between individual entities endowed with spatiotemporal coordinates and the abstract properties that these same entities instantiate. Herein lies our question. If entities and properties are not on the same ontological page, can they be situated in the same cognitive territory? Or do we need to organize so-called core systems into a multilayer cognitive architecture in which properties have not the same status as objects and agents? This question, both ontological and cognitive, also underlies the point we would like to make about social cognition.

Within Spelke’s framework, newborns, who are literally *infants*, that is, children who are “not able to speak,” have at their disposal two independent, automatic core knowledge systems dedicated to the nonphysical world: the “agent system” and the “social system.” *The agent system* allows infants to represent self-moving entities as *agents* who act in a causally effective and perceptually guided way. Present in human newborns and also in newly hatched chicks, the agent system consists in seeing self-propelled entities as goal-oriented beings that cause changes in other objects only on contact and through their motion. As for *the social system*, it targets *social beings*, that is, people “who endow one another with experiences like their own and who share their experiences in states of engagement.” To Spelke, evidence for a core system dedicated to “shareable experiences of known, individual people” is drawn from newborn’s sensitivity to mutual gazes and affective engagements with their caregivers. As core knowledge is modular, the distinction between the agent and the social systems has strange consequences for the first year of infants’ social life: People around them can be conceived either as entities sharing *phenomenal* experiences with them (social beings) or as *intentional* entities exerting their causal power over objects (agents). It is only at the end of the first year of life that these two central systems can be combined, and infants finally become able to see people as social agents, endowed with phenomenal *and* intentional properties. According to Spelke, this cognitive achievement is because of the progressive mastery of language and the access it provides to the plurality of others’ perspectives.

It is worth mentioning that Spelke used to have a different view of social core knowledge, which served to identify “members of one’s own social group” and “to guide social interactions with

in- and out-group members” (Kinzler & Spelke, 2007, p. 257). In her book, Spelke appears to have modified her perspective. Her definition of the social is primarily, if not solely, intersubjective, as evidenced by her insistence on mimicry, intimacy, proximity, imitation, and emotional sharing. Certainly, the prosocial inclination to build “We-ness” or “togetherness” is essential to individual survival, on both phylogenetic and ontogenetic scales. However, it is questionable whether the social is reducible to “like-ness” in both senses of *similarity* and *kindness*. Indeed, in Spelke’s core social system, the social world is a matter of affiliation or sociability: It is made up of entities “like-me,” then “like-us.” In this ego-centric model, newborns build their social world progressively, as if their knowledge followed a series of concentric circles. First exposed to intersubjective sharing and face-to-face interaction, *ego* bridges the gap between itself and the minds of those around it through mimicry and experience sharing, and then gradually learns to expand its social network.

The problem is that “being one of us” is not the only meaning of the social. Social groups are defined by a normative structure that enables their members to predict how others are likely to behave, given *their position in the social order*. Besides intersubjective, horizontal relationships between acquaintances, society is also based on vertical subordination to impersonal constraints and social hierarchies. In this sense, the mark of the social is not *We-ness* but the *impersonal* relations of interdependence between status bearers or role takers. Social structures are not a matter of voluntarily engagement but a matter of enlistment. As shown by ethology and primatology, others are not only benevolent social partners but also malevolent rivals. Social agents are not only those who voluntarily engage themselves in rewarding close relationships but also those who are enrolled, whether they want it or not, in impersonal scripts, situations, and roles. In short, human and nonhuman social life cannot be reduced to experience sharing.

Our proposal is to separate the components of social cognition that Spelke tends to intertwine, namely (a) the conception of “people as individuals with mental experiences and the simulation of shareable actions and experiences of known, individual people” and (b) “the relationships that connect them.” The first component involves the ability to understand and represent the mental states of others, often referred to as “naïve psychology.” The second component of social cognition focuses on “the kind of people” we are dealing with, and on the way we use group-level identification – be it race, gender, kinship, status, or occupation – as a basis for inference, prediction, and action (Hirschfeld, 2001). This ability, frequently referred to as “naïve” or “intuitive sociology,” might be viewed as a distinct core knowledge and could also be expanded. In fact, group membership, by definition, is a social relationship that influences and defines how group members relate to one another. Viewing group membership as a form of primitive social *relationship*, rather than merely a category-based perception of *individuals*, is theoretically heuristic: It directs attention toward other basic types of relationships, such as exchange, cooperation, competition, and dominance.

Various elements from Spelke’s impressive list of studies argue in favor of a relational conception of social knowledge, a conception that could well be integrated into her account. As mentioned above, she has defended few years ago an account of the core social system based on *in-group* and *out-group* reasoning. Another discussion in *What Babies Know* (Spelke, 2022) could pave the way for a relational view of social cognition, one (too)

quickly sidestepped: dominance. Indeed, 9- and 10-month-old infants expect small agents to bow and prostrate in subordination to others of more formidable physical size (Thomsen, Frankenhuys, Ingold-Smith, & Carey, 2011). Infants as young as 15-month-old demonstrate a strong sensitivity to third-party asymmetric relationships that they expect to remain stable from one conflict (e.g., when the dominant agent repeatedly pushes the subordinate to monopolize a specific area) to another conflict (e.g., when the two agents compete over a desired resource) (Mascaro & Csibra, 2012). When citing these studies, Spelke insists on the fact that these infants are a least 12 months old and are therefore “too old” for their cognitive performances to fall within the scope of core knowledge. Nevertheless, on p. 412, she does mention that 6-month-old infants, when they see a conflict between two individuals from separate social groups, expect the individual from the larger group to prevail (Pun, Birch, & Baron, 2016). In her persistent effort to exclude social relations from her scope, Spelke even downplays in a note the remarkable social abilities of nonhuman primates by attributing their behavior to a “mix of evolutionary adaptations and by a slow, associative learning” (note 5, p. 412).

And yet, studies on dominance clearly demonstrate that the social world cannot be reduced to mind interactions. From an evolutionary perspective, it makes sense. It would be surprising if humans evolved a finely tuned core knowledge dedicated to numbers or places while only having a single-core system devoted to social relations. To navigate their social environment, people and especially children must identify relevant social relationships, adjust to the norms that govern them, and anticipate the sanctions that reinforce them (Charafeddine et al., 2015; Clément, Bernard, & Kaufmann, 2011). They must assess, *within a given relationship*, who bears obligations toward whom, who possesses the authority to impose duties, who has the entitlement to claim certain goods, and what forms of retaliation are deemed appropriate for obligation violations (Jackendoff, 1999). Social relationships have a normative structure that, for instance, excludes the very possibility of experience sharing in situation of dominance. Another recent study might plead for a relational model of social cognition. In their wonderful experiments on the early concept of intimacy, Spelke and colleagues show that 8- to 10-month-old infants can identify the concrete features of an interaction (saliva sharing and food licking) to infer a type of social relationship (intimacy) and then to use it to expect some apparently unrelated behaviors to occur (helping a person in distress) (Thomas, Woo, Nettle, Spelke, & Saxe, 2022). Because infants set aside the individual traits of specific characters (such as the prosociality of the actress) to focus on the social relationship itself, this experiment supports the idea of a domain-specific relational processing.

Admitting social relationship within the realm of core knowledge brings us back to our initial, ontological comment about numbers and forms. What might be the place of relationship processing in our cognitive architecture? Is social relationship a quasi-perceptual entity that deserves an ontological status of its own or a second-order, language-infected property? We have argued elsewhere that social relationships might be an ontological primitive (Kaufmann & Clément, 2014). Patterns of relationships are recognizable *in situ* as a succession of constraining and enabling *affordances*, each action affording a set of possible

subsequent actions. Moreover, they have the strong inductive potential that characterizes abstract concepts. Finally, they are mostly shared with nonhuman primates, one of the main conditions for core knowledge according to Spelke.

Within this perspective, the phrase “navigating in the social world” must be taken literally. During their first months, infants must position themselves not only in physical space but also in social space. They must identify the normative boundaries, obstacles, and gaps that obstruct their path. They have also to map the *power dynamics* they are entangled in. Indeed, dominance is not solely a concern for animals in the wild or adults in culture; it is a vital concern for infants, who are constantly confronted with their complete state of dependency (Hrdy & Burkart, 2020). They have to evaluate their own power, in terms of capacity and possibility of acting as well as in terms of permission and constraint. To put it otherwise, the task of newborns is not primarily to *know* their surroundings, but rather to *find their place* among others and assume the position that awaits them within a pre-established kinship and community. As Rochat (2003) suggests it, the infant’s sense of self is not “*ego-logical*” but “*eco-logical*”: It is based on the proprioceptive cartography and relational mapping of her spatial and social environment. Of course, to better understand this self-in-the-making, Spelke’s brilliant odyssey of the human mind is indispensable.


**Financial support.** This research received no specific grant from any funding agency, commercial, or not-for-profit sectors.

**Competing interests.** None.

## References

- Charafeddine, R., Mercier, H., Clément, F., Kaufmann, L., Berchtold, A., Reboul, A., & Van der Henst, J.-B. (2015). How preschoolers use cues of dominance to make sense of their social environment. *Journal of Cognition and Development*, 16(4), 587–607.
- Clément, F., Bernard, S., & Kaufmann, L. (2011). Social cognition is not reducible to theory of mind: When children use deontic rules to predict the behaviour of others. *British Journal of Developmental Psychology*, 29, 910–928.
- Hirschfeld, L. A. (2001). On a folk theory of society: Children, evolution, and mental representations of social groups. *Personality and Social Psychology Review*, 5, 107–117.
- Hrdy, S. B., & Burkart, J. M. (2020). The emergence of emotionally modern humans: Implications for language and learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 375(1803), 20190499.
- Jackendoff, R. S. (1999). The natural logic of rights and obligations. In R. Jackendoff, P. Bloom, & K. Wynn (Eds.), *Language, logic, and concepts: Essays in memory of John Macnamara* (pp. 67–95). MIT Press.
- Kaufmann, L., & Clément, F. (2014). Wired for society: Cognizing pathways to society and culture. *Topoi*, 33, 459–475.
- Kinzler, K. D., & Spelke, E. S. (2007). Core systems in human cognition. *Progress in Brain Research*, 164, 257–264. [https://doi.org/10.1016/S0079-6123\(07\)64014-X](https://doi.org/10.1016/S0079-6123(07)64014-X)
- Mascaro, O., & Csibra, G. (2012). Representation of stable social dominance relations by human infants. *Proceedings of the National Academy of Sciences of the United States of America*, 109, 6862–6867.
- Pun, A., Birch, S. A., & Baron, A. S. (2016). Infants use relative numerical group size to infer social dominance. *Proceedings of the National Academy of Sciences of the United States of America*, 113(9), 2376–2381.
- Rochat, P. (2003). Five-levels of self-awareness as they unfold early in life. *Consciousness and Cognition*, 12, 717–731.
- Spelke, E. S. (2022). *What babies know: Core knowledge and composition*. Oxford University Press.
- Thomas, A. J., Woo, B., Nettle, D., Spelke, E. S., & Saxe, R. (2022). Early concepts of intimacy: Young humans use saliva sharing to infer close relationships. *Science (New York, N.Y.)*, 375(6578), 311–315.
- Thomsen, L., Frankenhuys, W. E., Ingold-Smith, M. C., & Carey, S. (2011). Big and mighty: Preverbal infants mentally represent social dominance. *Science (New York, N.Y.)*, 331, 477–480.

## More than language is needed to represent and combine different core knowledge components

Peter Krøjgaard\* , Trine Sonne and Osman S. Kingo

Department of Psychology and the Behavioural Sciences, Aarhus University, Aarhus, Denmark  
[peter@psy.au.dk](mailto:peter@psy.au.dk)  
[trines@psy.au.dk](mailto:trines@psy.au.dk)  
[osman@psy.au.dk](mailto:osman@psy.au.dk)

\*Corresponding author.

doi:10.1017/S0140525X23003242, e135

### Abstract

We question Spelke's key claim that the medium, in which contents from different core knowledge systems can be represented and combined, is language-based. Recalling an episodic memory, playing chess, and conducting mental rotation are tasks where core knowledge information is represented and combined. Although these tasks can be *described* by means of language, these tasks are not inherently language-based. Hence, language may be an important *subset* of an abstraction medium – not the medium as such.

In her book *What babies know*, Spelke (2022) presents the most recent and comprehensive account of her core knowledge theory on infant cognition. A central claim in Spelke's earlier accounts on her core knowledge approach (e.g., Spelke, Breinlinger, Macomber, & Jacobson, 1992) was that only human beings possessed domain-specific core knowledge systems. However, several of these core knowledge systems have subsequently been found to be present in other animals. Therefore, core knowledge systems *per se* could no longer explain why the world of human beings is so radically different from that of other species. For instance, only human beings invent and use smart phones, spaceships, and generic symbol systems such as the alphabet and the 10-digit systems. Hence, on accepting the core knowledge approach, it raises the key question of why human beings, relative to any other species, are so much more efficient in employing and combining their core knowledge domains. In short: What is the extra “secret sauce”? In *What babies know*, Spelke (2022) suggests that it is our *language* that sets human beings apart from all other creatures on Earth. First, Spelke presents the bold claim that language is the driving force when infants around their first birthday develop uniquely human concepts of other people and their mental states. Second, Spelke claims that it is language that provides the medium in which contents from different core knowledge systems can be represented and combined economically and efficiently. Here, we concentrate on the second claim.

In general, we find the book highly stimulating and a pleasure to read! We endorse the effort of presenting all the compelling evidence of core knowledge in both young human beings and animals – as well as the thought-provoking suggestion regarding language as the “secret sauce,” allowing core knowledge representations to meet and to be combined mainly because of the recursively combinatorial and compositional properties of

language. However, it may not be language as such, but rather the ability to abstract core knowledge features from a given domain and to combine these in an abstract medium that is key. This medium can, but does not, in our view, have to be, language-based *per se*. Language may thus be an important *subset* of this abstraction medium – not the medium as such.

Our reasoning for questioning that the hub for core knowledge information to meet and be combined should be exclusively language-based is the following: Although language is beyond any doubt a powerful abstract medium for representing and combining information from different core knowledge domains, language may not be the candidate medium for *all* such tasks. In fact, many intellectual tasks that most researchers would consider uniquely human, and in which core knowledge is combined, do not seem to be solved in a language-based medium. Consider the highly diverse cognitive tasks of recalling an episodic memory (e.g., Tulving, 2005), playing chess (e.g., Coates, 2013), or conducting mental rotation tasks (e.g., Shepard & Mentzler, 1971). Although each of these activities involves core knowledge components of at least the domains “object” and “space,” and can be, and often is, *described* or *reported* by means of language, none of them is predominantly *conducted* in a language-based medium. This can be elaborated as below.

### Recalling episodic memories

Recalling an episodic memory refers to the task of recalling specific, personally experienced events from the past (e.g., Tulving, 2005), as for instance the first author remembering the birth of his third child or recalling playing golf with his brother last Saturday. Although these recollections are very different, they share the features of concerning personally experienced unique events that took place at a certain place and at a certain point in time. Episodic memories are crucial for our sense of who we are and the history we share with other people (e.g., Harris, Rasmussen, & Berntsen, 2014). From an evolutionary perspective, episodic memories also have important directive functions by reminding us of successes and mistakes in the past, which helps us to fare better tomorrow than yesterday (e.g., Allen & Fortin, 2013).

A key feature of episodic memories is *autonoetic awareness*, that is, the “I-was-there-sense,” which reflects an inherent part of recall of episodic memories in our mind's eye (Tulving, 2005). For the sake of simplicity, we will use the “golf-memory” mentioned above as an example. When recalling playing golf with his brother last Saturday, the first author relives playing in the pine woods at Nordvestjysk Golf Club, Denmark. This recollection is filled with images (e.g., the beautiful view from the first tee), sensory input concerning the weather (e.g., it was remarkably warm and sunny for September in Denmark, although the forecast had predicted rain), and the setting (e.g., the distinct smell of pine trees in the forest, and experiencing that the greens were slower than they looked), as well as emotions concerning the quality of his play (he played terrible, with several lost balls and missed puts). Although the recollection involves language (e.g., conversations with his brother on the golf course) and can be reported by means of language as attempted here, the autonoetic awareness of recalling this (or any other) specific event is not primarily language-based. Rather the recollection is predominantly multimodal and constituted by mental images (e.g., recalling the vision of his drive on hole 2 disappearing in the woods), sensory information (e.g., the pleasant surprise of experiencing the warm sun), and emotions (e.g., being disappointed of the



terrible play, but enjoying the company of his brother; e.g., Rubin, 2005, 2006; Tulving, 2005).

The claim that language is not a crucial aspect of episodic memories is further supported by two key empirical findings. First, language impairments do typically not affect the ability to recall personally experienced events from the past (Rubin, 2006). If language had been key in episodic memory recollections, then language impairments should have affected recall. Second, the case of KC (Rosenbaum et al., 2005), who because of a motor-cycle accident had virtually no episodic memories, had no apparent language impairments. Again, if language had been the medium for episodic recollections, it becomes difficult to explain the virtually complete lack of episodic memories in KC, who at the same time had no obvious language impairments. Finally, some scholars claim that episodic memories as well as the ability to envision the future may have evolved *before* language, not after (Kellogg, 2023). In short: Episodic memories are central in human beings' mental life, and according to some scholars, uniquely human (e.g., Tulving, 2005). Although episodic memories typically involve representing and combining information from core knowledge systems (e.g., "object" and "space"), the defining auto-noetic awareness (the "I-was-there-sense") is not language-based in nature.

### Playing chess

Playing chess involves a range of cognitive skills, including pattern recognition, memory, imagery, and decision making (e.g., Campitelli, 2017). It is common to distinguish between tactical and strategical aspects of the game. Tactical operations concern forced sequences of moves leading to a clear advantage for one player. Strategical operations involve more long-term planning concerning positional aspects of the position such as for instance pawn structures and piece placement and activity (e.g., Coates, 2013). Because of space limitations, we here consider tactical aspects only. Occasionally, during a mating attack, there may for instance be a theoretical forced mate in three moves. This means that a player can win for certain, if she manages to find the correct sequence of three consecutive moves, regardless of which moves the opponent chooses as responses. To find such a sequence the player will have to, in her minds' eye, imagine all possible move orders until the end – a process at times (and slightly misleading) called *search* in the cognitive literature on chess (Coates, 2013). Finding a forced mate in three moves clearly draws on core knowledge information from at least the domains "object" and "space," as well as conforming to the rules of chess, of course. This information needs to be represented, combined, and processed in a mental space. However, when solving such a task, language plays a very minor role, if any at all. Rather, the process mainly draws on imagery in a mental space, in which the moves of the chess pieces are simulated one by one in the minds' eye of the attacker (e.g., Campitelli, 2017; Coates, 2013). The process of finding a forced mate in three moves (or any other tactical chess maneuver) can be *described* to others in language, but the process as such is not language-based.

The importance of imagery (but not language) when playing chess is further supported by the results from studies on blindfold chess (i.e., an especially challenging version of chess where you play *without* a visible chess board and pieces, and hence, in your mind's eye only). In one illustrative experiment (Saariluoma, 1991, exp. 1), chess players were given the task of following three already played chess games in their mind's eye.

The moves of the games were read out to the players by means of conventional algebraic chess annotation (i.e., a standard chess annotation system) one by one at the pace of one move (for both white and black) every fourth second. After 15 and 25 moves the players were asked about the position of the pieces. While attempting to follow the games, the players were exposed to one of the three different kinds of possible interferences in the form of additional tasks: (1) No interference (control), (2) articulatory interference (repetitive pronunciation of a syllable), and (3) imaging interference (to imagine the syllable). The results revealed that although the articulatory interference (and the control) had no diminishing effect on performance, imaging interference had (Saariluoma, 1991). Thus, even though the moves were presented in a language format, only imaging (but not articulatory) interference negatively affected performance. If language had been crucial for keeping track of the game, then articulatory interference should have affected performance, but it did not. Hence, the task of keeping track of an unfolding chess game by means of verbally presented chess moves in algebraic annotation does not seem to be conducted in a language-based mental medium, but predominantly in a medium based on visual imagination.

### Conducting mental rotation

In their seminal paper, Shepard and Mentzler (1971) asked participants to assess whether possibly identical three-dimensional (3D) objects presented at different angles were the same or different. To compare the two objects from the same angle, participants had to mentally rotate one of the objects. The results revealed a linear relation between the required rotation angle for direct comparison and response time. Mental rotation involves core knowledge information from the domains "object" and "space," but how are these bits of information represented and combined? It has been debated whether mental rotation is a modal process (e.g., Kosslyn, Ganis, & Thompson, 2001), or whether it may be conducted by means of abstract language-based propositions (e.g., Pylyshyn, 2003). However, the evidence seems to favor the modal approach. First, as mentioned above, Shepard and Mentzler (1971) found a linear relationship between rotation angle and response time, which is in accordance with the interpretation that mental rotation is a modal act. If mental rotation had been carried out by language-based propositions, there is no obvious reason why response time should be a linear function of rotation angle. Second, brain-imaging studies reveal that both the ventral and the dorsal streams are activated, when participants conduct mental rotation tasks (e.g., Milivojevic, Hamm, & Corballis, 2008). Again, these results are in accordance with the modal interpretation, whereas a language-based interpretation seems less straightforward.

### Conclusion

In summary, the above examples represent different tasks in which core knowledge information is represented and combined in an abstract medium that does not seem to be language-based as proposed by Spelke. Although we acknowledge that many cognitive tasks in which core knowledge information is represented and combined take place in a language-based medium (as for instance writing this commentary), there are important exceptions as exemplified above. Consequently, language only seem to constitute a subset of a candidate abstract medium serving to represent and combine core knowledge components when solving problems.


**Financial support.** This research received no specific grant from any funding agency, commercial, or not-for-profit sectors.

**Competing interest.** None.

## References

- Allen, T. A., & Fortin, N. J. (2013). The evolution of episodic memory. *Proceedings of the National Academy of Sciences of the United States of America*, *110*, 10379–10386.
- Campitelli, G. (2017). Cognitive processes in chess. In D. Z. Hambrick, G. Campitelli, & B. N. Macnamara (Eds.), *The science of expertise: Behavioral, neural, and genetic approaches to complex skill* (pp. 31–46). Routledge.
- Coates, P. (2013). Chess, imagination, and perceptual understanding. *Royal Institute of Philosophy Supplements*, *73*, 211–242.
- Harris, C. B., Rasmussen, A. S., & Berntsen, D. (2014). The functions of autobiographical memory: An integrative approach. *Memory (Hove, England)*, *22*, 559–581.
- Kellogg, R. T. (2023). More than language: Mental time travel, mentalizing, executive attention, and the left-hemisphere interpreter in human origins. *Psychological Review*, *130*, 1592–1611.
- Kosslyn, S. M., Ganis, G., & Thompson, W. L. (2001). Neural foundations of imagery. *Nature Reviews Neuroscience*, *2*, 635–642.
- Milivojevic, B., Hamm, J. P., & Corballis, M. C. (2008). Functional neuroanatomy of mental rotation. *Journal of Cognitive Neuroscience*, *21*, 945–959.
- Pylyshyn, Z. (2003). Return of the mental image: Are the really pictures in the brain? *Trends in Cognitive Sciences*, *7*, 113–118.
- Rosenbaum, R. S., Köhler, S., Schacter, D. L., Moscovitch, M., Westmacott, R., Black, S. E., Gao, F., & Tulving, E. (2005). The case of K.C.: Contributions of a memory-impaired person to memory theory. *Neuropsychologia*, *43*, 989–1021.
- Rubin, D. C. (2005). A basic systems approach to autobiographical memory. *Current Directions in Psychological Science*, *14*, 79–83.
- Rubin, D. C. (2006). The basic-systems model of episodic memory. *Perspectives on Psychological Science*, *1*, 277–311.
- Saariluoma, P. (1991). Aspects of skilled imagery in blindfold chess. *Acta Psychologica*, *77*, 65–89.
- Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science (New York, N.Y.)*, *171*, 701–703.
- Spelke, E. S. (2022). *What babies know: Core knowledge and composition* (Vol. 1). Oxford University Press.
- Spelke, E. S., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of knowledge. *Psychological Review*, *99*, 605–632.
- Tulving, E. (2005). Episodic memory and autoecesis: Uniquely human? In H. S. Terrace and J. Metcalfe (Eds.), *The missing link in cognition: Origins of self-reflective consciousness* (pp. 3–56). Oxford University Press.

## Not all core knowledge systems are created equal, and they are subject to revision in both children and adults

Rongzhi Liu and Fei Xu 

Department of Psychology, University of California, Berkeley, CA, USA  
[rongzhi\\_liu@berkeley.edu](mailto:rongzhi_liu@berkeley.edu)  
[www.babylab.berkeley.edu](http://www.babylab.berkeley.edu)

### Corresponding author:

Fei Xu; Email: [fei\\_xu@berkeley.edu](mailto:fei_xu@berkeley.edu)

doi:10.1017/S0140525X23003084, e136

### Abstract

Core knowledge systems play an important role in theories of cognitive development. However, recent studies suggest that fundamental principles of the object and agent systems can be revised by adults and preschoolers, when given small amounts

of counterevidence. We argue that not all core knowledge systems are created equal, and they may be subject to revision throughout development.

Spelke (2022) presents a comprehensive and elegant account of the origin of knowledge. She also presents strong arguments for the existence of six core knowledge systems that guide human learning and reasoning from infancy on. We are sympathetic to her view, but recent evidence also suggests a much more nuanced picture. We make two points in this commentary: (1) Some of the core knowledge systems are subject to revision, in children and adults. New studies show that surprisingly, even with just a small amount of counterevidence, adults and preschoolers readily revise their beliefs about core principles for both objects and agents (Liu & Xu, 2021, 2022, 2023). (2) Not all core knowledge systems are created equal. Given the body of evidence we have in hand, we argue that OBJECTS and NUMBER are perceptual systems, whereas AGENTS and SOCIAL BEINGS are more likely to be part of our belief system.

Recent studies investigated whether the core principles guiding our reasoning in the object and the agent systems are revisable (Liu & Xu, 2021, 2022, 2023). Adults and preschoolers observed a few pieces of evidence that violated the core principles of objects (e.g., a ball can go through a wall) and agents (e.g., an agent always takes an inefficient path to reach her goal). Then they made predictions about new events that were progressively more different from the events they observed. They were more likely to predict outcomes inconsistent with the core principles after observing the violations. Thus, both the object and the agent systems are subject to revision in adults and preschoolers. Furthermore, adults and preschoolers had stronger prior beliefs for objects than for agents, and the physical principles were harder to revise than the psychological principles in two ways: They were less likely to generalize the revised physical principles to new objects and new events; when they were asked to explain the violations, they were less likely to accept the counterevidence and more likely to try to explain it away (e.g., “there is a gap between the wall and the screen so the ball can go through”). In contrast, learners readily generalized the revised psychological principles to new agents and new events, and they accepted the counterevidence and generated plausible reasons for the agent’s unusual behavior (e.g., “the red child just likes to jump,” instead of taking the most efficient path to reach her goal). What explains these domain differences? One possibility is that infants are born with stronger prior beliefs about objects (i.e., the object system is more hard-wired to begin with); another possibility is that children and adults have observed more counterevidence about the psychological principles in everyday life, and therefore have weaker and more flexible beliefs about agents.

These findings also suggest that maybe not all core knowledge systems are created equal. We speculate that there might be two types of qualitatively different core knowledge systems – one type is more akin to perceptual systems, which are automatic, inflexible, and possibly encapsulated from conscious reasoning, and the other type resembles belief systems, which are more flexible and deliberate. We argue that the systems of objects and number (and perhaps space) may be of the first type, whereas the systems of agents and social beings (and perhaps form) are more likely to be of the second type.

A large body of research suggests that adults’ object representation depends on perceptual mechanisms (Scholl, 2001), and

perception of objects is disrupted when objects do not follow the core physical principles such as continuity and cohesion (Scholl & Pylyshyn, 1999; Scholl, Pylyshyn, & Feldman, 2001; vanMarle & Scholl, 2003). Furthermore, object perception seems to be unaffected by the top-down influences of cognition (Firestone & Scholl, 2016).

For the number system, past research has shown clear evidence that the approximate number system (ANS) activates automatically and unconsciously in all ages (Izard, Sann, Spelke, & Streri, 2009; Nieder & Dehaene, 2009). The precision of ANS increases during infancy, perhaps because of the improvement of visual acuity (Xu & Arriaga, 2007; Xu & Spelke, 2000). In addition, the neurological signatures of the ANS remain constant from infancy to adulthood, unaffected by years of mathematical education (Hyde & Spelke, 2009, 2011).

On the contrary, the systems of agents and social beings are less automatic and encapsulated, and more likely to be part of our belief systems. Three-month-old infants do not automatically expect agents' actions to be directed to objects; they flexibly learn the goal (objects or location) of an agent's actions based on the agent's previous behaviors (Woo, Liu, & Spelke, 2022). Although 1-year-old infants and children older than 4 years expect agents to take efficient paths to achieve their goals, 3-year-olds fail to show this expectation, suggesting that the development of the efficiency principle might be discontinuous (Gergely & Csibra, 2003; Gönül & Paulus, 2021).

Similarly, for the system of social beings, although expectations about how individuals interact and affiliate with one another emerge at a young age, these expectations are flexible and can be changed by infants' own social experiences. For instance, infants' social environments modulate their same-race preference – White and Black infants living in monoracial environments prefer faces of their own race, but Black infants living in predominantly White environments do not show a same-race preference (Bar-Haim, Ziv, Lamy, & Hodes, 2006). Infants' linguistic environments also change their expectations about social groups – monolingual infants expect individuals who speak different languages to have different food preferences, but bilingual infants expect them to share food preferences (Lieberman et al., 2016).

This distinction between perceptual versus conceptual core knowledge systems makes interesting predictions that can be tested in future research. For example, preschoolers' and adults' revision of the core physical principles in Liu and Xu (2021, 2022) may not affect the operation of these principles on the perceptual level – participants may revert to principle-consistent predictions about novel events when they are under cognitive load. More generally, learners may be more likely to accept the violations of the agent and social being systems compared to the object and number systems.

**Financial support.** This research received no specific grant from any funding agency, commercial, or not-for-profit sectors.



**Competing interests.** None.

## References

- Bar-Haim, Y., Ziv, T., Lamy, D., & Hodes, R. M. (2006). Nature and nurture in own-race face processing. *Psychological Science*, 17(2), 159–163.
- Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for “top-down” effects. *Behavioral and Brain Sciences*, 39, e229.
- Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naïve theory of rational action. *Trends in Cognitive Sciences*, 7(7), 287–292.
- Gönül, G., & Paulus, M. (2021). Children's reasoning about the efficiency of others' actions: The development of rational action prediction. *Journal of Experimental Child Psychology*, 204, 105035.

- Hyde, D. C., & Spelke, E. S. (2009). All numbers are not equal: An electrophysiological investigation of small and large number representations. *Journal of Cognitive Neuroscience*, 21(6), 1039–1053.
- Hyde, D. C., & Spelke, E. S. (2011). Neural signatures of number processing in human infants: Evidence for two core systems underlying numerical cognition. *Developmental Science*, 14(2), 360–371.
- Izard, V., Sann, C., Spelke, E. S., & Streri, A. (2009). Newborn infants perceive abstract numbers. *Proceedings of the National Academy of Sciences of the United States of America*, 106(25), 10382–10385.
- Lieberman, Z., Woodward, A. L., Sullivan, K. R., & Kinzler, K. D. (2016). Early emerging system for reasoning about the social nature of food. *Proceedings of the National Academy of Sciences of the United States of America*, 113(34), 9480–9485.
- Liu, R., & Xu, F. (2021). Revising core beliefs in young children. In T. Fitch, C. Lamm, H. Leder, & K. Tessmar-Raible (Eds.), *Proceedings of the 43rd annual conference of the Cognitive Science Society* (pp. 1677–1683). Cognitive Science Society.
- Liu, R., & Xu, F. (2022). Can adults revise their core beliefs about objects? In J. Culberston, A. Perfors, H. Rabagliati, & V. Ramenzoni (Eds.), *Proceedings of the 44th annual conference of the Cognitive Science Society* (pp. 1483–1489). Cognitive Science Society.
- Liu, R., & Xu, F. (2023). Can adults revise their core beliefs about agents? In M. Goldwater, F. Anggoro, B. Hayes, & D. Ong (Eds.), *Proceedings of the 45th annual conference of the Cognitive Science Society* (pp. 1637–1643). Cognitive Science Society.
- Nieder, A., & Dehaene, S. (2009). Representation of number in the brain. *Annual Review of Neuroscience*, 32, 185–208.
- Scholl, B. J. (2001). Objects and attention: The state of the art. *Cognition*, 80(1–2), 1–46.
- Scholl, B. J., & Pylyshyn, Z. W. (1999). Tracking multiple items through occlusion: Clues to visual objecthood. *Cognitive Psychology*, 38(2), 259–290.
- Scholl, B. J., Pylyshyn, Z. W., & Feldman, J. (2001). What is a visual object? Evidence from target merging in multiple object tracking. *Cognition*, 80(1–2), 159–177.
- Spelke, E. S. (2022). *What babies know: Core knowledge and composition volume 1* (Vol. 1). Oxford University Press.
- vanMarle, K., & Scholl, B. J. (2003). Attentive tracking of objects versus substances. *Psychological Science*, 14(5), 498–504.
- Woo, B. M., Liu, S., & Spelke, E. (2022, June 30). Infants rationally infer the goals of other people's reaches in the absence of first-person experience with reaching actions. *PsyArXiv*. <https://doi.org/10.31234/osf.io/dx2er>
- Xu, F., & Arriaga, R. I. (2007). Number discrimination in 10-month-old infants. *British Journal of Developmental Psychology*, 25(1), 103–108.
- Xu, F., & Spelke, E. S. (2000). Large number discrimination in 6-month-old infants. *Cognition*, 74(1), B1–B11.

## Concepts, core knowledge, and the rationalism–empiricism debate

Eric Margolis<sup>a\*</sup>  and Stephen Laurence<sup>b</sup> 

<sup>a</sup>Department of Philosophy, University of British Columbia, Vancouver, BC, Canada and <sup>b</sup>Department of Philosophy and Hang Seng Centre for Cognitive Studies, University of Sheffield, Sheffield, UK  
[eric.margolis@ubc.ca](mailto:eric.margolis@ubc.ca)  
[S.Laurence@sheffield.ac.uk](mailto:S.Laurence@sheffield.ac.uk)  
<https://philosophy.ubc.ca/profile/eric-margolis/>  
<https://www.sheffield.ac.uk/philosophy/people/academic-staff/stephen-laurence>

\*Corresponding author.

doi:10.1017/S0140525X23003072, e137

### Abstract

While Spelke provides powerful support for concept nativism, her focus on understanding concept nativism through six innate core knowledge systems is too confining. There is also no reason to suppose that the *curse of a compositional mind* constitutes a principled reason for positing less innate structure in explaining the origins of concepts. Any solution to such problems must take into account poverty of the stimulus considerations, which argue for postulating more innate structure, not less.



*What Babies Know* is a landmark achievement, consolidating and developing Spelke's many important contributions to the rationalism–empiricism debate concerning the origins of concepts. The account it offers has two major parts. One is that children's learning begins with six innate core knowledge systems. This provides them with an initial starting point comprised, not simply of low-level sensorimotor representations, but also of abstract concepts, allowing them to form a richer initial conception of the world around them. The other is that, unlike other animals (many of whom share a broadly similar representational starting point), children have language, which both allows and encourages them to form concepts that can range over, combine, and extend the content domains that are associated with their individual core knowledge systems.

While Spelke's focus is not on how the rationalism–empiricism debate about the origin of concepts should be understood, the broad outlines of how her account fits into the larger landscape of positions in this debate is clear: It's meant to offer a rationalist theory of the origins of concepts that avoids the excesses of Fodor's notorious *radical concept nativism*, which takes virtually all lexical concepts to be innate (Fodor, 1981). Spelke's six innate core knowledge systems, and the idea that evolution is likely to have favored learning mechanisms that employ highly abstract forms of representation, sharply differentiate her account from standard empiricist accounts. At the same time, she is equally clearly not a Fodorian radical concept nativist, as she refuses to postulate further innate concepts on the grounds that they aren't seen in newborns and (specifically with reference to Fodor) because she thinks that having too much innate psychological structure makes learning harder, not easier – a problem that she calls “the curse of a compositional mind.”

There is much that we admire and agree with in the immensely important body of work by Spelke and others that this book synthesizes. In particular, while we have a different understanding of some of the core knowledge systems that she posits,<sup>1</sup> we think that there are overwhelming grounds for accepting innate abstract representations in each of the six areas of core knowledge discussed. We also agree that language plays a vital role in conceptual development. Despite these substantial points of agreement, we will argue that it's too confining to think of rationalist views in this debate, as Spelke seems to, only in terms of how many and which core knowledge systems there are.

To see why, we first need to clarify what is and isn't at stake in the rationalism–empiricism debate about the origins of concepts, since the range of options available in the debate depend in part on how it is framed. On some ways of framing the debate, Spelke's view turns out to not be a version of concept nativism at all. On such views, a minimal requirement for being a version of concept nativism is positing some innate concepts. And while it's clear that Spelke's core knowledge systems involve abstract representations of some sort or another, these representations would not count as concepts on a number of different accounts of what makes a representation conceptual. For example, some accounts distinguish concepts from nonconceptual representations in terms of something like the *generality constraint*, according to which a representation is a concept only if it can be flexibly and freely combined with all other concepts (e.g., Evans, 1982). Others hold that concepts are constitutively linked to capacities for conscious rational reflection, including the ability to justify one's use of a concept (e.g., McDowell, 1994). On either of these approaches, the representations in Spelke's core knowledge systems clearly aren't concepts. These systems are supposed to be, for the most part, Fodorian modules which automatically

respond to very specific types of input and that operate independently of one another, competing for cognitive resources such as attention. The fact that they aren't involved in general forms of rational reflection and that their representations aren't capable of being freely combined with conceptual representations in other parts of the mind is built into the whole idea of a core knowledge system, as Spelke understands this construct. It's also what explains why, on her account, animals (who share these systems with humans) lack the form of flexible cognition that begins to appear once children's linguistic abilities are sufficiently developed.

This is not to say that these representations *aren't* concepts. It only means that it isn't uncontroversial to suppose that they are. There are many different ways of drawing the conceptual/nonconceptual distinction, and on other accounts of what concepts are, these representations turn out to be conceptual.<sup>2</sup> The point is that this matter remains enormously controversial (see Laurence & Margolis, 2012, for an analysis of this complex debate). So if the rationalism–empiricism debate is taken to be about whether there are innate concepts, it will be unclear whether or not Spelke's position should be understood as a form of concept nativism.

In our view, however, this isn't how the debate about the origins of concepts should be understood, and so none of this should be seen as an objection to Spelke. The requirement that concept nativists must posit innate concepts stems from one of a number of common misunderstandings about the rationalism–empiricism debate. What's at stake in this debate is not whether or not there are innate concepts. It's whether the foundational psychological basis for acquiring concepts – what we call the *acquisition base* for learning – is rationalist or empiricist. The acquisition base is the collection of all psychological structures that are not themselves learned (or otherwise the product of any psychological-level process of acquisition) and which form the ultimate psychological basis for acquiring all concepts (Margolis & Laurence, 2013, 2023). On an empiricist account, the acquisition base is largely confined to sensorimotor representations and domain-general learning mechanisms. On a rationalist account, it also includes such things but also various types of abstract representations and domain-specific learning mechanisms. Spelke's account is clearly rationalist – a version of concept nativism – in light of the fact that she effectively holds that the acquisition base includes six domain-specific core knowledge systems, each of which encompasses a set of interrelated abstract representations. Whether these representations are concepts per se (according to some preferred account of the conceptual/nonconceptual distinction) may be difficult to say. But on our understanding of the rationalism–empiricism debate about the origins of concepts, this is irrelevant. All that matters is whether the acquisition base for acquiring concepts is rationalist.

However, while this objection doesn't pose a direct problem for Spelke's account, once we see that what matters to this debate is the character of the acquisition base, it immediately follows that a rationalist account of the origins of concepts can be grounded in many different types of rationalist psychological structures in the acquisition base above and beyond any core knowledge systems that are part of the acquisition base. Seen in this light, Spelke's focus on core knowledge is too confining. Concept nativists should be prepared for the existence of many different types of structures in the acquisition base that don't fit the mold of core knowledge systems. They may well not be evidenced as being present at birth – as in the case of universal grammar, for example. They may be relatively isolated structures, not part of a tightly interrelated set of concepts. They may involve attentional biases of various kinds, domain-specific heuristics, links between innate

cognitive systems, domain-specific dispositions in memory, particular types of motivations, mechanisms for causal-explanatory reasoning, representations for basic metaphysical distinctions (e.g., events, individuals, and kinds), and much else. So, this debate should not be understood as being about how many or which core knowledge systems exist.

Finally, what should we make of the curse of a compositional mind (Spelke's claim that too much innate psychological structure makes learning harder, not easier)? We agree that the possibility of there being too many hypotheses to entertain in acquiring a new concept (or any type of psychological structure) is an important problem. But as we see it, this problem doesn't stem from the postulation of a large number of innate abstract representations and is not solved by substantially reducing the initial representational resources that a learner has access to.

First, as Chomsky noted in connection with poverty of the stimulus arguments, learners face the problem of there being indefinitely many mistaken or unproductive hypotheses to entertain even if they lack the resources to formulate the correct hypothesis. In fact, even a learner with only concepts or representations pertaining to basic sensory or perceptual properties might endlessly entertain different hypotheses involving just combinations of such representations without ever hitting on the correct hypothesis for acquiring some more abstract concept. Second, the fact that adults possess a vastly greater number of abstract concepts than infants but aren't crippled by the curse of a compositional mind shows that it is not the addition of abstract concepts per se that creates the problems here. So, there is no reason to think that the curse of a compositional mind is a principled objection to having far more innate resources in the acquisition base than Spelke's six core knowledge systems. Third, poverty of the stimulus considerations also suggests that the acquisition of interestingly new psychological traits can crucially depend on having a richer initial representational starting point (Laurence & Margolis, *in press*). For example, how could a learner who didn't possess any prior capacity for representing modality acquire wholly new concepts pertaining to what *might* have been or what *must* be? Without some initial glimmer that there is more to the world than how things actually are, learners wouldn't be in a position to formulate such modal concepts and wouldn't even see the point of doing so.

Taken together, these considerations argue that the solution to the curse of a compositional mind should be one that simultaneously helps infants avoid problems tied to poverty of the stimulus considerations. For this reason, it is likely that the solution involves adding *more* to the acquisition base, not pairing down its representational resources. Like adults, children might also have domain-specific principles, heuristics, architectural constraints on inferences, or other cognitive resources that allow and encourage them to apply representations in a restricted way. The organization of core knowledge systems provides one way of constraining the hypothesis space in this way. But there are many other possibilities as well. Accordingly, the curse of a compositional mind doesn't give any reason at all to suppose that the number of innate abstract representations (or for that matter, the number of core knowledge systems) is bound to be small.

Spelke is right that having a large number of innate representations or potential combinations of innate representations doesn't necessarily make learning easier; but likewise, it doesn't necessarily make learning harder either. The question is simply how much such structure needs to be postulated to explain the origins of all the concepts we can acquire. So, while we agree with Spelke that some form of concept nativism provides the

best account of the origins of concepts, we think that this account will ultimately prove to be one that involves a considerably richer acquisition base than the account grounded in Spelke's six core knowledge systems.<sup>3</sup>

**Financial support.** This work was supported by Canada's Social Sciences and Humanities Research Council (grant number 435-2019-0002).

**Competing interest.** None.

## Notes

1. For example, we have argued that Spelke's account of the origins of natural number concepts doesn't posit enough innate structure specific to this domain. What's missing is a system that represents small precise numerical quantities as such – representations for *one*, *two*, and *three* (see Laurence & Margolis, 2005, 2007; Margolis, 2020; Margolis & Laurence, 2008; for a related account, see Leslie, Gallistel, & Gelman, 2007).
2. For example, they would arguably come out as conceptual on Fodor's account, according to which the conceptual/nonconceptual distinction roughly corresponds to the distinction between iconic and discursive representations (Fodor, 2008).
3. This article was fully collaborative; the order of the authors' names is arbitrary.

## References

- Evans, G. (1982). *The varieties of reference*. Oxford University Press.
- Fodor, J. A. (1981). The present status of the innateness controversy. In J. A. Fodor (Ed.) *Representations: Philosophical essays on the foundations of cognitive science* (pp. 257–316). MIT Press.
- Fodor, J. A. (2008). *LOT 2: The language of thought revisited*. Oxford University Press.
- Laurence, S., & Margolis, E. (2005). Number and natural language. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The innate mind: Structure and contents* (pp. 216–235). Oxford University Press.
- Laurence, S., & Margolis, E. (2007). Linguistic determinism and the innate basis of number. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The innate mind: Foundations and the future* (pp. 139–169). Oxford University Press.
- Laurence, S., & Margolis, E. (2012). The scope of the conceptual. In E. Margolis, R. Samuels, & S. Stich (Eds.), *The Oxford handbook of philosophy of cognitive science* (pp. 291–317). Oxford University Press.
- Laurence, S., & Margolis, E. (in press). *The building blocks of thought*. Oxford University Press.
- Leslie, A. M., Gallistel, C. R., & Gelman, R. (2007). Where integers come from. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The innate mind: Foundations and future* (pp. 109–138). Oxford University Press.
- Margolis, E. (2020). The small number system. *Philosophy of Science*, 87, 1–22.
- Margolis, E., & Laurence, S. (2008). How to learn the natural numbers: Inductive inference and the acquisition of number concepts. *Cognition*, 106(2), 924–939.
- Margolis, E., & Laurence, S. (2013). In defense of nativism. *Philosophical Studies*, 165(2), 693–718.
- Margolis, E., & Laurence, S. (2023). Making sense of domain specificity. *Cognition*, 240, 105583.
- McDowell, J. (1994). *Mind and world*. Harvard University Press.

## How do babies come to know what babies know?

David S. Moore<sup>a\*</sup>  and David J. Lewkowicz<sup>b</sup> 

<sup>a</sup>Psychology Field Group, Pitzer College, Claremont, CA, USA and <sup>b</sup>Child Study Center, Yale University School of Medicine, New Haven, CT, USA

[dmoore@pitzer.edu](mailto:dmoore@pitzer.edu); [david.lewkowicz@yale.edu](mailto:david.lewkowicz@yale.edu)  
<http://pzacad.pitzer.edu/~dmoore/>; <https://medicine.yale.edu/childstudy/research/autism-and-neurodevelopment/llamb/>

\*Corresponding author.

doi:10.1017/S0140525X23003102, e138

### Abstract

Elizabeth Spelke's *What Babies Know* is a scholarly presentation of core knowledge theory and a masterful compendium of empirical evidence that supports it. Unfortunately, Spelke's principal theoretical assumption is that core knowledge is simply the innate product of cognitive evolution. As such, her theory fails to explicate the developmental mechanisms underlying the emergence of the cognitive systems on which that knowledge depends.

Elizabeth Spelke's (2022) *What Babies Know* is scholarly, erudite, and often insightful; it is an intriguing, thought-provoking book packed with research results. It begins with an interesting question: "what do human infants know ... when their learning begins?" (p. xv). However, no clear definition of "learning" is offered, an omission that proves problematic. Learning entails functional changes that result from experiences, and because embryonic tissues – including ectodermal cells that become the first neurons – are functionally changed by the contexts they experience (Spemann & Mangold, 1924/2001), "learning" arguably begins before infants *know* anything at all. Although there is value in this book's collation of experimental data, the theoretical scaffolding that serves as the work's glue is nondevelopmental and outdated. Spelke's efforts to build a theory on the notion of "core" knowledge have done her empirical work a disservice by situating that work in a nondevelopmental theoretical framework.

Spelke's conceptualization of certain cognitive systems as "core" suggests that she is in thrall to an old *Weltanschauung* that sees some capabilities as inevitable outcomes of prenatal development, capabilities nineteenth-century theorists would have attributed to "nature" rather than "nurture" (Moore, 2001, 2013). But this dichotomous way of thinking about phenotype origins has been rendered obsolete by the work of developmental psychobiologists (Gottlieb, 2007; Michel & Moore, 1995) and molecular (Lewontin, 2000; Strohmman, 2003), physiological (Noble & Noble, 2023), and developmental (Gilbert & Epel, 2015) biologists. All have conclusively established that phenotypes are emergent products of a probabilistic, multifactorial, and context-dependent developmental process that depends on the interaction and coaction of genetic and nongenetic factors. (Nongenetic factors reside in multiple places: In the environment outside the body *and* both inside and outside of cells but still inside the body.) The consensus of developmental systems theorists is that cognitive systems *emerge* from complex, dynamic interactions between – and coactions of – these genetic and nongenetic factors, where emergence is likely to be characterized by reorganization of component systems during development. Clearly, this developmental systems perspective differs radically from the more predetermined view held by nativists like Spelke, a view that overlooks the need for investigations of the mechanisms underlying the developmental emergence of early-appearing cognitive skills.

The finding that some competencies are present at birth carries great significance for Spelke because it suggests to her that these competencies are not learned. But learning is only one component of experience, and starting with developmental processes that commence at conception, *any* experiential factor can potentially have profound effects on the developmental emergence of a cognitive/behavioral skill (Gottlieb, 1991; Lehrman, 1953). Indeed, several experience-dependent cognitive competencies have been detected in fetuses, as Spelke discusses in chapter 9;

such findings are consistent with the understanding that development is both a *continuous* process beginning at conception and one that normally entails reorganization.

Furthermore, some species-typical competencies are *not* present at birth but are no less foundational for normal functioning. Indeed, the timing of appearance of a cognitive competence need not bear any relation to how foundational it is. Therefore, developmental scientists' principal aim should not be to identify cognitive functions present at birth in order to declare them "core" and fundamental to all that emerge later. Rather, once particular functions are discovered, our job is to explicate the processes underlying their emergence in development *regardless of when they emerge*.

One entrenched, old idea is that there are two different processes responsible for phenotypes, one that relies on experiences and one that yields experience-independent "evolved behaviors." This idea is known as the "phylogeny fallacy" (Lickliter & Berry, 1990) simply because phenotype emergence in each generation is the product of development, and *all* development involves both genetic and experiential factors. To be sure, Spelke acknowledges that so-called "innate" cognitive processes must develop prenatally and that the experiences that give rise to them might involve activity generated in "subcortical or older cortical regions...that propagates to the plastic [neocortex]" (p. 195). But what brings about such prenatal activity and why should the resultant cognitive processes be considered any more "core" than earlier- or later-emerging processes? After all, the development of every higher cognitive process reflects neocortical cells' experiences with incoming stimulation. Consequently, a developmental analysis must *a priori* consider all stimulation as potentially crucial unless its role has been empirically ruled out. Such analyses do not need to distinguish between stimulation that arises outside of the neocortex but still within the brain (e.g., in a subcortical region) and stimulation that arises outside of the brain entirely (e.g., circulating hormones in a fetus's body, auditory stimuli that flood a fetus's brain with neurotransmitters, or a newborn's first whiff of its mother). Given this, Spelke's categorical distinction between different kinds of stimuli can be considered arbitrary and offers few insights into the emergence of cognitive skills.

At the heart of Spelke's theoretical argument are these statements: "the core systems ... have been shaped by hundreds of millions of years of cognitive evolution. Some core systems are shared by animals as remotely related to us as fish, and aspects of these systems are shared by flies and worms" (p. xx). For example, "the place system is innate. Cognitive and brain scientists have studied the mechanisms and processes by which place representations arise in infant minds, through research on animals who have been reared under systematically controlled conditions" (p. 139). These statements reflect at least three significant misunderstandings.

First, *every* cognitive system in extant organisms reflects millions of years of evolution and relies on subsystems that have survived natural selection (e.g., functioning neurons, sodium-potassium pumps, etc.). Phenotypes with long evolutionary histories should not be considered any more "core" – in the sense of "atomic" – than more recently evolved phenotypes that contribute to organisms' survival. *All* cognitive systems are built from smaller components that have long evolutionary histories.

Second, Spelke writes that core systems "have been shaped by...evolution," a claim that reflects a misunderstanding common among psychologists. This idea conflates the dynamics of populations, which can be influenced by natural selection, with the dynamics of individuals, which cannot (Witherington, Lickliter,



& Moore, 2023). Natural selection does not have the creative power to “shape” individuals’ phenotypes (Sober, 1984). Although natural selection can, over time, affect the distribution of phenotypes in a population, individuals’ phenotypes always reflect the *developmental* process that is of paramount concern to developmental scientists (Moore, 2008).

Finally, investigating how representations arise in *human* infant minds by doing experiments on nonhuman animals is fraught, because homology is not identity. Bird wings are homologous with human arms, but the bones in bird wings are different in key respects from the bones in our arms (Dumont, 2010). Thus, Spelke is on shaky ground when she claims that the human navigational system is “core” because other animals possess such a system at birth.

Spelke is correct that her “core” systems support children’s later learning, and her claim that it is advantageous to develop the capacity for abstraction early in life is insightful; abstraction need not await arrival at some developmental pinnacle. Still, there is no reason to argue that experiences are unnecessary as we develop the ability to abstract information. Although Spelke argues that infants “are predisposed to learn to ... categorize objects by using shape descriptions that capture the characteristic forms of plants and animals” (p. 202), there is strong experimental evidence that some of these forms are experience-dependent. For example, the basic categories of faces (e.g., same-race vs. other-race) and language (native vs. nonnative) are acquired through perceptual narrowing during infancy (Lewkowicz & Ghazanfar, 2009). Initially, these forms are so broadly specified that newborns do not distinguish human from nonhuman faces (Di Giorgio, Leo, Pascalis, & Simion, 2012) and integrate nonhuman faces with temporally synchronized tones (Lewkowicz, Leo, & Simion, 2010). These findings challenge Spelke’s view that infants are predisposed to recognize “people as social beings” (p. 301). Clearly, that level of cognitive specificity emerges after months of everyday postnatal experience.

Is it worth knowing that a particular ability is functional at birth? Absolutely. But that does not tell us (1) how the ability develops, (2) how important it is, (3) that it develops independently of contextual factors, or (4) that it is any more “core” than later-appearing abilities. Abilities that are important to normal human functioning appear at various times in development: The ability to hear and to suck is present in utero, the ability to navigate is present at birth, and the ability to walk is present only after about a year of postnatal development. There is little about appearing at birth that makes a characteristic special. Furthermore, it is not yet possible to do fair tests of human newborns’ competencies in most of Spelke’s “core” domains. Consequently, Spelke draws many of her conclusions from studies of older infants, toddlers, or young children. The problem is that all these studies reflect *postnatal* experiences and therefore do not convincingly demonstrate that these competencies are “innate” in humans.

When Spelke’s arguments invoke the commonalities we share with fish, she claims that it is “the same cognitive system, inherited from a distant common ancestor, [that] underlies their performance” (p. 129). However, vertebrates and other complex organisms inherit only developmental resources (including genes used in protein production), resources that we use to *build* our phenotypes anew in each generation via probabilistic developmental processes; we cannot inherit full blown phenotypes (Lewkowicz, 2011; Moore & Lickliter, 2023). In discussing the navigational abilities of fish, Spelke supports her arguments with the results of experiments that deny fish certain experiences.

These are the same sorts of “isolation” experiments that Lehrman (1953) warned can only reveal if withheld environmental factors are probably *not* directly involved in the development of a behavior; isolation experiments do not license the conclusion that the behavior under investigation is innate. Spelke’s conclusion that such experiments reveal innateness only stymies true developmental analysis.

Spelke believes “the core systems provide the foundations for the abstract concepts at the center of all our explicit knowledge” (p. xix). Perhaps. Nevertheless, this does not free developmental scientists from the responsibility of investigating the developmental emergence of such systems. Of course, it is not Spelke’s *personal* responsibility to do so. Nonetheless, it is one thing to choose a starting point and try to explain what develops from that point, but it is quite another to imply that the developmental processes responsible for bringing about that starting point do not warrant study. Spelke effectively does the latter when she labels these competencies “core.”

The renowned Swiss professor of anatomy and physiology Wilhelm His wrote in 1888:

The single word “heredity” cannot dispense science from the duty of making every possible inquiry into the mechanism of organic growth... . To think that heredity will build organic beings without mechanical means is a piece of unscientific mysticism... . A direct explanation ... [of the emergence of phenotypes] can only come from the immediate study of the different phases of individual development. (p. 295)

His highlighted the critical role of developmental analysis, regardless of whether the phenotype has a long or short evolutionary history and regardless of when in ontogeny the phenotype appears. Accordingly, we can reasonably ask nativists why we should explore how children accomplish tasks that confront them at 1 year of age (as Spelke will do in *How children learn*), but merely accept earlier-appearing competencies as somehow “core” and consequently not requiring developmental analysis. Scientists interested in development should resist the notion of “core” competencies because it short-circuits developmental analysis, leaving us ignorant of the factors that lead to the emergence of these competencies in the first place. Because knowledge of these factors could very well be important for understanding why adult-like competence in domains such as social communication fail to develop (as in autism), understanding their developmental origins should be a central goal of our discipline.

**Financial support.** This research received no specific grant from any funding agency, commercial, or not-for-profit sectors.

**Competing interest.** None.

## References

- Di Giorgio, E., Leo, I., Pascalis, O., & Simion, F. (2012). Is the face-perception system human-specific at birth? *Developmental Psychology*, 48(4), 1083–1090.
- Dumont, E. R. (2010). Bone density and the lightweight skeletons of birds. *Proceedings of the Royal Society B: Biological Sciences*, 277, 2193–2198.
- Gilbert, S. F., & Epel, D. (2015). *Ecological developmental biology: The environmental regulation of development, health, and evolution* (2nd ed.). Sinauer.
- Gottlieb, G. (1991). Experiential canalization of behavioral development: Theory. *Developmental Psychology*, 27, 4–13.
- Gottlieb, G. (2007). Probabilistic epigenesis. *Developmental Science*, 10(1), 1–11.
- His, W. (1888). On the principles of animal morphology. *Proceedings of the Royal Society of Edinburgh*, 15, 287–298.
- Lehrman, D. S. (1953). A critique of Konrad Lorenz’s theory of instinctive behavior. *The Quarterly Review of Biology*, 28, 337–363.

- Lewkowicz, D. J. (2011). The biological implausibility of the nature–nurture dichotomy and what it means for the study of infancy. *Infancy*, 16(4), 331–367.
- Lewkowicz, D. J., & Ghazanfar, A. A. (2009). The emergence of multisensory systems through perceptual narrowing. *Trends in Cognitive Sciences*, 13(11), 470–478. <https://doi.org/10.1016/j.tics.2009.08.004>
- Lewkowicz, D. J., Leo, I., & Simion, F. (2010). Intersensory perception at birth: Newborns match nonhuman primate faces and voices. *Infancy*, 15(1), 46–60.
- Lewontin, R. C. (2000). *The triple helix: Gene, organism, and environment*. Harvard University Press.
- Lickliter, R., & Berry, T. D. (1990). The phylogeny fallacy: Developmental psychology's misapplication of evolutionary theory. *Developmental Review*, 10, 348–364.
- Michel, G. F., & Moore, C. L. (1995). *Developmental psychobiology: An interdisciplinary science*. MIT.
- Moore, D. S. (2001). *The dependent gene: The fallacy of "nature vs. nurture"*. Henry Holt.
- Moore, D. S. (2008). Individuals and populations: How biology's theory and data have interfered with the integration of development and evolution. *New Ideas in Psychology*, 26, 370–386.
- Moore, D. S. (2013). Current thinking about nature and nurture. In K. Kampourakis (Ed.), *The philosophy of biology: A companion for educators* (pp. 629–652). Springer.
- Moore, D. S., & Lickliter, R. (2023). Development as explanation: Understanding phenotypic stability and variability after the failure of genetic determinism. *Progress in Biophysics and Molecular Biology*, 178, 72–77. <https://doi.org/10.1016/j.pbiomolbio.2023.01.003>
- Noble, R., & Noble, D. (2023). *Understanding living systems*. Cambridge University Press.
- Sober, E. (1984). *The nature of selection: Evolutionary theory in philosophical focus*. University of Chicago Press.
- Spelke, E. S. (2022). *What babies know: Core knowledge and composition*. Oxford University Press.
- Spemann, H., & Mangold, H. (1924/2001). Induction of embryonic primordia by implantation of organizers from a different species. *International Journal of Developmental Biology*, 45, 13–38.
- Strohman, R. C. (2003). Genetic determinism as a failing paradigm in biology and medicine: Implications for health and wellness. *Journal of Social Work Education*, 39, 169–191.
- Witherington, D. C., Lickliter, R., & Moore, D. S. (2023). What taking development seriously means for evolutionary theory. *Human Development*, 1–10. <https://doi.org/10.1159/000535723>

## The key to understanding core knowledge resides in the fetus

Vincent M. Reid\* 

School of Psychology, The University of Waikato, Hamilton, New Zealand  
[vincent.reid@waikato.ac.nz](mailto:vincent.reid@waikato.ac.nz)  
<https://profiles.waikato.ac.nz/vincent.reid/>

\*Corresponding author.

doi:10.1017/S0140525X23003175, e139

### Abstract

*What Babies Know* outlines a compelling case for why infancy research is fundamental for conceptualizing what it is to be human. There is another period in human development that is relatively inaccessible, yet is more important. In order to truly understand the nature of core knowledge, perception, and cognition, we must start not with the infant, but with the fetus.

In *What Babies Know*, Spelke (2022) uses the framework of early human development to uncover the perceptual and cognitive architecture of humanity. Not only is this a compelling argument. It is also very clearly the right position to take when confronted with the empirical evidence from across all domains of

psychology. Her very first aim is to explore what infants know at the time when their learning begins. At this initial hurdle, *What Babies Know* immediately becomes contentious. Learning does not begin during infancy. There is now a wealth of data to show that learning across several cognitive domains is present during the third trimester of pregnancy. When viewed through the lens of the fetus, *What Babies Know* offers some intriguing possibilities for the future of research during this earlier period of development.

Thankfully, the fetus has not been entirely ignored in *What Babies Know*. When the topic of language processing is broached, seminal work such as DeCaspar and Fifer (1980) is acknowledged and at the forefront of the chapter. What is surprising is that given the aims of the book, the fetus is mystifyingly absent when an in-depth exploration of this topic would enrich the arguments present throughout. Given that Spelke is aware of the fetus, how does the fetus interplay with the thesis that “experiments on infants provide the most direct access to the earliest emerging cognitive capacities at the foundations of our knowledge?” (Spelke, 2022, prologue xxi). It is the fetus, not the infant, that aligns with this statement. Given the carefully constructed positions that are developed throughout the book, there must be a rationale for the exclusion of evidence in multiple areas, despite fetal data having direct relevance to several topics in the volume.

The third trimester is so important for the development of perception and cognition that it has clear implications for a number of sections in *What Babies Know*. For example, number processing has been explored via neuroscience methods in the fetus (Schleger et al., 2014). Fetal vision is also advancing as a field, now that technical impediments are being overcome (e.g., Reid et al., 2017). In the domain of object processing, eye movements related to tracking an object in space, in this case light on the maternal abdomen, have also been observed via ultrasound (Donovan, Dunn, Penman, Young, & Reid, 2020). In the area of vision and object processing, Spelke does outline some of the seminal work on the biological basis of fetal development (e.g., as reviewed in Ackman & Crair, 2014) and provides suggestions on why the fetus may have object processing capacities and emerging spatial function. Work with the fetus nonetheless shines a clearer light on the parameters of other domains covered in the book. With speech detection and early learning, for example, the melody of a neonate cry is shaped by auditory experiences prior to birth (Mampe, Friederici, Christophe, & Wermke, 2009). There is a role for experience to shape the parameters of speech detection during the fetal period. There has been an exciting growth in our knowledge related to the perceptual and psychological characteristics of the fetus (Reid & Dunn, 2021), with recent models of sound processing (Vogelsang, Vogelsang, Diamond, & Sinah, 2023) demonstrating the impact of auditory processing on the spoken word during pregnancy. Clearly the fetus has much to offer when considering the themes that are outlined by Spelke.

*What Babies Know* also maps the evolution of thinking surrounding core knowledge since initial proposals several decades ago through to the present day. Within this framework, the changing definitions of the term *innate* are of particular interest. The current position, namely “they are present and functional on an infant’s first informative encounters with the entities that they serve to represent,” (Spelke, 2022, prologue xix) actively incorporates aspects of constructivist viewpoints. Through so doing, Spelke clearly seeks to remove fuel from the fire of the historically heated debate on the role of experience in shaping development (e.g., Spelke & Kinzler, 2009; Spencer et al., 2009).

If encountering information prior to when it is required is an essential aspect of cognitive development, then under this definition of innate, the nature of acquisition is incorporated into the framework of innate processes. With this premise, learning and the seeking out of specific forms of information to refine neural development are the natural unfolding of genetic codes. This approach effectively signals an acknowledgement that learning is inherent in human development and much of the neural systems that underpin cognitive development are not hard wired in genetics. Rather, it is the pathway to learning itself that is now defined in nativist terms. Such an approach is capable of encompassing everything from the formation of initial visual experience via retinal waves occurring before a fully functional visual system (e.g., Firth, Wang, & Feller, 2005) which is beautifully described by Spelke, through to the proposals that teaching is an inherent part of being human (Csibra & Gergely, 2011). There has been a significant accumulation of evidence in recent decades that supports a neuroconstructivist approach to development. For example, computational modeling on the development of processing the past tense (Westermann & Ruh, 2012) presents a strong argument for this position. The approach by Karmiloff-Smith (2009), when exploring developmental disorders is equally compelling. Spelke's modification to the innate position will undoubtedly produce much thought and dialogue in the field.

From the perspective of fetal research, it is important to recognize that the methods outlined by Spelke for application with infant populations are almost certainly those best suited for research with the fetus. Fetal research was once almost exclusively the domain of observation (e.g., Zoia et al., 2007) or correlations with aspects of fetal development with later behavioral outcomes, such as prenatal head circumference and later educational achievement (e.g., DiPietro, Costigan, & Voegtline, 2015). Due to the advent of noninvasive brain imaging and high-quality ultrasound, experimental approaches are increasingly common with the fetus. The paradigms utilized are exclusively derived from infancy research (e.g., Johnson & Morton, 1991) or from the neurosciences (e.g., oddball paradigms, for a review, see Dunn, Reissland, & Reid, 2015). Fetal research is by no means a fully established field. There are currently no mapped parameters related to fetal attention to auditory, tactile, or visual stimulation. There is little point in utilizing a paradigm where your sample is unable to process the stimuli due to fatigue or inattention. The consequence of this is that current experiments are short in nature and typically feature only two conditions. What is also unknown are the specific infancy paradigms that are best suited to the fetal population. Spelke highlights successful and illuminating experiments with the infant that utilize peripheral stimulus presentation, paired preference, sequential presentations, and classic habituation amongst even more paradigms. The volume reads like a shopping list of experimental methods that are yet to be employed and fully understood with a fetal sample.

The arguments presented in *What Babies Know* raises a number of important questions related to the fetal period. Can motion be discriminated during the third trimester? If so, will there be preferential engagement with agentive motion as is the case with the newborn (e.g., Di Giorgio, Lunghi, Simion, & Vallortigara, 2017)? It is tempting to determine that fetal processing of information must derive from genetic predispositions. Such a conclusion does not account for fetal experiences that must be considered. For example, it is possible that self-produced actions (Myowa-Yamakoshi & Takeshita, 2006; Wilkinson, Paikan, Gredebäck, Rea, & Metta, 2014) and the visual

capacity to monitor them (Del Giudice, 2011) could lay the foundation that drives early psychological systems, including a preference for agentive motion. It is also feasible that contingently responsive maternal behavior produces environmental change that is detected as causal by the fetus. Supporting this notion, prior work has shown that fetal heart rate is modified based on maternal swinging movements (Lecanuet & Jacquet, 2002). Based on these studies, an experimental framework for social-cognitive development is possible, even during the fetal period. From this specific topic alone, it is clear that in order to understand the issue of core knowledge, we must start not with the infant, but with the fetus. Spelke is unquestionably one of the most preeminent scientists of the past century. It is my hope that the equivalent to Spelke born in the 2000s will, after many decades of research, produce a volume entitled *What the Fetus Knows*.

**Financial support.** This work was supported by Te Apārangi, The Royal Society of New Zealand, Marsden grant 20-UOW-22.

**Competing interest.** None.

## References

- Ackman, J. B., & Crair, M. C. (2014). Role of emergent neural activity in visual map development. *Current Opinion in Neurobiology*, 24, 166–175. doi: [10.1016/j.conb.2013.11.011](https://doi.org/10.1016/j.conb.2013.11.011)
- Csibra, G., & Gergely, G. (2011). Natural pedagogy as evolutionary adaptation. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 366, 1149–1157. doi: [10.1098/rstb.2010.0319](https://doi.org/10.1098/rstb.2010.0319)
- DeCaspar, A. J., & Fifer, W. P. (1980). Of human bonding: Newborns prefer their mothers' voices. *Science*, 208(448), 1174–1176. doi: [10.1126/science.7375928](https://doi.org/10.1126/science.7375928)
- Del Giudice, M. (2011). Alone in the dark? Modeling the conditions for visual experience in human fetuses. *Developmental Psychobiology*, 53, 214–219. doi: [10.1002/dev.20506](https://doi.org/10.1002/dev.20506)
- Di Giorgio, E., Lunghi, M., Simion, F., & Vallortigara, G. (2017). Visual cues of motion that trigger animacy perception at birth: The case of self-propulsion. *Developmental Science*, 20, e12394. doi: [10.1111/desc.12394](https://doi.org/10.1111/desc.12394)
- DiPietro, J. A., Costigan, K. A., & Voegtline, K. M. (2015). Studies in fetal behavior: Revisited, renewed and reimagined. *Monographs of the Society for Research in Child Development*, 80(3). doi: [10.1111/mono.v80.3](https://doi.org/10.1111/mono.v80.3)
- Donovan, T., Dunn, K., Penman, A., Young, R. J., & Reid, V. M. (2020). Fetal eye movements in response to a visual stimulus. *Brain and Behavior*, 10(8), e01676. <https://doi.org/10.1002/brb3.1676>
- Dunn, K., Reissland, N. N., & Reid, V. M. (2015). The functional foetal brain: A systematic preview of methodological factors in reporting foetal visual and auditory capacity. *Developmental Cognitive Neuroscience*, 13, 43–52. doi: [10.1016/j.dcn.2015.04.002](https://doi.org/10.1016/j.dcn.2015.04.002)
- Firth, S. L., Wang, C.-T., & Feller, M. B. (2005). Retinal waves: Mechanisms and function in visual system development. *Cell Calcium*, 37(5), 425–432. doi: [10.1016/j.ceca.2005.01.010](https://doi.org/10.1016/j.ceca.2005.01.010)
- Johnson, M. H., & Morton, J. (1991). *Biology and cognitive development: The case of face recognition*. Blackwell.
- Karmiloff-Smith, A. (2009). Nativism versus neuroconstructivism: Rethinking the study of developmental disorders. *Developmental Psychology*, 45(1), 56–63. doi: [10.1037/a0014506](https://doi.org/10.1037/a0014506)
- Lecanuet, J.-P., & Jacquet, A.-Y. (2002). Fetal responsiveness to maternal passive swinging in low heart rate variability state: Effects of stimulation direction and duration. *Developmental Psychobiology*, 40(1), 57–67. doi: [10.1002/dev.10013](https://doi.org/10.1002/dev.10013)
- Mampe, B., Friederici, A. D., Christophe, A., & Wermke, K. (2009). Newborns' cry melody is shaped by their native language. *Current Biology*, 19, 1994–1997. doi: [10.1016/j.cub.2009.09.064](https://doi.org/10.1016/j.cub.2009.09.064)
- Myowa-Yamakoshi, M., & Takeshita, H. (2006). Do human fetuses anticipate self-oriented actions? A study by four-dimensional (4D) ultrasonography. *Infancy*, 10(3), 289–301. [https://doi.org/10.1207/s15327078in1003\\_5](https://doi.org/10.1207/s15327078in1003_5)
- Reid, V. M., & Dunn, K. (2021). The fetal origins of human psychological development. *Current Directions in Psychological Science*, 30(2), 144–150. <https://doi.org/10.1177/0963721420984419>
- Reid, V. M., Dunn, K., Young, R. J., Amu, J., Donovan, T., & Reissland, N. (2017). The human fetus preferentially engages with face-like stimuli. *Current Biology*, 27(12), 1825–1828. doi: [10.1016/j.cub.2017.05.044](https://doi.org/10.1016/j.cub.2017.05.044)
- Schleger, F., Landerl, K., Muenssinger, J., Draganova, R., Reinl, M., Kiefer-Schmidt, I., ... Preissl, H. (2014). Magnetoencephalographic signatures of numerosity discrimination in fetuses and neonates. *Developmental Neuropsychology*, 39(4), 316–329. doi: [10.1080/87565641.2014.914212](https://doi.org/10.1080/87565641.2014.914212)



- Spelke, E. (2022). *What babies know*. Core Knowledge and Composition Volume 1. Oxford University Press.
- Spelke, E. S., & Kinzler, K. D. (2009). Innateness, learning, and rationality. *Child Development Perspectives*, 3(2), 96–98. doi: [10.1111/j.1750-8606.2009.00085.x](https://doi.org/10.1111/j.1750-8606.2009.00085.x)
- Spencer, J. P., Samuelson, L. K., Blumberg, M. S., McMurray, B., Robinson, S. R., & Tomblin, J. B. (2009). Short arms and talking eggs: Why we should no longer abide the nativist-empiricist debate. *Child Development Perspectives*, 3(2), 79–87. doi: [10.1111/j.1750-8606.2009.00081.x](https://doi.org/10.1111/j.1750-8606.2009.00081.x)
- Vogelsang, M., Vogelsang, L., Diamond, S., & Sinah, P. (2023). Prenatal auditory experience and its sequelae. *Developmental Science*, 26, e13278. doi: [10.1111/desc.13278](https://doi.org/10.1111/desc.13278)
- Westermann, G., & Ruh, N. (2012). A neuroconstructivist model of past tense development and processing. *Psychological Review*, 119, 649–667. doi: [10.1037/a0028258](https://doi.org/10.1037/a0028258)
- Wilkinson, N., Paikan, A., Gredebäck, G., Rea, F., & Metta, G. (2014). Staring us in the face? An embodied theory of innate face preference. *Developmental Science*, 17(6), 809–825. doi: [10.1111/desc.12159](https://doi.org/10.1111/desc.12159)
- Zoia, S., Blason, L., D'Ottavio, G., Bulgheroni, M., Pezzetta, E., Scabar, A., & Castello, U. (2007). Evidence of early development of action planning in the human foetus: A kinematic study. *Experimental Brain Research*, 176, 217–226. doi: [10.1007/s00221-006-0607-3](https://doi.org/10.1007/s00221-006-0607-3)

## Perceptual (roots of) core knowledge

Brian J. Scholl\* 

Department of Psychology, Yale University, New Haven, CT, USA  
[brian.scholl@yale.edu](mailto:brian.scholl@yale.edu)  
<https://perception.yale.edu/>

\*Corresponding author.

doi:10.1017/S0140525X23003023, e140

### Abstract

Some core knowledge may be rooted in – or even *identical to* – well-characterized mechanisms of mid-level visual perception and attention. In the decades since it was first proposed, this possibility has inspired (and has been supported by) several discoveries in both infant cognition and adult perception, but it also faces several challenges. To what degree does *What Babies Know* reflect how babies *see* and *attend*?

### Introduction: What babies see?

As the various subfields of cognitive science have become ever more distinct and specialized, the notion of *core knowledge* has acted as a sort of intellectual glue – synergizing research from its intellectual origins in developmental psychology, to studies of animal cognition, adult visual perception, linguistic representation, computational modeling and AI, and beyond. Here I focus on one particular form of synergy, between *What Babies Know* (Elizabeth Spelke's brilliant and groundbreaking book summarizing one of the most productive research programs in all of science; Spelke, 2022; henceforth *WBK*) and the study of what and how we *see* (as explored in studies of adult visual perception and attention).

What kinds of mental representations and processes characterize core knowledge? Once upon a time, the answer was unambiguous: *Higher-level thought*. As Spelke once suggested, “Humans come to know about an object's unity, boundaries, and persistence in ways like those by which we come to know about its material composition or its market value” (Spelke, 1988, p. 198). This

view was inspired by a (now-obsolete) characterization of perception as relatively unsophisticated. Chapter 2 of *WBK*, for example, involves discrete *objects*, but as Spelke once suggested: “Perceptual systems do not package the world into units.... The parsing of the world into things may point to the essence of thought and to its essential distinction from perception” (1988, p. 229). And since infants continue to represent objects that are not currently in view, the responsible mechanisms must therefore “carry infants beyond the world of immediate perception” (p. 172). By the 1990s, however, advances in the study of adult perception had made it clear that visual processing does in fact “package the world into units” on its own, independent of higher-level thought – into representations of both surfaces (for an early review, see Nakayama, He, & Shimojo, 1995) and objects (for an early review, see Scholl, 2001), which then persist through time, occlusion, and featural change (for a review, see Scholl, 2007).

These discoveries led to a proposal, first articulated in the late 1990s and early 2000s (Leslie, Xu, Tremoulet, & Scholl, 1998; Scholl, 2001, Section 7.2; Scholl & Leslie, 1999; see also Carey & Xu, 2001) that at least some types of core knowledge may be rooted in the mechanisms and representations of mid-level visual processing and object-based attention (henceforth *mid-level vision*). Infants may have expectations about the behaviors of objects not because of considered deliberation or conceptual theories, but because that is simply how they experience the world in the first place, in terms of their brute visual percepts. “[S]urprising parallels between recent results in cognitive developmental psychology and the study of object-based visuospatial attention suggest that the two areas of inquiry may have something to do with each other” (Scholl & Leslie, 1999, p. 60) – and although “visual processing in adults may seem relatively unrelated to the study of core knowledge in infant cognition, ... recent work has suggested that these two seemingly different fields may in fact be studying the same underlying representations and constraints” (Strickland & Scholl, 2015, p. 571).

### Progress: Sophisticated seeing!

The ultimate value of any theoretical proposal lies in the concrete progress it inspires. How has the proposal that core knowledge is rooted in mid-level vision fared in the decades since it was first introduced? Here are three examples of how this view has fueled new discoveries in both domains:

**Cohesion and persistence:** Many core knowledge principles apply to objects but not non-solid substances, and early work showed how cohesion violations (failures to maintain rigid boundaries and internal connectedness) frustrate infants' object tracking (e.g., Huntley-Fenner, Carey, & Solimando, 2002; Spelke & Van de Walle, 1993). This led directly to the discoveries in adult vision that cohesion violations also frustrate attentional tracking (vanMarle & Scholl, 2003) and the maintenance of object-file representations – even when just viewing a single object split into two (Mitroff, Scholl, & Wynn, 2004). And this adult vision work then directly inspired the demonstration that even a single object (e.g., a cracker) splitting into two destroys infants' ability to track quantity (Cherries, Mitroff, Wynn, & Scholl, 2008).

**Attentional prioritization:** Categorizing a stimulus into a particular “event type” (such as *occlusion* or *containment*) biases infants to remember features that are especially diagnostic for that type (such as the width of an object, with a vertical container; e.g., Hespos & Baillargeon, 2001; Wang, Baillargeon, & Brueckner, 2004). This led directly to the discovery that such

prioritization also occurs spontaneously in adults' visual working memory: While viewing dynamic containment (but not occlusion) events, change detection is better for those changes in that affect whether objects will "fit" (Strickland & Scholl, 2015) – and the subtle details of this were then subsequently also seen in infants' object tracking (Goldman & Wang, 2019).

*Seeing agency:* Chapter 7 of *WBK* reviews many studies showing how infants automatically treat certain motion patterns (e.g., involving pursuit) as cues to agency and intentionality – and how they expect agents to behave rationally, for example by following direct paths (e.g., Gergeley, Nadasdy, Csibra, & Biro, 1995; Southgate & Csibra, 2009). This led directly to the discovery that adults' mid-level vision also spontaneously (and even irresistibly) extracts properties such as agency and goal-directedness when viewing "chasing" displays (Gao, Newman, & Scholl, 2009, 2010; van Buren, Uddenberg, & Scholl, 2016) – and that violations of rational action similarly destroy adults' ability to spontaneously see chasing (Gao & Scholl, 2011).

These examples demonstrate how taking connections between infant cognition and adult perception seriously can drive empirical progress – showing how these two domains employ similar representations (e.g., of agency), are constrained by similar principles (e.g., of cohesion), and have similar downstream consequences (e.g., of orienting attention). At the least, such parallels suggest that one domain may help to fuel the other – that core knowledge may be rooted in, and partially grow out of, mid-level vision. At the most extreme, such connections suggest that these two domains could be one and the same.

### Challenges: Prosociality, language, and beyond

The essence of the progress reviewed above is a striking *match* between the results of experiments in infant cognition and adult perception. And such matches may go far beyond these three case studies (Bai, 2023), extending even into the nuances and mechanics of habituation itself (Turk-Browne, Scholl, & Chun, 2008). But just how close is this match? The biggest challenges to the view sketched above may lie in cases where the match is imperfect, in either direction.

The suggestion that core knowledge in infancy transcends mid-level vision in adults seems especially salient for at least two domains, each of which is the focus of a key chapter of *WBK*. First, as reviewed in Chapter 8, young infants may already have expectations and preferences related to prosociality – as when they observe one shape help (or hinder) another shape from climbing an incline (Hamlin, Wynn, & Bloom, 2007). But no work has yet suggested that visual processing itself directly extracts representations of helping, hindering, or prosociality in general (even though properties such as [im]morality in certain visual scenes may be correlated with lower-level cues; De Freitas & Alvarez, 2018). Second, as reviewed in Chapter 9, several aspects of core knowledge seem intimately related to *language*. Even infants' object tracking, for example, can depend on how people linguistically refer to the objects (Dewar & Xu, 2007; Xu, 2002). But mid-level vision seems largely encapsulated from linguistic processing, and vice versa (Firestone & Scholl, 2016) – and so if core knowledge reflects the operation of mid-level vision, then such linguistic connections may be rendered mysterious or inexplicable.

Potential mismatches may also loom large in the other direction – when adults' mid-level vision seems *more* sophisticated than infants' core knowledge. The studies reviewed in the previous

section are all examples in which visual representations have been found to be especially *sophisticated* – encompassing properties and constraints (such as agency and cohesion) more closely associated with higher-level thought. But this trend in vision research goes far beyond the classical domains of core knowledge. Additional work, for example, has suggested that mid-level vision automatically and spontaneously extracts representations of *causal history* (i.e., of how objects came to look the way that they do; Chen & Scholl, 2016), *soft-material intuitive physics* (e.g., inferring the shape of objects under cloths; Wong, Bi, Soltani, Yildirim, & Scholl, 2023), and even *unfinishedness* (as when an object appears not to have ended its movement; Ongchoco, Wong, & Scholl, 2023). But competence involving such seemingly sophisticated domains is nowhere to be found in most characterizations of core knowledge.

On one hand, some of these challenges could be dissolved with further research. After all, when the current proposal was first articulated in the late 1990s, nobody yet suspected that mid-level vision might match infant cognition in the ways reviewed in the previous section. And so we might still discover that mid-level vision extracts representations of prosociality, or that infant core knowledge also encompasses representations of causal history. On the other hand, some of these challenges remain despite having been recognized long ago (e.g., Scholl & Leslie, 1999, Section 5.5) – and without a principled way to demarcate which results we expect to "match" and which we do not (e.g., only spatiotemporal processing, but not contact-mechanical processing; Cheries, Mitroff, Wynn, & Scholl, 2009; Scholl & Leslie, 1999), these challenges continue to be acute.

### The state of the art

The view of core knowledge sketched here contrasts in some ways with that from *WBK*. On one hand, Spelke notes that "I believe there is truth to this view" that "the object representations [involved in core knowledge] are the products of perceptual processes" (p. 78) – and throughout the book she masterfully reviews relevant work on adults' mid-level vision (including much of the work discussed here). She also notes that her views on these issues have changed over time: "I once proposed, wrongly, that objects are not grasped by a perceptual system but by ... a system of central cognition.... Research ... provided decisive evidence against this proposal...: Adults were found to share the representational system found in infants" (precis, sect. 1). As such, the view sketched here is meant as more of a friendly extension than a criticism – perhaps just placing a sharper focus on certain themes from *WBK*.

On the other hand, *WBK* also provides several additional arguments against this view, which seem less compelling. Spelke suggests that some aspects of core knowledge cannot have perceptual origins because (a) core knowledge representations are *abstract* (p. xxi) – but many aspects of mid-level vision also abstract over many surface variables (Scholl, 2007); (b) perception involves "detectable surfaces" rather than "the entities that those surfaces belong to" (p. 198) – but recent work in mid-level vision argues for exactly the opposite view (Wong et al., 2023); and (c) core knowledge representations have a time-course that transcends momentary perception (p. 78) – but at least some object-file representations in mid-level vision have been shown to persist throughout interruptions for at least 8 seconds, and possibly much longer (a result that was also explicitly motivated by connections to infant cognition; Noles, Scholl, & Mitroff, 2005).

Some of the most recent discussions of core knowledge also still seem to veer away from the possibility of substantive connections to mid-level vision in other ways. In *WBK*, for example, Spelke champions the notion of the infant mind as implementing a type of “physics engine” (Ullman, Spelke, Battaglia, & Tenenbaum, 2017) – but such frameworks do not necessarily abide by the constraints of encapsulated mid-level vision, as they also readily accommodate higher-level knowledge (e.g., about how the colors of blocks may arbitrarily signal their masses; Battaglia, Hamrick, & Tenenbaum, 2013). As a result, the physics-engine framework may be true and important – but also simply orthogonal to the distinction between mid-level vision and higher-level thought. And whereas *WBK* ultimately characterizes core knowledge (somewhat ambiguously) as occupying “a middle ground between perceptual systems and belief systems” (p. xxi), I hope that we might also continue to take seriously the more direct possibility that core knowledge is (rooted in) mid-level vision – and that *What Babies Know* may largely reflect how babies *see* and *attend*.

**Financial support.** This research received no specific grant from any funding agency, commercial, or not-for-profit sectors.

**Competing interest.** None.

## References

- Bai, D. (2023). Intuitive physics in visual perception. PhD dissertation, l'École Normale Supérieure.
- Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences*, *110*, 18327–18332.
- Carey, S., & Xu, F. (2001). Infants knowledge of objects: Beyond object-files and object tracking. *Cognition*, *80*, 179–213.
- Chen, Y.-C., & Scholl, B. J. (2016). The perception of history: Seeing causal history in static shapes induces illusory motion perception. *Psychological Science*, *27*, 923–930.
- Cherries, E. W., Mitroff, S. R., Wynn, K., & Scholl, B. J. (2008). Cohesion as a constraint on object persistence in infancy. *Developmental Science*, *11*, 427–432.
- Cherries, E. W., Mitroff, S. R., Wynn, K., & Scholl, B. J. (2009). Do the same principles constrain persisting object representations in infant cognition and adult perception?: The cases of continuity and cohesion. In B. Hood & L. Santos (Eds.), *The origins of object knowledge* (pp. 107–134). Oxford University Press.
- De Freitas, J., & Alvarez, G. (2018). Your visual system provides all the information you need to make moral judgments about generic visual events. *Cognition*, *178*, 133–146.
- Dewar, K., & Xu, F. (2007). Do 9-month-old infants expect distinct words to refer to kinds? *Developmental Psychology*, *43*, 1227–1238.
- Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for “top-down” effects. *Behavioral and Brain Sciences*, *39*(E229), 1–77.
- Gao, T., McCarthy, G., & Scholl, B. J. (2010). The wolfpack effect: Perception of animacy irresistibly influences interactive behavior. *Psychological Science*, *21*, 1845–1853.
- Gao, T., Newman, G. E., & Scholl, B. J. (2009). The psychophysics of chasing: A case study in the perception of animacy. *Cognitive Psychology*, *59*, 154–179.
- Gao, T., & Scholl, B. J. (2011). Chasing vs. stalking: Interrupting the perception of animacy. *Journal of Experimental Psychology: Human Perception & Performance*, *37*, 669–684.
- Gergeley, G., Nadasdy, Z., Csibra, G., & Biro, S. (1995). Taking the intentional stance at 12 months of age. *Cognition*, *56*, 165–193.
- Goldman, E. J., & Wang, S.-h. (2019). Comparison facilitates the use of height information by 5-month-olds in containment events. *Developmental Psychology*, *55*, 2475–2482.
- Hamlin, J. K., Wynn, K., & Bloom, P. (2007). Social evaluation by preverbal infants. *Nature*, *450*, 557–559.
- Hespos, S. J., & Baillargeon, R. (2001). Infants’ knowledge about occlusion and containment events: A surprising discrepancy. *Psychological Science*, *12*, 141–147.
- Huntley-Fenner, G., Carey, S., & Solimando, A. (2002). Objects are individuals but stuff doesn’t count: Perceived rigidity and cohesiveness influence infants’ representations of small numbers of discrete entities. *Cognition*, *85*, 203–221.
- Leslie, A. M., Xu, F., Tremoulet, P., & Scholl, B. J. (1998). Indexing and the object concept: Developing “what” and “where” systems. *Trends in Cognitive Sciences*, *2*, 10–18.
- Mitroff, S. R., Scholl, B. J., & Wynn, K. (2004). Divide and conquer: How object files adapt when a persisting object splits into two. *Psychological Science*, *15*, 420–425.
- Nakayama, K., He, Z. J., & Shimojo, S. (1995). Visual surface representation: A critical link between lower-level and higher-level vision. In S. Kosslyn (Ed.), *Visual cognition* (pp. 1–70). Volume 2 of *An Invitation to Cognitive Science*, 2nd ed. MIT Press.
- Noles, N. S., Scholl, B. J., & Mitroff, S. R. (2005). The persistence of object file representations. *Perception & Psychophysics*, *67*, 324–334.
- Ongchoco, J. D. K., Wong, K., & Scholl, B. J. (2023). The “unfinishedness” of dynamic events is spontaneously extracted in visual processing: A new “Visual Zeigarnik Effect”. *Journal of Vision*, *23*, 4974.
- Scholl, B. J. (2001). Objects and attention: The state of the art. *Cognition*, *80*, 1–46.
- Scholl, B. J. (2007). Object persistence in philosophy and psychology. *Mind & Language*, *22*, 563–591.
- Scholl, B. J., & Leslie, A. M. (1999). Explaining the infant’s object concept: Beyond the perception/cognition dichotomy. In E. Lepore & Z. Pylyshyn (Eds.), *What is cognitive science?* (pp. 26–73). Blackwell.
- Southgate, V., & Csibra, G. (2009). Inferring the outcome of an ongoing novel action at 13 months. *Developmental Psychology*, *45*, 1794–1798.
- Spelke, E. (1988). Where perceiving ends and thinking begins: The apprehension of objects in infancy. In A. Yonas (Ed.), *Perceptual development in infancy* (pp. 197–234). Erlbaum.
- Spelke, E. (2022). *What babies know*. Oxford University Press.
- Spelke, E., & Van de Walle, G. (1993). Perceiving and reasoning about objects: Insights from infants. In N. Eilan, R. McCarthy & W. Brewer (Eds.), *Spatial representation* (pp. 132–161). Basil Blackwell.
- Strickland, B., & Scholl, B. J. (2015). Visual perception involves “event type” representations: The case of containment vs. occlusion. *Journal of Experimental Psychology: General*, *144*, 570–580.
- Turk-Browne, N. B., Scholl, B. J., & Chun, M. M. (2008). Babies and brains: Habituation in infant cognition and functional neuroimaging. *Frontiers in Human Neuroscience*, *2*, Article 16.
- Ullman, T. D., Spelke, E. S., Battaglia, P., & Tenenbaum, J. B. (2017). Mind games: Game engines as an architecture for intuitive physics. *Trends in Cognitive Sciences*, *21*, 649–665.
- van Buren, B., Uddenberg, S., & Scholl, B. J. (2016). The automaticity of perceiving animacy: Goal-directed motion in simple shapes influences visuomotor behavior even when task-irrelevant. *Psychonomic Bulletin & Review*, *23*, 797–802.
- vanMarle, K., & Scholl, B. J. (2003). Attentive tracking of objects versus substances. *Psychological Science*, *14*, 498–504.
- Wang, S., Baillargeon, R., & Brueckner, L. (2004). Young infants’ reasoning about hidden objects: Evidence from violation-of-expectation tasks with test trials only. *Cognition*, *93*, 167–198.
- Wong, K. W., Bi, W., Soltani, A., Yildirim, I., & Scholl, B. J. (2023). Seeing soft materials draped over objects: A case study of intuitive physics in perception, attention, and memory. *Psychological Science*, *34*, 111–119.
- Xu, F. (2002). The role of language in acquiring object kind concepts in infancy. *Cognition*, *85*, 223–250.

## Core knowledge and its role in explaining uniquely human cognition: Some questions

Armin W. Schulz\* 

Department of Philosophy, University of Kansas, Lawrence, KS, USA  
 awschulz@ku.edu  
<http://people.ku.edu/~a382s825/>

\*Corresponding author.

doi:10.1017/S0140525X23003199, e141

### Abstract

Questions can be raised about the central status that evolutionarily ancient core knowledge systems are given in Spelke’s otherwise very compelling theory. So, the existence of domain-general learning capacities has to be admitted, too, and no clear reason is provided to doubt the existence of uniquely human cognitive adaptations. All of these factors should be acknowledged when explaining human thought.



*What Babies Know* (Spelke, 2022) is a big book – some 532 pages. As is to be expected from one of the founders of cognitive science, it is uniformly well-written, thoroughly well-researched, and chock-full of fascinating, original, and compelling arguments. It is also only half of Spelke's theory of human cognition: *What Babies Know* sets out the theoretical foundations for the next volume *What Babies Learn*, and together, the two books explain how humans come to be the kinds of thinkers that they are. Because the latter book is not part of this symposium, assessing the overall picture that emerges from Spelke's masterful synthesis of decades worth of work on the peculiarities of human cognition is not trivial. Still, the very fact of this division of her overall framework is meaningful – and raises some questions. To get at these questions, though, a brief background of the cornerstones of Spelke's theory of core knowledge is necessary.

According to this theory, human cognition is built on the foundation of a number of “core knowledge systems.” These are evolutionarily ancient, domain-specific sets of mental representations that structure how we interpret the world we live in. For example, there is a core knowledge system centering on *objects*: Beyond just carving the world into edges and corners (as on Marrian theories of vision), we are also born expecting the world to contain *objects* that have various features, such as the fact that they exclude each other (they cannot both occupy the same part of space). Further core knowledge systems concern the domains of *space*, *number*, *form* (roughly, what objects are plausible biological entities), *agency*, and *sociality*. At heart of Spelke's theory then is the claim that these core knowledge systems make for the basis of human cognition: They structure our thoughts, and as their outputs get further integrated through our linguistic capacities – another core knowledge system – they build up the kinds of minds we have. For present purposes, two aspects of this theory are crucial. (Spelke's theory is built on several decades' worth of empirical studies, most of which she was instrumental in conducting. Although the results of these studies can be interpreted in different ways – a point often noted explicitly in the book – I shall leave this aside here.)

First, it is an explicitly nativist picture of human cognition, and thus contrasts with, say, the account of Heyes (2018), according to which human cognition is the result of a handful of general purpose learning tools that are then culturally harnessed to yield human knowledge structures. As a nativist account, Spelke's picture has to grapple with exactly what it means for something – like a knowledge system – to be innate. To address this question, Spelke employs the following characterization of innateness: a “cognitive system is innate if it is *not learned*: that is, if it is present and functional on the infant's first effective perceptual encounters with the entities to which it applies” (p. 71).

As an account of innateness, though, the part behind the colon here is a bit tricky. If kids *culturally learn* about something – for example, that big snakes are dangerous – then this knowledge can be “present and functional on the infant's first effective perceptual encounters with the entities to which it applies.” However, we would not want to see this as a case of innateness: It is learned, after all (as also pointed out by the first part of the above characterization). On the flipside of this, if an infant encounters something – such as agents acting on false beliefs – several times, but if these encounters do not result in the infant *learning* about the entity, then a later capacity for dealing with the entity could still be a case of innateness, even though this capacity is not “present and functional on the infant's first effective perceptual encounters with the entities to which it applies.”

It is thus really the first part of the sentence that does the work here: The issue is that core knowledge systems are *not* individually, culturally, or socially learned. Here it is noteworthy that Spelke of course accepts that such learning does occur – we are not born knowing how to play piano, say, or how to read music. A key tenet of her theory is just that many of our major representational expectations are *not* learned. This will become important again below.

This leads to the second and most important point to note here: The fact that the focus of Spelke's explanatory framework is squarely on the ancient, innate knowledge structures that are shared with many nonhuman animals. The question I want to raise in this commentary is whether this is the best way to get at distinctively human cognition. Put differently: Even if we accept that human minds contain core knowledge systems, which have evolved because they prepared many organisms (including humans) for successful interactions with the world, the question remains of what makes these systems so special that they should be taken to be *fulcrum* on which human cognition rests.

To make this question clearer, it is useful to note that, in the précis of the book, there is a somewhat misleading statement about the evolutionary presuppositions of core knowledge systems. So, in the précis, it is said: “An ancient system that first emerged in highly distant ancestors is likely to center on abstract content, because it had to be applicable to the diverse environments that the descendants of that last common ancestor came to inhabit.” This is misleading, as it seems to imply that processes of biological evolution have foresight about which traits will be adaptive in the future – which is not the case. A psychological system cannot be selected for conferring fitness benefits to the descendants of the bearer of the system (in whatever environment they live): Somehow, these benefits have to come back to the bearer itself. In the book itself, Spelke is clearer in noting that the issue is that psychological traits that are consistently adaptive over many generations and in different environments are more likely to go into fixation in the population than traits that are not adaptive in these ways. According to Spelke, core knowledge systems are such traits. Still, there is an important point buried here.

On the one hand, many of the kinds of cognitive traits that are consistently adaptive over many generations and in different environments would seem to be *extremely* “abstract.” Indeed, some key examples of such traits would seem to be powerful abilities for individual and cultural learning. This is because of the fact that former can be used in many different environments and circumstances to acquire locally adaptive ways of acting: How to be socially successful, say, or which foods are edible in which ways. As a matter of fact, this is also precisely the reasoning underlying the common accounts of the evolution of cultural learning (see, e.g., Boyd & Richerson, 2005; Henrich, 2015; Heyes, 2018). By its own logic, Spelke's theory thus needs to acknowledge the existence of these abilities for learning.

On the flipside of this, it is also plausible that there are psychological traits that have been specifically adaptive in the human lineage. Now, Spelke appears reluctant to admit this, but it is unclear exactly why that is. For example, in response to Gergely and Csibra's idea (see, e.g., their Csibra & Gergely, 2011) that humans, specifically, have evolved to be teachers and learners, Spelke writes (p. 428): “Good teachers cannot efficiently fill the gaps in students' knowledge unless they are aware of what their students do and do not understand” – that is, unless they rely on a core knowledge system of other agents. This, though, seems false:

We could be born with certain expectations of what is important to teach; this could still be adaptive even if it is not fully accurate. For example, even if teachers unnecessarily belabor points, teaching and learning can be adaptive if there are enough morsels of insight in the teaching.

Something similar goes for a uniquely human talent for symbol cognition. In response, Spelke writes (pp. 432–433): “A further reason for doubting that a species-unique talent for learning and using spatial symbols underlies our cognitive accomplishments, including our construction of social agent concepts, comes from studies of symbol learning in nonhuman animals. Some chimpanzees, monkeys, and parrots have been trained to use pictures or objects as symbols since infancy, and they developed impressive abilities both to communicate by means of those symbols and to use the symbols as tools for thought.” However, this clearly seems to overstate the case here. Despite much training, no nonhuman animal has been shown to be able to read and write. In general, the issue here is not whether nonhuman animals have *some* ability for symbol cognition; the question is whether humans have a unique adaptation for symbol cognition. Overall, it is just not clear what the case against human-specific cognitive adaptations here is meant to be, or what the theoretical reasons are for thinking there are no such adaptations.

Putting all of this together, this leads to the following question. Why should we think that at the heart of human cognition are core knowledge systems, rather than *a combination* of our capacities for cultural or individual learning, our species-specific cognitive adaptations, and more widely shared innate representational endowments? Put differently, given that we should accept that *all* of (cultural and individual) learning, core knowledge, and uniquely human cognitive structures are parts of our innate psychological endowments, why give core knowledge such a prominent pride of place?

In response, the book makes a somewhat cryptic remark (p. 446): “I present this hypothesis as an alternative to theories that root our uniquely human cognitive achievements in our innate propensities for shared intentionality, for pedagogical learning, or for symbolic thought, but the hypothesis suggests there is truth to all these theories. Infants learn language from the members of their social world, and they put language to use, first and foremost, to share their experiences with the people they care about.” This is a little puzzling, though. Given that, as argued earlier, we need to accept that many organisms – including humans – have found individual and social learning to be adaptive, the latter cannot be seen as derived from systems of core knowledge. Similarly, because Spelke has not provided a compelling reason to doubt the existence of some human-specific cognitive adaptations, the latter cannot be explained as the result of the workings of core knowledge systems either.

Importantly, accepting this does not entail giving up on the importance of core knowledge systems. Rather, it just means acknowledging that cultural and individual learning, uniquely human cognitive adaptations, symbols and other forms of cognitive technology, and core knowledge *all* matter to human cognition. Indeed, I defend an account of distinctively human cognition of exactly this type: According to this, human cognition is the result of a variety of innate representations – some, but not all of which are integrated into sets of core knowledge systems – that are expanded and used through cultural learning and technology (see, e.g., my Schulz, 2020).

In a nutshell, the question at the heart of this commentary is: Why emphasize core knowledge so strongly – why make *that* its

own book? I look forward to hearing the answer to this question; it will undoubtedly be illuminating and lead to fruitful exchange.

**Competing interest.** None.

## References

- Boyd, R., & Richerson, P. (2005). *The origin and evolution of cultures*. Oxford University Press.
- Csibra, G., & Gergely, G. (2011). Natural pedagogy as evolutionary adaptation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1567), 1149–1157. <https://doi.org/10.1098/rstb.2010.0319>
- Henrich, J. (2015). *The secret of our success: How culture is driving human evolution, domesticating our species, and making us smarter*. Princeton University Press.
- Heyes, C. M. (2018). *Cognitive gadgets: The cultural evolution of thinking*. Harvard University Press.
- Schulz, A. (2020). Enhancing thoughts: Culture, technology, and the evolution of human cognitive uniqueness. *Mind & Language*, 37, 465–484. <https://doi.org/10.1111/mila.12320>
- Spelke, E. S. (2022). *What babies know: Core knowledge and composition*. Oxford University Press.

## Questioning the nature and origins of the “social agent” concept

Denis Tatone<sup>a\*</sup>  and Barbara Pomiechowska<sup>b</sup>

<sup>a</sup>Department of Cognitive Science, Central European University, Vienna, Austria and <sup>b</sup>Center for Developmental Science, Centre for Human Brain Health, School of Psychology, University of Birmingham, Birmingham, UK  
[TatoneD@ceu.edu](mailto:TatoneD@ceu.edu); [b.pomiechowska@bham.ac.uk](mailto:b.pomiechowska@bham.ac.uk)  
<https://denistatone.wixsite.com/my-site>; <https://bpomie.github.io/>

\*Corresponding author.

doi:10.1017/S0140525X23003011, e142

### Abstract

Spelke posits that the concept of “social agent,” who performs object-directed actions to fulfill social goals, is the first noncore concept that infants acquire as they begin to learn their native language. We question this proposal on empirical grounds and theoretical grounds, and propose instead that the representation of object-mediated interactions may be supported by a dedicated prelinguistic mechanism.

In chapter 10, Spelke (2022) articulates what is, in our view, the most thought-provoking contribution in her expanded review of the core knowledge (CK) theory. She argues that the concept of “social agent,” that is, an agent who interacts with others through object-directed actions (e.g., giving an object or demonstrating its function to someone), is not a part of CK. Instead, infants acquire this concept through exposure to their native language. By participating in a community of relevant and efficient speakers, they discover how to combine the representations of two distinct CKs: An *agent* system that interprets agents as acting on their environment to bring about instrumental goals, and a *social cognition* system that interprets agents as interacting with others through shared experiences. Spelke argues that prior to this discovery, occurring at around 10 months of age, infants fail to interpret social interactions involving object-directed actions because they lack the means to integrate the aforementioned systems. In

this commentary, we express empirical and theoretical reservations regarding this proposal.

To begin with, the developmental literature reveals several exceptions to the proposed timeline, suggesting that, at least in some settings, young infants are capable of interpreting the instrumental goals of “social agents.” For instance, by 4 months of age, infants begin to form expectations of distributive fairness, which requires representing interactions based on resource transfer (Buyukozer Dawkins, Sloane, & Baillargeon, 2019; Geraci & Surian, 2023). Around 6 months of age, infants infer social dominance on the basis of relative group size in scenarios involving agents with conflicting instrumental goals (Pun, Birch, & Baron, 2016; a finding that Spelke explains away by suggesting that infants construed the agents’ movements in the presence of group members as a “social gesture”: p. 412). By this age, infants also begin to show a manual preference for prosocial characters (Hamlin, Wynn, & Bloom, 2007; Kanakogi et al., 2017), suggesting an incipient understanding of instrumental actions directed at increasing the utility of other agents. Even if the interpretation of these findings as providing evidence for the representation of second-order social goals remains contested (Powell & Spelke, 2018), the success of the manual-choice measure, which entails selecting one of two ostensibly presented characters, shows that young infants can appropriately respond to acts of offering. Furthermore, by 9 months of age, infants can already represent the joint goals of agents involved in simple forms of collaboration (moving synchronously toward a common target; Begus, Curioni, Knoblich, & Gergely, 2020), and start to leverage ostensive signals not only to guide their attention to objects, but also to learn about their featural information (Okumura, Kanakogi, Kobayashi, & Itakura, 2020; Thiele, Hepach, Michel, & Haun, 2021).

Even discounting such exceptions, the failure of younger infants in interpreting the goals of “social agents” does not necessarily indicate a lack of corresponding conceptual frame. As Spelke points out, CKs compete for attentional resources. The task of identifying a teleological relation between the participants of an object-mediated interaction is bound to be more challenging than in the case of instrumental actions (or acts of social engagement) simply because the former, being a more structurally complex event, generates more concurrent goal hypotheses. For instance, typical helping actions (e.g., A helps B to open a box) contain cues of social affiliation (A approaches B and mirrors some of their movements) and first-order instrumental goals (A brings about a change of state in the box), either of which may prime well-formed goal hypotheses that prevent observers from recognizing the teleological dependency between the helper’s object-directed efforts and the facilitation of the helpee’s goal fulfillment (Schlingloff-Nemecz, Tatone, & Csibra, 2023). Similarly, instances of taking can be reduced to nonsocial acts of object acquisition, disregarding the effects that the action has on the original resource possessor (the takee). Supporting this possibility, infants and adults spontaneously interpret taking actions involving unreactive takees as nonsocial instances of resource acquisition (Tatone, Geraci, & Csibra, 2015; Yin, Csibra, & Tatone, 2022; Yin, Tatone, & Csibra, 2020), leading to the omission of the takee from the event frame.

The aforementioned evidence cautions against the claim that the composite concept of “social agent” does not compete for attention with concepts from other CKs (p. 422). To date, we do not know of any experimental evidence attesting to this claim (for a similar argument, see Revencu & Csibra, 2023). In fact, the above findings suggest that even adults, who should

presumably leverage such a “social agent” concept (and thus bypass the hurdle of competing available frames), similarly experience the conflict between interactive and instrumental construals. Under our account, the existence of such conflict well beyond infancy is evidence that arbitrating among competing frames is a challenge inherent to the process of action interpretation (Tatone, 2022), and not specific to developmentally early instances of attentional-resource competition among CKs. Importantly, unlike Spelke’s, this account does not entail any strong kind of dispositional ascription: it is not the agent, but their episodic behavior, that is being construed as instrumental or social, based on the available goal hypotheses. Because of this, infants (and adults) should be able to entertain agents as having goals pertaining to different CKs over time (e.g., engaging with a partner now, and pursuing their own instrumental goals later). This does not seem to be the case for Spelke, who characterizes CK frames as an instance of kind categorization, which, upon being deployed, constrains the type of potential goals that an agent may in principle pursue (e.g., acting on the environment vs. establishing social engagement).

The notion that the concept of “social agent” may derive from the infants’ pragmatic interpretation of object-directed communicative acts raises further concerns. To begin with, it is not quite clear how the linguistic combination of two distinct relations (one toward an object, the other toward a partner) could guide infants to discover specific utility functions underpinning a given social behavior (Powell, 2022). Even if infants could linguistically expand the interpretation of an agent’s object-directed actions to include social goals based on “engagement, shareability, and experience” (p. 349), this affiliative motive does not seem sufficiently precise to distinguish between different types of interdependence (Aktipis et al., 2018), such as prosocial versus mutualistic interactions. Arguing that the relevant distinctions may be acquired by participating in the linguistic community is rather problematic, for two reasons. First, it would require linguistic inputs to exhibit a tight mapping between thematic and beneficiary roles (i.e., who acts vs. who stands to gain from the action), which it is not always the case (Newman, 1996). Relatedly, linguistic descriptions of social interactions do not reveal the relational concepts (e.g., dominance, reciprocity) underlying these episodes (Fiske, 1992), despite these concepts being nonetheless inferred from the occurrence of specific social behaviors already by prelinguistic infants (Mascaro & Csibra, 2012; Tatone & Csibra, 2020).

More generally, Spelke’s reliance on natural language (as opposed to a language of thought) for combining core representations seems to run into an evolutionary conundrum: if the acquisition of the concept of “social agents” indeed depends on interactions with efficient and relevant speakers geared with the appropriate conceptual frame, how is this process originally bootstrapped? That is, what kind of combinatorial resources could the first language users tap into? Furthermore, if, as Spelke seems to recognize, concepts of “social agents” capture universal and species-wide experiences (e.g., giving, communicating), unlike a myriad of other concepts that are the product of socioculturally contingent traditions (e.g., integers), what accounts for this critical distinction, considering their surface similarity (i.e., being all discovered by participating in a linguistic community)? The reliable acquisition (across ontogenesis and different sociocultural niches) of the “social agent” concept, one may concede, attests to its adaptive relevance and centrality to human livelihood. But if that is the case (as we believe), we should expect concept learning to have been canalized through the emergence



of dedicated mechanisms over evolutionary time (Baldwin effect; Barrett, 2005). It is hard to assess whether the developmental timeline that Spelke sketches – with “social agents” being systematically the first noncore concept that infants acquire – bears evidence of such preparedness. Spelke does entertain the possibility of adaptive canalization when discussing our predisposition to associate snake avoidance with fear displays (p. 416), but considers specialized mechanisms to be mostly the province of the cognition of nonhuman animals.

With respect to nonhuman species, the proposal that the understanding of object-mediated interaction is heralded by the acquisition of natural language leaves us with two possible scenarios. One is that nonhuman animals, equipped with the same CKs as humans (as per Spelke’s criteria), lack a concept of “social agent,” and would thus be forced to navigate interactions based on food provisioning by construing them as either instances of nonsocial resource collection or goal-demoted forms of social bonding. This possibility seems unlikely, given that these interactions are among the first instances of caregiving in species with dependent offspring (from birds to social carnivores; Jaeggi & Van Schaik, 2011), and are thus pivotal to growth and survival. A more plausible alternative is that, in light of the fitness-relevant role that these interactions have, nonhuman animals evolved specialized mechanisms for recognizing and appropriately partaking in them (e.g., food offering calls or gestures; Jaeggi & Gurven, 2013; Scheid, Schmidt, & Noë, 2008). But if these arguments can be appealed to for explaining the proficiency of nonhuman animals in identifying certain interactions, what prevents a similar logic from being applied to our species?

Proposals in this direction have already been put forth (Frankenhuis & Barrett, 2013), also with regard to object-mediated interactions. In the domain of giving, for instance, it has been argued that infants are prepared to treat giving-like outcomes as potential goal states (Tatone et al., 2015) and to readily structure event representations in the service of capturing such goals. Buttressing this claim is the evidence that infants ascribe giving goals based on minimal cues of possession transfer, even when other perceptual indicators of engagement (e.g., social approach or receipt acknowledgment) are amiss (Tatone, Hernik, & Csibra, 2019). Such preparedness, we maintain, is to be expected because a generative model of action understanding governed solely by assumptions of instrumental rationality would be unable to compute the utility of actions, such as resource transfer, which inherently entail the voluntary self-imposition of costs. Since Spelke’s model relies on similar assumptions, it also requires the repertoire of instrumental goals that infants are originally geared with to be expanded to include the suite of social goals directed at increasing or decreasing another agent’s utility. How linguistic scaffolding enables the discovery of such goal states and their integration however, is presently not discussed.

Spelke’s theory of CK represents a trailblazing milestone in the study of human cognitive development. Drawing on an impressive body of experimental evidence, her latest book cements the tenets of this theory and expands its scope in daring new ways. Our commentary provided critical counterpoints to the idea that learning a natural language would allow infants to combine instrumental and social construals in the service of a new concept of “social agents.” In our view, the present proposal falls short of explaining with sufficient precision how this concept is discovered, how its utility function is derived and distinguished from the core concept of social engagement, and what accounts for its systematic acquisition. We hope our critical analysis may help spur further elaborations.


**Financial support.** This work was supported by a European Research Council grant (No. 742231 PARTNERS).

**Competing interest.** None.

## References

- Aktipis, A., Cronk, L., Alcock, J., Ayers, J. D., Baciu, C., Balliet, D., ... Winfrey, P. (2018). Understanding cooperation through fitness interdependence. *Nature Human Behaviour*, 2(7), 429–431.
- Barrett, H. C. (2005). Cognitive development and the understanding of animal behavior. In B. J. Ellis & D. F. Bjorklund (Eds.), *Origins of the social mind: Evolutionary psychology and child development* (pp. 438–467). Guilford.
- Begus, K., Curioni, A., Knoblich, G., & Gergely, G. (2020). Infants understand collaboration: Neural evidence for 9-month-olds’ attribution of shared goals to coordinated joint actions. *Social Neuroscience*, 15(6), 655–667.
- Buyukozer Dawkins, M., Sloane, S., & Baillargeon, R. (2019). Do infants in the first year of life expect equal resource allocations?. *Frontiers in Psychology*, 10, 116.
- Fiske, A. P. (1992). The four elementary forms of sociality: Framework for a unified theory of social relations. *Psychological Review*, 99(4), 689.
- Frankenhuis, W. E., & Barrett, H. C. (2013). Design for learning: The case of chasing. In M. D. Rutherford & V. A. Kuhlmeier (Eds.), *Social perception. Detection and interpretation of animacy, agency, and intention* (pp. 171–198). MIT.
- Geraci, A., & Surian, L. (2023). Intention-based evaluations of distributive actions by 4-month-olds. *Infant Behavior and Development*, 70, 101797.
- Hamlin, J. K., Wynn, K., & Bloom, P. (2007). Social evaluation by preverbal infants. *Nature*, 450(7169), 557–559.
- Jaeggi, A. V., & Gurven, M. (2013). Natural cooperators: Food sharing in humans and other primates. *Evolutionary Anthropology: Issues, News, and Reviews*, 22(4), 186–195.
- Jaeggi, A. V., & Van Schaik, C. P. (2011). The evolution of food sharing in primates. *Behavioral Ecology and Sociobiology*, 65, 2125–2140.
- Kanakogi, Y., Inoue, Y., Matsuda, G., Butler, D., Hiraki, K., & Myowa-Yamakoshi, M. (2017). Preverbal infants affirm third-party interventions that protect victims from aggressors. *Nature Human Behaviour*, 1(2), 37.
- Mascaro, O., & Csibra, G. (2012). Representation of stable social dominance relations by human infants. *Proceedings of the National Academy of Sciences of the United States of America*, 109(18), 6862–6867.
- Newman, J. (1996). *Give: A cognitive linguistic study* (No. 7). Walter de Gruyter.
- Okumura, Y., Kanakogi, Y., Kobayashi, T., & Itakura, S. (2020). Ostension affects infant learning more than attention. *Cognition*, 195, 104082.
- Powell, L. J. (2022). Adopted utility calculus: Origins of a concept of social affiliation. *Perspectives on Psychological Science*, 17(5), 1215–1233.
- Powell, L. J., & Spelke, E. S. (2018). Third-party preferences for imitators in preverbal infants. *Open Mind*, 2(2), 61–71.
- Pun, A., Birch, S. A., & Baron, A. S. (2016). Infants use relative numerical group size to infer social dominance. *Proceedings of the National Academy of Sciences of the United States of America*, 113(9), 2376–2381.
- Revenu, B., & Csibra, G. (2023). The missing link between core knowledge and language: Review of Elizabeth Spelke’s *What Babies Know*, Volume 1 (2022). *Mind & Language*, 38(5), 1314–1322. doi: 10.1111/mila.12482
- Scheid, C., Schmidt, J., & Noë, R. (2008). Distinct patterns of food offering and co-feeding in rooks. *Animal Behaviour*, 76(5), 1701–1707.
- Schlingloff-Nemecz, L., Tatone, D., & Csibra, G. (2023). The representation of third-party helping interactions in infancy. *Annual Review of Developmental Psychology*, 5, 67–88. doi: 10.1146/annurev-devpsych-120321-033548
- Spelke, E. S. (2022). *What babies know: Core knowledge and composition*. Oxford University Press.
- Tatone, D. (2022). More than one way to skin a cat: Addressing the arbitration problem in developmental science. *Behavioral and Brain Sciences*, 45, e123.
- Tatone, D., & Csibra, G. (2020). Infants infer different types of social relations from giving and taking actions. In M. L. Denison, M. L. Mack, Y. Xu, & B. Armstrong (Eds.), *Proceedings of the 42nd Annual Conference of the Cognitive Science Society* (pp. 2981–2987). Cognitive Science Society.
- Tatone, D., Geraci, A., & Csibra, G. (2015). Giving and taking: Representational building blocks of active resource-transfer events in human infants. *Cognition*, 137, 47–62.
- Tatone, D., Hernik, M., & Csibra, G. (2019). Minimal cues of possession transfer compel infants to ascribe the goal of giving. *Open Mind*, 3, 31–40.
- Thiele, M., Hepach, R., Michel, C., & Haun, D. (2021). Observing others’ joint attention increases 9-month-old infants’ object encoding. *Developmental Psychology*, 57(6), 837.
- Yin, J., Csibra, G., & Tatone, D. (2022). Structural asymmetries in the representation of giving and taking events. *Cognition*, 229, 105248.
- Yin, J., Tatone, D., & Csibra, G. (2020). Giving, but not taking, actions are spontaneously represented as social interactions: Evidence from modulation of lower alpha oscillations. *Neuropsychologia*, 139, 107363.

## Early pragmatic expectations in human infancy

Tibor Tauzin<sup>a</sup>, Pierre Jacob<sup>b</sup> and György Gergely<sup>c\*</sup> 

<sup>a</sup>Institute of Linguistics, University of Vienna, Wien, Austria; <sup>b</sup>Institut Jean Nicod, CNRS UMR 8129, Paris, France and <sup>c</sup>Department of Cognitive Science, Central European University, Wien, Austria  
[tibor.miklos.tauzin@univie.ac.at](mailto:tibor.miklos.tauzin@univie.ac.at); [jacobpiotr11@gmail.com](mailto:jacobpiotr11@gmail.com); [gergelygy@ceu.edu](mailto:gergelygy@ceu.edu)

\*Corresponding author.

doi:10.1017/S0140525X23003230, e143

### Abstract

There is no room for pragmatic expectations about communicative interactions in core cognition. Spelke takes the combinatorial power of the human language faculty to overcome the limits of core cognition. The question is: Why should the combinatorial power of the human language faculty support infants' pragmatic expectations not merely about speech, but also about nonverbal communicative interactions?

Spelke's picture of early infant cognition in her monumental (2022) book, *What Babies Know*, involves two basic cognitive mechanisms: core cognitive systems (none of which is unique to humans) and the human language faculty (which uniquely enables human infants to speak their native tongue). The agents core system generates infants' expectations about the efficiency of an agent's instrumental action directed toward a physical target. The core social cognition system enables infants to represent people as social beings who interactively engage and share phenomenal experiences with others. Initially the core systems are independent, compete for attention, and cannot interact with each other. Therefore, young infants cannot simultaneously attend to people as agents who efficiently pursue instrumental goals and as social beings who engage and share phenomenal experiences with others through coordinated social actions.

On Spelke's account, the human language faculty does not merely enable preverbal infants to learn the words and grammar of their native language; it also enables them to overcome the initial representational and attentional limitations of their core cognitive systems. On Spelke's view, infants' pragmatic expectations about communication are derived from their social interactions with speakers whose utterances are used to "express their thoughts to others, in accord with pragmatic principles of economy, informativeness, and relevance" (p. 435). Spelke suggests that these pragmatic expectations may arise automatically from the agents core system. As a result of language acquisition that allows the integration of the agents core system with the core social cognition system, infants' expectations about the efficiency of an agent's instrumental action are transferred from the agents core system to the core social cognition system. Nonverbal communicative interactions between social agents, however, are not supposed to induce pragmatic expectations in young infants.

Therefore, without the benefits of language acquisition, infants should be unable to construe an agent's gaze shift toward one of a pair of physical targets as an object-directed action whose social goal is to share attention to the referent object with a social partner. Nor can they understand such a nonverbal ostensive

referential gesture in terms of the social agent's communicative and informative intentions to convey to his recipient relevant information about the indicated referent. In contrast to this account, we shall argue for an alternative view according to which significantly earlier than 12 months young infants exhibit special sensitivity to nonverbal behavioral cues of ostensive communication that generate their pragmatic expectations about the referentiality, informativity, and relevance of the communicative actions of social agents. As evidence for this view, we shall focus on four relevant studies that demonstrate early sensitivity to nonverbal ostensive cues of communicative intention, and indicate referential expectations as well as pragmatic expectations about the relevance of the ostensively communicated information.

In a recent study by Tauzin and Gergely (2019), 10-month-olds watched videos depicting the interactions between two unfamiliar agents whose sole activity consisted in exchanging unfamiliar nonlinguistic sound triplets in a turn-taking manner. In the identical signals condition, the sequence of sounds produced by the second agent was fully predictable from the first agent's sound signals because the second agent strictly replicated the sound triplet emitted by the first agent. In the variable signals condition, the second agent's sound triplets were not fully predictable because they replicated only the initial sound of the first agent's sound triplets while the second and third sounds of the second agent's sound triplets were freely varied. In the subsequent test phase, only one of the two agents was present with two laterally positioned objects on its two sides, and the agent turned toward one of the two objects. Tauzin and Gergely found that the agent's target-oriented movement induced a gaze-following response in the 10-month-olds only in the variable signals condition, not in the identical signals condition. This finding does not seem compatible with the assumptions of the core social cognition system which, according to Spelke, would be expected to generate imitative attunement and gaze-following of the entity's attentional orienting response in both conditions. Arguably, however, new information can only be conveyed by (partially) unpredictable but not fully predictable signal sequences (cf. Shannon, 1948). If so, then gaze-following responses subsequent to the social agent's target-oriented movement only in the variable signal condition indicate that infants are sensitive to the pragmatic principle of informativity and they rely on it to identify communicative information transfer between agents.

Recently, Okumura, Kanakogi, Kobayashi, and Itakura (2020) reported a gaze-following study with 9-month-olds to investigate the predictions of the theory of natural pedagogy (Csibra & Gergely, 2011). They tested whether ostensive cues (e.g., eye-contact or infant-directed speech) induce pragmatic expectations of referentiality, informativity, and relevance in infants in contrast to noncommunicative attention-grabbing behaviors (e.g., shivering or uttering beeping sounds). The results showed that 9-month-olds followed the model's gaze shift and spent an equal amount of time looking at the target object in both conditions. However, in a further test, infants were presented with pairs of pictures one of which depicted the previously fixated target and the other depicted a novel object. Okumura and colleagues hypothesized that ostensive cues were likely to boost the infants' processing, encoding, and retention of the properties of the fixated object. As a result, they expected a novelty effect and predicted that infants would look longer at the picture of the previously nonfixated object. The results showed that the infants looked reliably longer at the picture of the novel than of the fixated object in the ostensive, but not in the nonostensive attention-grabbing condition in accordance

with their hypothesis. Moreover, in a further object choice test the infants were given the opportunity to choose between a three-dimensional (3D) replica of the previously fixated object and a 3D replica of an unfamiliar novel object. Okumura et al. (2020) argued that the selective preference toward the target object indicated by the model's gaze shift can be regarded as "evidence that the actor's gaze impacts the affective appraisal of objects." In line with this assumption, infants showed selective preference for choosing the previously fixated object in the ostensive condition whereas no selective object choice was found in the noncommunicative attention enhancing behaviors condition. These findings suggest that the model's ostensive communicative signals induced not only referential expectations in the 9-month-olds but also pragmatic expectations of informativity and relevance about the referent object that was communicatively manifested for them by the model's referential gesture.

In an object individuation study with 10-month-olds, Futó, Téglás, Csibra, and Gergely (2010) demonstrated that in an ostensive cuing context, manual demonstrations of the functions of two novel artifacts can induce kind-based object individuation even in the absence of naming the objects with words. In the ostensive condition, the infants were first addressed by infant-directed speech before a hand brought out the novel artifacts on either sides of an occluder and performed different means actions on them that generated either a sound or a visual effect. After the occluder was removed the 10-month-olds looked reliably longer when one rather than two objects were present. In a nonostensive attention induction condition (using a mechanical sound), however, the very same goal-directed action manipulations did not induce longer looking at one versus two objects. This study indicates that similarly to the use of word labels in an ostensive context, nonverbal action demonstrations are interpreted by young infants as communicative manifestations of kind-relevant informational properties of novel objects, such as their functions.

A groundbreaking study by Vouloumanos, Martin, and Onishi (2014) provides further evidence that in many relevant respects even 6-month-old infants process verbal conversations between speakers of their linguistic community the way monolingual adults process conversations between speakers of a foreign language. They can recognize the presence of a speaker's communicative intention before they can speak their native tongue. Infants were first familiarized with a single agent (the speaker) who repeatedly showed her preference to play with one of a pair of toys in front of her. In the next scene, the speaker appeared behind a tiny window and could not reach the toys anymore. Opposite to her, however, appeared someone else (the recipient) who could both see and act on the toys located between them. The speaker turned toward the recipient and either uttered the (novel) word "koba" or produced a coughing sound. When the speaker uttered "koba," but not when she coughed, 6-month-olds looked reliably longer if the recipient picked up the toy that was not the speaker's preferred toy rather than when she picked up the speaker's favorite toy. Although coughing unquestionably drew infants' attention toward the speaker, only the speaker's utterance of "koba," not coughing, triggered their referential and communicative expectations. Arguably, infants ascribed a communicative intention to the speaker and used the context to fill in the content of her informative intention. They knew about the speaker's preference for one of the toys and could see that she was patently unable to satisfy her preference by her own bodily action. In this context, they expected the speaker to make a request and not an assertion that could only be fulfilled

by the recipient if the latter took the speaker's tokening of "koba" to refer to her favorite toy.

In an important recent study, Neff and Martin (2023) replicated these findings and provided further evidence showing that 6-month-olds do not assume that a verbal utterance is a sufficient condition to ensure the success of a speaker's communicative action. Neff and Martin found that only in an ostensive context, in which the speaker and the recipient are in face-to-face contact during speech (not if either is looking elsewhere) do infants expect the recipient to pick up the speaker's favorite toy and thereby fulfill her informative intention. These results show that even preverbal infants possess a sensitivity to nonverbal cues of ostensive communication which induces their pragmatic expectations about the relevance of communicated information.

The evidence reviewed above suggests that in response to ostensive stimuli, infants form pragmatic expectations about an agent's verbal or nonverbal communicative action. These include expectations about an agent's referential action as well as the expectation that the communicative agent is seeking to convey information relevant to her recipient.

There is, however, no room for pragmatic expectations in either the agents core system or the core social cognition system. On Spelke's account, pragmatic expectations arise from the human language faculty. There are, however, two possible interpretations of the role of the human language faculty. One possibility is that infants' expectations about the efficiency of an agent's instrumental action generated by the core agents system are converted by the combinatorial power of the human language faculty into pragmatic expectations about communicative actions. It is unclear, however, how the combinatorial power of the language faculty could fill the gap between expectations about the efficiency of an agent's instrumental action and pragmatic expectations of relevance and informativity in communicative interactions. A second possibility is that speakers of a natural language express their thoughts in accordance "with pragmatic principles of economy, informativeness, and relevance" because the pragmatic principles of communication are built into the human language faculty itself. If so, then communication would likely be a major function of the human language faculty – a view adamantly rejected by Chomsky (Bolhuis, Tattersall, Chomsky, & Berwick, 2014). In this case, however, pragmatic expectations about nonverbal communicative actions in human infancy would be puzzling.

The alternative not explored by Spelke's (2022) monumental book is that preverbal human infants are innately prepared to form pragmatic expectations about an agent's verbal or nonverbal communicative acts (see Gergely & Jacob, 2012). So far as we know, unlike nonhuman great apes, humans are uniquely disposed to provide information relevant to others and conversely to extract information relevant to themselves from others' ostensive-communicative displays (cf. Tomasello, 2014). This mutual adjustment suggests a biological adaptation rather than an ontogenetic explanation in terms of learning processes and poses difficulties for Spelke's developmental account.

**Financial support.** This research received no specific grant from any funding agency, commercial, or not-for-profit sectors.

**Competing interest.** None.


## References

- Bolhuis, J. J., Tattersall, I., Chomsky, N., & Berwick, R. C. (2014). How could language have evolved? *PLoS Biology*, 12(8), e1001934.



- Csibra, G., & Gergely, G. (2011). Natural pedagogy as evolutionary adaptation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1567), 1149–1157.
- Futó, J., Téglás, E., Csibra, G., & Gergely, G. (2010). Communicative function demonstration induces kind-based artifact representation in preverbal infants. *Cognition*, 117(1), 1–8.
- Gergely, G., & Jacob, P. (2012). Reasoning about instrumental and communicative agency in human infancy. *Advances in Child Development and Behavior*, 43, 59–94.
- Neff, M. B., & Martin, A. (2023). Do face-to-face interactions support 6-month-olds' understanding of the communicative function of speech? *Infancy*, 28(2), 240–256.
- Okumura, Y., Kanakogi, Y., Kobayashi, T., & Itakura, S. (2020). Ostension affects infant learning more than attention. *Cognition*, 195, 104082.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27(3), 379–423.
- Spelke, E. S. (2022). *What babies know: Core knowledge and composition* (Vol. 1). Oxford University Press.
- Tauzin, T., & Gergely, G. (2019). Variability of signal sequences in turn-taking exchanges induces agency attribution in 10.5-mo-olds. *Proceedings of the National Academy of Sciences of the United States of America*, 116(31), 15441–15446.
- Tomasello, M. (2014). *A natural history of human thinking*. Harvard University Press.
- Vouloumanos, A., Martin, A., & Onishi, K. H. (2014). Do 6-month-olds understand that speech can communicate? *Developmental Science*, 17(6), 872–879.

## Core knowledge as a neuro-ethologist views it

Giorgio Vallortigara\* 

Centre for Mind/Brain Sciences, University of Trento, Rovereto, TN, Italy  
[Giorgio.vallortigara@unitn.it](mailto:Giorgio.vallortigara@unitn.it)  
<http://r.unitn.it/en/cimec/abc>

\*Corresponding author.

doi:10.1017/S0140525X23003035, e144

### Abstract

Innateness of core knowledge mechanisms (in the form of “cognitive priors”) can be revealed by proper comparisons of altricial and precocial species. Cognitive priors and sensitive periods in their expression may also provide clues for the development of plausible artificial intelligence systems.

Elizabeth Spelke (Spelke, 2022) championed in her book – and throughout her scientific career – the idea that research on infants and (non-human) animals should synergize and thrive. In particular, research on human newborns and animals of precocial species (which, differently than those of altricial species, are at mature sensory and motoric levels at birth) would allow a direct approach to fundamental questions: What do organisms know at the time when their learning begins? And what makes their learning go so well, that is, so quickly?

Here I want to remark on some crucial insights that arose from a comparative approach such as this. The first two examples show up as a sort of clash of evidence from human and non-human studies, and their resolution yields, I believe, real progress in scientific understanding.

The first example is described by Spelke in detail in the book, and it is related to recognition of partly occluded objects, but I will consider it as regards one neglected aspect of it. When we published, more than 35 year ago (Regolin & Vallortigara, 1995), our results showing that newborn chicks are capable of

recognizing partly occluded objects (or “amodal completion” as Michotte, Thines, & Crabbé [1964] had dubbed the phenomenon), and then Lea, Slater, and Ryan (1996) duplicated our results using a procedure more similar to that employed by developmental psychologists, we were faced with an apparent paradox. Even taking into account that chicks are precocial and human infants altricial, it appeared that at least 4 months were needed for young humans to develop a capacity for amodal completion that chicks showed at the onset of life (Vallortigara, 2021). It is tempting in these cases to argue that humans, differently than other animals, rely more on learning rather than on innate predispositions. But this is a weird argument in my view: If the machinery for amodal completion is available at the start as an evolutionarily granted mechanism (based on rules associated with the way in which junctions of objects' boundaries occur in visual scenes), why should an organism accept the cost of acquiring such a capacity by trial and error learning, which is likely to be very long and is open to risks of mistakes? The solution, as Spelke tells us in the book, was provided by Valenza, Leo, Gava, and Simion (2006), who showed that using stroboscopic motion, rather than slow, gradual motion, as human newborns do, shows evidence for amodal completion at birth. Claiming innateness is obviously not the same as claiming that a certain capacity is operational at birth, for this would depend on the overall pattern of development of a species. Precocial species such as chicks have mechanisms for motion perception which are mature soon after hatching, whereas in humans, an altricial species, these mechanisms will mature later on. In the latter species, it may be difficult to reveal the presence of amodal completion at birth, in spite of the fact that the mechanism is already there, and it is indeed innate, for it simply does not show up until proper testing conditions are used (such as those measuring stroboscopic motion).

The second example is more recent. Starting from the classical work of Francis Galton (1880), the idea that humans have a mental number line has obtained several confirmations (e.g., Dehaene, Bossini, & Giraux, 1993). Number and space seem to be inherently associated, and these kinds of phenomena (also referred to as SNAs, as in Space-Number Associations) have been extensively investigated. Yet the debate regarding their nature and origin remains hot. For example, in recent years, two papers have been published, among others, trying to unveil the origin of SNAs, that reached diametrically opposite conclusions: One showing that chicks exhibit something similar to an ordered mental number line (Rugani, Vallortigara, Priftis, & Regolin, 2015), the other showing that a traditional human population lacking any formal arithmetic does not show any significant left-to-right bias (Pitt et al., 2021). On the basis of second paper and several other anthropological studies, it can be concluded that SNAs are determined by cultural habits dependent upon literacy (e.g., the direction of reading/writing), as they were absent in preschoolers and indigenous people from oral tribes. On the basis of the first paper, one can conclude, on the contrary, that SNAs (and in particular, a left-to-right mapping) are biologically, and not culturally, determined; indeed, such associations are observed in pre- or non-verbal subjects, such as human newborns (Di Giorgio et al., 2019) and in several non-human animal species (e.g., monkeys: Drucker & Brannon, 2014; honeybees: Giurfa, Marcout, Hilpert, Thevenot, & Rugani, 2022).

Recently, we aimed to resolve these apparently conflicting findings (Eccher et al., 2023). We conducted two behavioral experiments in three populations of different ages and cultural

backgrounds: Italian adults, Italian preschoolers, and adults from the Himba tribe (an indigenous African tribe with no writing system). Our results showed that when tested with explicit tasks, only Italian adults show a consistent SNA, while when tested with implicit tasks, all the three populations exhibit a common and consistent left-to-right-oriented SNA.

These results support the hypothesis that the SNA phenomenon is dissociable into two different components: One which is acquired and cultural-dependent, and one which is biologically predisposed (note that the underlying mechanisms can be nonetheless quite different from those of a proper number line and rely instead on brain asymmetry, see, for a specific hypothesis, Vallortigara, 2018). But apart from the particular case in point, what seems interesting to me is that there is a special value in this sort of comparative research, namely the fact that evidence in non-human species *forced us* to reconsider our hypotheses on human nature.

The third insight concerns the proper way to build up intelligence in artificial systems. Plasticity seems to be a magic word nowadays in neuroscience and also in artificial intelligence (AI; but see Marcus, 2018). Comparative research on core knowledge tells us a different story. Consider what we learned from newly hatched chicks. These animals seem to be predisposed to orient toward objects that possess features associated with animate objects, such as biological motion, changes in speed, and face-like configurations (review in Di Giorgio et al., 2017; Vallortigara, 2021 see also Vallortigara, Regolin, & Marconato, 2005). These are unlearned priors that help chicks to orient toward the mother hen and their siblings, thus facilitating and guiding a robust learning process called filial imprinting. (Similar mechanisms have been documented in human newborns [Lorenzi & Vallortigara, 2021; Vallortigara, 2012, 2021], even though in this species, proper control of past experience and access to neural substrates are limited for obvious reasons.)

The issue then is how can young organisms orient toward the “right” stimulus in the absence of any previous experience? In contrast to machine-learning systems, biological organisms do not require explicit reinforcement, supervised learning, or thousands/millions of examples to feed learning. They are equipped with dedicated orienting mechanisms that work as adaptive priors that imply some assumptions about the external world that guide learning (Versace, Martinho-Truswel, Kacelnik, & Vallortigara, 2018). The priors are sufficiently general to allow errors. For instance, early preferences of chicks are not strictly species-specific but apply equally to hen face-like or polecat face-like features or to the biological-motion appearance of either a hen or a cat (review in Vallortigara, 2021). There is a profound biological reason for that. The predisposed orienting mechanisms cannot be too specific for the individual features, given that these are to some extent unpredictable from the genetic repertoire (because of variability between adults within a species and due to changes in the appearance of even a single individual).

However, high plasticity coupled with prior assumptions is not enough. In biological organisms, both early predispositions and high plasticity are transient phenomena that end either with some maturational processes or when the necessary information has been acquired. The existence of critical and sensitive periods has been documented in several domains and functions (Hensch & Quinlan, 2018). However, a very important recent finding is that critical periods do not apply only to the plasticity associated with learning but also to the periods of expression of the priors themselves. For instance, cues of animacy associated with speed changes

are expressed in chicks during a restricted period in early life and can be reopened by the administration of certain substances (Lorenzi, Lemaire, Versace, Matsushima, & Vallortigara, 2021).

There are costs associated with neural plasticity, and this is the reason why, after a certain age, learning new languages and solving amblyopia is so hard, and why early experiences are important for subsequent stages of life. Thus, the plasticity of the nervous system is actively reduced by molecular “brakes” that promote the stabilization of mature brain function (Hensch & Quinlan, 2018).

Elizabeth Spelke’s book seems to suggest, among other things, that AI systems could benefit from being equipped with a set of priors that offer a guidance system and a way to speed up plasticity associated with learning mechanisms. I believe, however, that plasticity without critical periods of expression for these priors might have costs that prevent effective learning and cognitive functions.

**Financial support.** This work was supported by funding from the European Research Council under the European Union’s Horizon 2020 research and innovation program (Grant Agreement 833504 SPANUMBRA).


**Competing interest.** None.

## References

- Dehaene, S., Bossini, S., & Giraux, P. (1993). The mental representation of parity and number magnitude. *Journal of Experimental Psychology: General*, 122, 371–396.
- Di Giorgio, E., Loveland, J. L., Mayer, U., Rosa-Salva, O., Versace, E., & Vallortigara, G. (2017). Filial responses as predisposed and learned preferences: Early attachment in chicks and babies. *Behavioural Brain Research*, 325, 90–104.
- Di Giorgio, E., Lunghi, M., Rugani, R., Regolin, L., Dalla Barba, B., Vallortigara, G., & Simion, F. (2019). A mental number line in human newborns. *Developmental Science*, 22, 1–10.
- Drucker, C. B., & Brannon, E. M. (2014). Rhesus monkeys (*Macaca mulatta*) map number onto space. *Cognition*, 132, 57–67.
- Eccher, E., Jossierand, M., Caparos, S., Boissin, E., Buiatti, M., Piazza, M., & Vallortigara, G. (2017). A universal left-to-right bias in number-space mapping across ages and cultures. *PsyArXiv*. <https://doi.org/10.31234/osf.io/w2st6>
- Galton, F. (1880). Visualised numerals. *Nature*, 21, 252–256.
- Giurfa, M., Marcourt, C., Hilpert, P., Thevenot, C., & Rugani, R. (2022). An insect brain organizes numbers on a left-to-right mental number line. *Proceedings of the National Academy of Sciences of the USA*, 119, 2203584119.
- Hensch, T. K., & Quinlan, E. M. (2018). Critical periods in amblyopia. *Visual Neuroscience*, 35, E014.
- Lea, S. E. G., Slater, A. M., & Ryan, C. M. E. (1996). Perception of object unity in chicks: A comparison with human infant. *Infant Behavior and Development*, 19, 501–504.
- Lorenzi, E., Lemaire, B. S., Versace, E., Matsushima, T., & Vallortigara, G. (2021). Resurgence of an inborn attraction for animate objects via thyroid hormone T3. *Frontiers in Behavioral Neuroscience*, 15, 675994.
- Lorenzi, E., & Vallortigara, G. (2021). Evolutionary and neural bases of the sense of animacy. In A. Kaufman, J. Call & J. Kaufman (Eds.), *The Cambridge handbook of animal cognition* (pp. 295–321). Cambridge University Press.
- Marcus, G. (2018). Innateness, alphazero, and artificial intelligence. *arXiv*.
- Michotte, A., Thinés, G., & Crabbé, G. (1964). *Les compléments amodaux des structures perceptives [Amodal complements of perceptual structures]*. Publications Universitaires de Louvain.
- Pitt, B., Ferrigno, S., Cantlon, J. F., Casasanto, D., Gibson, E., & Piantadosi, S. T. (2021). Spatial concepts of number, size, and time in an indigenous culture. *Science Advances*, 7(33), 1–7.
- Regolin, L., & Vallortigara, G. (1995). Perception of partly occluded objects by young chicks. *Perception and Psychophysics*, 57, 971–976.
- Rugani, R., Vallortigara, G., Priftis, K., & Regolin, L. (2015). Number-space mapping in the newborn chick resembles humans’ mental number line. *Science*, 347, 534–536.
- Spelke, E. (2022). *What babies know*. Oxford University Press.
- Valenza, E., Leo, I., Gava, L., & Simion, F. (2006). Perceptual completion in newborn human infants. *Child Development*, 77, 1810–1821.
- Vallortigara, G. (2012). Core knowledge of object, number, and geometry: A comparative and neural approach. *Cognitive Neuropsychology*, 29, 213–236.

- Vallortigara, G. (2018). Comparative cognition of number and space: The case of geometry and of the mental number line. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 373, 20170120.
- Vallortigara, G. (2021). *Born knowing: Imprinting and the origins of knowledge*. MIT Press.
- Vallortigara, G., Regolin, L., & Marconato, F. (2005). Visually inexperienced chicks exhibit a spontaneous preference for biological motion patterns. *PLoS Biology*, 3, 1312–1316 (e208).
- Versace, E., Martinho-Truswel, A., Kacelnik, A., & Vallortigara, G. (2018). Priors in animal and artificial intelligence: Where does learning begin? *Trends in Cognitive Sciences*, 22, 963–965.

## Developmental origin of a language–cognition interface in infants: Gateway to advancing core knowledge?

Sandra R. Waxman 

Department of Psychology, Northwestern University, Evanston, IL, USA  
<https://psychology.northwestern.edu/people/faculty/core/profiles/sandra-waxman.html>  
[www.childdevelopment.northwestern.edu](http://www.childdevelopment.northwestern.edu)

### Corresponding author:

Sandra R. Waxman;  
 Email: [s-waxman@northwestern.edu](mailto:s-waxman@northwestern.edu)

doi:10.1017/S0140525X23003096, e145

### Abstract

Spelke's sweeping proposal requires greater precision in specifying the place of language in early cognition. We now know by 3 months of age, infants have already begun to forge a link between language and core cognition. This precocious link, which unfolds dynamically over development, may indeed offer an entry point for acquiring higher-order, abstract conceptual and representational capacities.

Spelke's (2022) sweeping proposal for core knowledge is breathtaking for its theoretical depth and empirical reach. And like any comprehensive theory, it raises new challenges. A foundational challenge for this proposal, as currently construed, is articulating with greater precision the place of language in the creation of knowledge. Spelke acknowledges the gravity of this challenge, especially because it is essential to her argument that language is the “glue” that permits human infants to weave together distinct pieces of core knowledge from distinct cognitive domains.

Meeting this challenge will require rigorous assessment of how – and how early – an interface between language and core knowledge unfolds in infants. Questions concerning this interface have had a long history of spirited debate. But the question is no longer *whether* language and cognition are linked, but *how* this quintessentially human link emerges. How do infants begin to forge an interface between language and cognition in the first place? And what advantages does this link afford the developing infant mind? Questions like these have taken center stage in the developmental sciences. The accumulating evidence, exciting in its own right, bears directly on Spelke's provocative proposal. Spelke

herself makes a start, marshaling evidence from infants as young as 9 months of age. But we now know that a language–cognition interface is already in place, at least in rudimentary form, far earlier.

Here, my goal is to shed light on the developmental origin and rapid unfolding of a language–cognition link in very young infants, to trace its increasing power and precision over the first year of development, and to underscore that it is imperative that we consider this link as a dynamic one that unfolds over the first year. To foreshadow, listening to language supports core cognitive capacities in infants as young as 3 months; at issue is whether and how this precocious language–cognition link ultimately serves as a gateway for a suite of conceptual and representational capacities that are the signatures of human cognition.

### A surprisingly early link

By 3 months of age, before they can roll over in their own cribs or recognize the sound of their own names, simply listening to language supports infants' core cognitive capacities, including object categorization and abstract rule-learning, and does so in a way that other carefully matched acoustic signals (e.g., sine-wave tone sequences, reversed speech) do not. Initially, this link is quite broad. Listening to vocalizations of nonhuman primates confers precisely the same cognitive benefit as listening to their native language (Ferry, Hespos, & Waxman, 2010, 2013); this initial link is also sufficiently broad to include signed (American Sign Language) as well as spoken language (Novack, Brentari, Goldin-Meadow, & Waxman, 2021). But by 6 months, infants have tuned this link specifically to their native language (s) (for a review, see Perszyk & Waxman, 2018).

This precocious link between language and cognition, and its rapid tuning, reveals unique contributions of innate capacities and infants' experience. The initially broad link cannot be built on experience alone: By 3 months, infants have gained ample exposure to English, but virtually none to nonhuman primate vocalizations or sign language. Yet these signals confer the same cognitive advantage at 3 and 4 months.

### An increasingly precise and powerful link

A few months later, the consequences of listening to language become considerably more nuanced and powerful. By 7 months, even before infants say their first words, naming a set of objects (e.g., a dog, a horse, a duck) with the same, consistently applied name focuses infants' attention on their commonalities and supports the formation of an object category (animal). Yet providing a distinct name for each object has a very different effect, focusing infants' attention on distinctions among the objects and supporting their representation of each as a unique individual (LaTourrette, Chan, & Waxman, 2023; LaTourrette & Waxman, 2020; Xu, Cote, & Baker, 2005). This provides infants, like adults, with exceptional conceptual and representational flexibility. For example, they can represent any object (e.g., the family dog) *either* as a unique individual (Rover) *or* as a member of an object category (e.g., a dog). This flexibility is supported by language: How an object is *named* – either as a unique individual or a member of a category – is instrumental to how we *mentally represent* that object. And this representation, in turn, has powerful downstream consequences, guiding their learning and reasoning about objects. Future work will be required to discover whether



7-month-old infants' individual and object kind representations, guided by naming, are sufficiently robust to support higher order, and perhaps even combinatorial, processes.

### *A dynamic link, shaped by cascading effects of infants' advances in language and cognition*

Perhaps most importantly from a developmental vantage point, the language–cognition link is not a steady state. What an infant gleans from listening to language at 3 months will vary considerably from what they will glean at 7 months and later. Therefore, a strong developmental approach is required if we are to trace the unfolding of this link, one that takes into account the cascading and dynamic effects of infants' linguistic and cognitive advances as they unfold. This is especially important because infants' linguistic and cognitive capacities are certainly not fixed; their perception of language and ability to learn from it evolve dramatically over the first year. By 7 months, infants deftly identify distinct words in continuous speech and link them to distinct kinds of representations. Put differently, at this time, they begin to establish reference. By roughly their first birthdays, infants begin to combine property and kind concepts, evoked by predicates and nouns, respectively. But at 3 months, when infants do not yet even parse individual words from the continuous speech-stream, this kind of precision is far beyond reach.

How, then, might listening to language promote core cognition in infants so young? We have proposed that for such very young infants, listening to language engages systems of arousal and attention, heightening infants' attention to the objects in their environment, and that this, in turn, facilitates downstream core cognitive processes, like object categorization and rule-learning (Woodruff Carr et al., 2021).

### *In closing*

Very young infants reveal a surprisingly precocious interface between language and cognition, one that becomes increasingly powerful and precise over the first year of life. This interface, which evolves considerably over infants' first year, may provide a gateway for advancing core systems of knowledge and establishing the higher-order, abstract and combinatorial representations that distinguish human thought from that of our nearest evolutionary relatives.

**Financial support.** This work was supported by the National Institutes of Health (R01 HD088310) and Northwestern's Institute for Policy Research.

**Competing interests.** None.


### References

- Ferry, A. L., Hespos, S. J., & Waxman, S. R. (2010). Categorization in 3- and 4-month-old infants: An advantage of words over tones. *Child Development*, 81(2), 472–479. <https://doi.org/10.1111/j.1467-8624.2009.01408.x>
- Ferry, A. L., Hespos, S. J., & Waxman, S. R. (2013). Nonhuman primate vocalizations support categorization in very young human infants. *Proceedings of the National Academy of Sciences of the United States of America*, 110(38), 15231–15235. <https://doi.org/10.1073/pnas.1221166110>
- LaTourrette, A., Chan, D., & Waxman, S. (2023). A principled link between object naming and representation is available to infants by seven months of age. *Scientific Reports*, 13(1), 14328. <https://doi.org/10.1038/s41598-023-41538-y>
- LaTourrette, A., & Waxman, S. (2020). Naming guides how 12-month-old infants encode and remember objects. *Proceedings of the National Academy of Sciences of the United States of America*, 117(35), 21230–21234. <https://doi.org/10.1073/pnas.2006608117>

- Novack, M. A., Brentari, D., Goldin-Meadow, S., & Waxman, S. (2021). Sign language, like spoken language, promotes object categorization in young hearing infants. *Cognition*, 215, 104845.
- Perszyk, D. R., & Waxman, S. R. (2018). Linking language and cognition in infancy. *Annual Review of Psychology*, 69, 231–250.
- Spelke, E. S. (2022). *What babies know: Core knowledge and composition*. Oxford University Press.
- Woodruff Carr, K., Perszyk, D. R., Norton, E. S., Voss, J. L., Poeppel, D., & Waxman, S. R. (2021). Developmental changes in auditory-evoked neural activity underlie infants' links between language and cognition. *Developmental Science*, 24(6). <https://doi.org/10.1111/desc.13121>
- Xu, F., Cote, M., & Baker, A. (2005). Labeling guides object individuation in 12-month-old infants. *Psychological Science*, 16(5), 372–377. <https://doi.org/10.1111/j.0956-7976.2005.01543.x>

## Author's Response

### Response to commentaries on *What Babies Know*

Elizabeth S. Spelke\* 

Department of Psychology, Harvard University, and Center for Brains, Minds & Machines, Cambridge, MA, USA  
[spelke@wjh.harvard.edu](mailto:spelke@wjh.harvard.edu)

\*Corresponding author.

doi:10.1017/S0140525X24000049, e146

#### Abstract

Twenty-five commentaries raise questions concerning the origins of knowledge, the interplay of iconic and propositional representations in mental life, the architecture of numerical and social cognition, the sources of uniquely human cognitive capacities, and the borders among core knowledge, perception, and thought. They also propose new methods, drawn from the vibrant, interdisciplinary cognitive sciences, for addressing these questions and deepening understanding of infant minds.

Thanks to all the commentators for reading my book and sharing their thoughts. I wish that I could engage with all the points made by each commentary individually, but space constraints require condensing. This response begins with preliminaries (sect. R1) and then addresses five questions: Does the book's account of the origins of knowledge hit the right balance between innate structure and learning (sect. R2)? Are there core knowledge systems that it neglected (sect. R3)? How does language relate to core knowledge, learning, and other symbol systems (sect. R4)? How does core knowledge interface with perception and thought (sect. R5)? Going forward, what kinds of research might best advance understanding of the origins, growth, and nature of human minds (sect. R6)?

#### R1. Preliminaries

I begin by correcting some confusions that I caused. As Schulz notes, the book's definition of innateness is inadequate, so here is a better one: A cognitive system is innate if it is present and functional on the infant's first effective *informational* [formerly,

perceptual] encounters with the entities to which it applies. The corrected definition makes room for learning from text, television, or testimony and also for innate knowledge of entities for which prior perceptual encounters yielded no usable information. Schulz also notes an ambiguity in my argument for the abstractness of concepts with a long evolutionary history. I should have written that a core cognitive system that first emerged in distant ancestors *and is present in their distantly related, contemporary descendants* is more likely to center on abstract, broadly applicable concepts (like *plant*) than on more specific concepts (like *coconut*), because the former concepts will be useful for all descendant species, however much their environments differ.

**Moore & Lewkowicz** note that the précis contains no definition of learning. The book does (p. 71), but its definition is flawed by its focus on perceptual rather than informational encounters. A better definition considers a representational system to be learned if it is shaped by the learner's informational encounters with entities in its domain. Both revised definitions serve as invitations to research, especially research on controlled-reared animals and randomized experiments evaluating interventions to promote learning. In contrast, Moore & Lewkowicz's definition of learning, as "functional changes that result from experiences," is too broad to guide research in developmental cognitive science: Falls can cause functional changes to cognition but concussions are not learned.

Finally, in the book's chapter on number, I characterized infants' number representations as approximate but took no stand on whether their imprecision inheres in the core number system or in the systems with which it interfaces. I was less careful in the précis and used the expression, "the approximate number system," as if it were synonymous with the core number system. This usage is misleading, as research has not revealed whether the imprecision of numerical comparisons, made without counting, stems from limits to the core number system, to the perceptual systems that serve as its inputs, or to the systems for action planning, memory, and thought that support counting and other forms of exact enumeration. I will be more careful in this response.

## R2. Core knowledge, nativism, and empiricism

Learning is impossible without cognitive systems that support it. For classical empiricists, human learning depended on innate sensory systems and mechanisms for learning by association. *What babies know* proposes a richer initial structure, with domain-specific cognitive systems providing abstract concepts that guide learning and reasoning throughout life. Many commentaries argue this account is too lean, and three argue it is too rich. In this section, I respond to two commentaries in the first category and all the commentaries in the second, deferring the remaining calls for more inborn structure to sections R3 and R5.

According to **Margolis & Laurence**, an abundance of innate cognitive capacities is needed to account for the development of knowledge in our species, including a universal grammar, representations of basic metaphysical distinctions, modal operators that support planning and counterfactual reasoning, representations of exact small numbers, and cognitive mechanisms that regulate attention and memory. **Carey** further argues that the core knowledge systems are limited by their iconicity; she calls for one or more innate representational systems that are propositional and language-like. I begin with one of the capacities that Margolis

& Laurence call for – an innate modal operator representing possible states of affairs – and suggest that it builds on core knowledge.

Both in humans and in rats, place cells are activated during action planning, and research has focused on the planning process: When navigating rats arrive at an unexpected blockade in a familiar environment, they pause while place cells fire in sequences that correspond to different traversable paths to their goal. After generating representations of several paths, the animal chooses one and heads for the goal, as the précis describes.

Do the imagined place cell sequences represent *possible* routes to the navigator's goal? Consistent with that suggestion, distinctive neural activity accompanies the sequences that are generated before people or animals choose their next move. When rats pause at a choice point, the activation of place cell sequences occurs against a background of bursts of neural activity throughout the hippocampus, called sharp-wave ripples. In contrast, travel down real paths activates place cell sequences against a background of rhythmic theta waves. The coupling of sharp-wave ripples to the place cell sequences may allow animals to evaluate different action plans without creating false memories of travel on paths not taken. In humans, sharp-wave ripples also occur during recall of past events in episodic memory tasks (e.g., Norman et al., 2019), so they may serve to mark all simulated events that are detached from current reality, including but not limited to possible actions or states of affairs. If so, then further patterns of activity in the hippocampus or other brain structures might distinguish between representations of possible and past events.

These findings suggest, first, that representations of possibility indeed may be innate, as **Margolis & Laurence** argue, because they are evident in animals that have never taken any action (Farooq & Dragoi, 2019). Second, these representations build on core knowledge – the core place system, in this example – not on representations of a different kind. Third, contra **Carey**, the core place system is not wholly iconic, and representations of possibility are not wholly language-like: Both appear to straddle the divide between iconic and propositional representations, because each action plan has discrete components – representations of each place along the way – that together yield a holistic simulation of a course of action. However, action planning also requires more cognitive systems than the book discussed: In this example, one or more simulation engines, a process for searching through available options, and a process for distinguishing simulations from reality.

Three commentaries find the core knowledge hypothesis to be too strong. **Reid** claims that the capacities that I attribute to innate core knowledge can be explained, in part, by learning during the fetal period, in response to light sources outside the womb, but one consideration suggests only a limited role for visual learning during this period. Dark-reared rats and controlled-reared chicks perceive the world at birth at least as well as newborn human infants do, as **Vallortigara** described, but in most studies, they are confined to darkness prior to as well as after birth. Nevertheless, I salute Reid's efforts to bring cognitive research during the fetal period to bear on questions concerning the origins of knowledge. My hunch is that experience-dependent changes indeed shape the fetal brain, but the most important experiences are generated internally, by the systems of core knowledge.

Second, **Adolph & Schmuckler** reject the methods and conclusions of almost all the research discussed in the book. Our disagreements begin with their first sentence: "Inferences about

infant perception and cognition must be based on their observable behaviors because babies can't talk, they don't follow instructions, and their physiological responses may reflect different psychological processes than adults." With this dictum, they exclude all the methods and findings of research on animal behavior and in systems, cognitive and computational neuroscience. Their last phrase is true but argues for, rather than against, interdisciplinary research on the origins of knowledge. Functional brain imaging experiments reveal both common and diverging patterns of neural activity in humans, rodents, and chicks, as **Vallortigara** writes, and also in human infants, children, and adults. All these findings shed light on the origins and growth of human knowledge, as does work in many other fields, including philosophy, computer science, linguistics, and anthropology. The convergence in findings across these and other disciplines has been the single most exciting scientific development of my lifetime.

It is behavioral research in developmental psychology, considered on its own, that has sometimes suggested a misleading picture of infants' knowledge. For example, infants may careen down a sloping surface because they misperceive the slope; because they perceive the slope but don't know how to traverse it; because they know what actions are needed to traverse it but lack the strength or skill to perform them; or even because they know they lack these skills, as **Gweon & Zhu** or **Berke & Jara-Ettinger** might propose, and opt for trial-and-error learning. Decades of behavioral research, documenting infants' misadventures on slopes but committed to the position in the commentators' first sentence, have not distinguished these possibilities. The good news is that all these questions can be addressed by research in the interdisciplinary, developmental cognitive sciences.

I also dispute **Adolph & Schmuckler's** characterization of the arguments in the book as "rich interpretations based on lean 'looking-time' behaviors." First, looking isn't lean: As Gibson taught us, it is an exploratory behavior, guided by what infants know and aimed at extending their knowledge. It is highly useful for behavioral research on infants, who are limited actors but eager explorers. Second, the book's conclusions depend on evidence that extends well beyond any single behavior: All the cited findings concerning object representations, for example, are corroborated by research using exploratory reaching as the outcome measure (e.g., Clifton, Rochat, Litovsky, & Perris, 1991; von Hofsten & Spelke, 1985). Moreover, many of its findings are supported by research on precocious animals, using locomotion and other actions as measures. Third, infants' knowledge is abstract but it isn't rich: Their object representations, for example, fail to distinguish a toy duck from a shoe.

Nevertheless, I am grateful to my academic siblings for introducing the concept of affordances into this discussion. I also treasure their account of Gibson's approach to their own work and fully endorse her reaction: You should do more experiments and figure out what is going on! Does blinking at an approaching object reflect infants' perception of impending or possible events, or is it a reflex, triggered by detection of patterns in the 2D optic array? Do infants topple down slopes because they misperceive them, misjudge the actions needed to traverse them, or revel in trial-and-error learning? Decades of research on infants, conducted with an exclusive focus on overt behavior, have not distinguished these possibilities. Multidisciplinary research could do so.

Finally, psychophysical experiments differ from other experiments not by the outcome measures they use but by the research

strategy they follow: They leverage mathematical relationships between physical stimulation and perceptual experience to gain insight into perceptual mechanisms and processes. As a full-time faculty member, with her own infant laboratory, for only 3 years, Gibson lacked the time to extend her psychophysical studies of adults and children to infants. Soon after her mandatory retirement, however, Held, Teller, Banks, and others published studies using the simplest measures of selective looking – for example, did the baby look left or right? – in psychophysical experiments on visual development, and Keen (then Clifton), Werker, Kuhl, Aslin, and others used a related measure – head-turning – in psychophysical experiments on auditory localization and speech perception. Important insights will be lost if these rich literatures are rejected because of prejudices against the exploratory behaviors on which they rest.

**Moore & Lewkowicz** advance an analogy between the development of knowledge and the growth of organisms. I applaud their openness to insights from biology, but their use of this analogy obscures the specific (and considerable) challenges faced by every investigation into the nature, origins, and development of capacities to represent the world. Their claim that "...a developmental analysis must *a priori* consider all stimulation as potentially crucial unless its role has been empirically ruled out" sets an impossible standard for research: Because beating hearts are necessary for animal life, no studies on animals can remove all sources of stimulation. Instead, good experiments evaluate hypotheses against counterhypotheses that are reasonable, given the current state of knowledge. Setting a skewed standard, whereby empiricist claims are taken to be true unless proven false, introduces a bias that is especially harmful to empiricists, because experiments probing infants' learning become superfluous if their claims are deemed to be true in the absence of empirical support.

Nevertheless, the definitions and methodological standards that **Moore & Lewkowicz's** commentary articulates have been highly popular over the history of developmental psychology, and I thank them for giving me the chance to push back against them. The dismissal of all claims concerning the inherent, endogenously generated cognitive processes that give rise to innate knowledge, and guide learning, impeded the study of cognitive development for many decades. Fortunately, the interdisciplinary study of the origins and early growth of knowledge is now at an exciting time, as psychology, neurobiology, and other fields interact to discover how infants represent the world. I hope Moore & Lewkowicz enjoy the methods and insights that these synergies bring to our field.

### R3. Are there more systems of core knowledge?

Nine commentators argue for additional core knowledge systems. The elegant experiments by **Hespos & Rips** show that infants distinguish nonsolid substances from objects. Their findings raise the possibility that infants have core knowledge of substances, but they also are compatible with the hypothesis that knowledge of substances is learned. Research could distinguish these accounts by probing knowledge of nonsolid substances in controlled-reared animals or newborn infants. If infants are endowed with core knowledge of substances, interesting questions arise concerning its content and its relation to the core object system, as Hespos & Rips note. Strictly speaking, objects do not pass through nonsolid substances, they displace them. Do infants expect the height of a liquid to rise, or the shape of a sandpile to change, when an



object moves into the space that it occupies? Experiments on infants, animals, children, and adults, using multiple methods, could address these questions.

**Duval** argues for a system of place representation, focused on the appearances and affordances of landmarks, that is distinct from the core place system on which I focused. Indeed, animals and humans navigate by landmarks and use them to correct errors that will otherwise accumulate during path integration, as he notes. It is not clear whether the ability to navigate by landmarks arises from a domain-specific core knowledge system or from domain-general systems of perception and memory, but landmark-based navigation deserved more respect than my book gave it. Learning to recognize places by their distinctive landmarks is a key task for all navigators, and beautiful research illustrates how landmarks interplay with geometric place representations to enhance navigation performance, in creatures as distantly related as ants (Muller & Wehner, 1988) and human adults (e.g., Doeller & Burgess, 2008).

Drawing on the rich, subtle, and fascinating literature on infants' social reasoning, **Hamlin** and **Tatone & Pomiechowska** argue that infants have early emerging knowledge of social goals: Instrumental actions undertaken to benefit another agent. They point to studies in which younger infants selectively look or reach to characters who helped other characters without imitating them. In these studies, however, the reasons for infants' actions are not clear. First, infants may infer that such characters have helpful intentions, as Hamlin argues. Second, after seeing a protagonist repeatedly try and fail to achieve its goal, infants may view the character whose action completes the goal as a better agent, not a better social partner, because it is willing and able to move the action along. Third, infants may develop a social preference for characters who help others without understanding their acts of helping. Infants may deem such characters to be interesting, because they sometimes act on objects and sometimes engage with other characters, before they represent social actions like giving or helping.

Several considerations suggest that moral evaluations emerge through learning, rather than from an innate moral core (see also Hamlin, 2023; Spelke, 2023). First, infants' evaluations of helpers develop gradually in many experiments: For example, negative evaluations often emerge before positive ones. Moreover, as Hamlin notes, moral intuitions depend not on the acts we see but on actors' intentions: Pushing an agent up a hill is morally praiseworthy if the actor aims to fulfill the unsuccessful climber's goal, but not if she aims to push the climber out of her way. Because core knowledge builds on perception, it cannot readily distinguish between distinct construals of a single perceived action. Nevertheless, learning about predictive relationships between individual people's actions and their social behaviors may well influence infants' social evaluations, prior to the development of social agent concepts or true moral evaluation. Research could test these suggestions (see sect. R6).

Three commentaries argue for early emerging knowledge of triadic interactions involving the infant, a person who engages with him, and an object toward which the infant and his partner share attention. My book benefitted greatly from **Grossmann's** findings concerning young infants' brain responses to social others: His research provides evidence, I believe, for core knowledge of people as sentient beings who share their experiences in states of engagement. Does a further social distinction, between social engagements that are dyadic versus triadic, also emerge early in infancy? In the behavioral experiments he cites, young infants'

longer looking when a person's gaze shifts back and forth between them and an object could occur for multiple reasons, because objects and hands are potent elicitors of attention. The cited neuroimaging experiments make a stronger case for core knowledge of triadic interactions, but their findings also could depend on the distinct, attention-eliciting effects of direct gaze, gaze shifts to objects, and object motion. These possibilities may be hard to distinguish using fNIRS, given its relatively low spatial and temporal resolution, but future neuroimaging experiments using other methods could test them (see Ellis and sect. R6).

**Tauzin, Jacob, & Gergely** argue that infants endow agents with pragmatic communicative intentions in the absence of language. Two of their cited studies focus on young infants but are open to multiple interpretations. First, infants may view face-to-face interactions as social engagements rather than communicative exchanges: A speaker's utterance may be interpreted as a social overture (like "hello") and the listener's response, reproducing the speaker's action, may be viewed as social imitation (as in Powell & Spelke, 2013; although see Neff & Martin, 2023). Second, infants may learn, by 6 months, that a single speaker can be informative, efficient, or relevant to the current situation at different times within an episode, without integrating these three goals into a unitary communicative intention. These competing claims also should be testable, using methods that track the time course of signatures of action understanding and social engagement.

**Kaufmann & Clément** argue for an early-emerging system of intuitive sociology, focused on relations like shared group membership or dominance. I find this hypothesis highly worthy of testing. Kaufmann & Clément also note that learning about the social world bears a strong resemblance to learning about the navigable world. In both cases, the fundamental entities that populate these worlds are not geographical or social groups but individual places and people. Moreover, infants' learning of the specific social relationships that connect people resembles the learning, by navigating animals, of specific paths that connect places. This learning therefore could arise either from a core naïve sociology or from core knowledge of objects and places, together with a domain-general mechanism for representing the networks that connect distinct individuals within a domain (see sect. R5).

Finally, the commentaries by **Berke & Jara-Ettinger** and **Gweon & Zhu** point to an important gap in my book's coverage: Nowhere do I ask what babies know about themselves. This omission is not minor, because *all* the core systems apply to infants themselves. Infants cannot pass through walls or act on objects at a distance; they are living, agentive, and social beings; and their representations of number, geometry, and social relationships apply to themselves: Young children use their fingers in counting, retrieve information about their changing location as they navigate, and learn about relationships connecting them to other members of their social world (Thomas, Saxe, & Spelke, 2022). I am grateful to these commentators for signaling this gap in the book's coverage.

Conceptions of the self are hard to study: When the infants in Thomas' experiments see one of the two puppets comforting their mother, for example, do they learn something about themselves (*I am close to this puppet*), about the puppet (*this puppet is close to me*), about their family as a group (*we are close*), or all three? **Berke & Jara-Ettinger** suggest, however, that self-knowledge can be studied by focusing on situations that distinguish our experience of the world from the world itself, such as perceptual illusions, and **Gweon & Zhu** appeal to research by Stahl and Feigenson as

evidence for infants' knowledge of their own violated expectations concerning the behavior of objects. Do these situations reveal unlearned representations of the self from the beginning?

Against that possibility, 1-year-old infants' exploration of objects that have behaved in surprising ways suggests that they are more focused on learning about the objects than on learning about themselves. When adults are presented with magic shows, we assume that we failed to detect the sleight of hand, and we may increase or redirect our attention during the next act. The infants in Stahl and Feigenson's studies don't respond in this manner, however; instead, they vary their actions on the object to test for its properties. They may do this for good reasons: Infants are great explorers and learners, and so their own minds are moving targets, given the rapid growth of their knowledge and skills. Thus, I lean toward the view that explicit knowledge of the self is learned, and it grows as infants discover their own, and other people's, diverse and changeable perspectives on the world. But this topic deserved discussion in *What babies know*, and I'm inspired, for the sequel, to think more about it. As long-standing research by Rochat (2018) and others has shown, implicit self-knowledge either begins to be learned in infancy or is present from the beginning.

#### R4. Core knowledge and language

Goldin-Meadow appeals to research on deaf children of hearing parents, who communicate by homesign, to challenge the book's claims for language learning as the primary process that carries infants beyond core knowledge and toward richer conceptions of people as social agents. Although I took no stand on how infants learn their language, studies of homesign accord with Carey's suggestion that language acquisition depends on a seventh domain-specific system of core knowledge. If homesigners' invented language is accompanied by the emergence of knowledge of individual actions as simultaneously social and object-directed, like gift-giving, then Goldin-Meadow wins, and language attainment, not language learning, underlies infants' concepts of social actions.

However, I argue (in ch. 10) that children's fast and flexible cultural learning depends on learning of an established conventional language, because such languages are shaped by generations of speakers who have aimed to communicate efficiency, informatively, and relevantly. In such a language, words that convey the culture-specific concepts that provide the most useful perspectives on the world will tend to be short and frequent, and general learning mechanisms, biased to learn from simpler and more frequent events, will orient children toward their culture's most useful concepts in ways that language invention cannot. I predict, therefore, that homesigning children will be slower, less effective cultural learners than are children who learn an established conventional language. Consistent with this prediction, some studies of homesigning adults, and of adult speakers of emerging sign languages, provide evidence for limited mastery of culturally variable concepts of exact numbers (Spaepen, Coppola, Spelke, Carey, & Goldin-Meadow, 2011, 2013), spatial relationships (Pyers, Shusterman, Senghas, Spelke, & Emmory, 2010), and mental states (Pyers & Senghas, 2009).

My chapters on language and the construction of new concepts appealed to Waxman's research with older infants and toddlers, showing that language allows them to characterize the same perceived entity in diverse ways: For example, as a car, a Fiat, or a blue thing. She also finds effects of language on 3-month-old

infants' categorization of diverse pictures of dinosaurs, but this finding is open to multiple interpretations, as she notes. For example, speech and other vocalizations, like lemur calls, may alert infants to the presence of something worth attending to, allowing focused perceptual analysis of object properties. I hope future research will chart the emergence of sensitivity to the different perspectives on objects that language can convey. I hypothesized, in chapter 10, that infants begin to distinguish these perspectives at about 12 months, when they first come to view others' mental states as both phenomenal and intentional. Consistent with this hypothesis, word learning accelerates and becomes more confident in the first months of the second year (Bergelson, 2020).

Against my account of core knowledge and language as the sources of our uniquely human concepts and cognitive skills, Schulz argues that talents for cultural learning and symbol use are the foundations of uniquely human cognitive accomplishments. I agree that these talents are key signatures of distinctively human intelligence, but what inherent talents underlie them? Young children readily learn their native language, but other symbol systems, including pictures, toy cars, maps, and graphs, are harder for children to master (DeLoache, 2004), unless the symbols are accompanied by language (Winkler-Rhoades, Carey, & Spelke, 2013). Humans learn diverse symbol systems, I suggest, because we are endowed with core knowledge and with an inherent attunement to the symbolic functions of language.

Margolis & Laurence pushed back against my argument that language learning reverses the curse of a compositional mind and claim there is no such curse: The richer the child's innate conceptual repertoire, the easier concept learning will be. I do not claim that possession of a combinatorial system that generates all humanly attainable concepts will make concept learning impossible in the absence of language learning: A machine with infinite time will find every concept in it. Concept learning will be slow if concepts are drawn from a rich mental language in a culture with no public language, however, because if they are sampled in a manner that allows learning in any culture, then children will need to search through a vast space of combinations to find the right ones. In contrast, finding the right concepts in a vast repertoire becomes more manageable if concept learning is informed both by core knowledge, which applies to all cultures and environments, and by speech in the conventional language of one's culture, which provides multiple cues to the concepts that its members find most useful.

In the book, I argued that core knowledge is never expressed in language, because it is universal (and therefore can be left unsaid) and unconscious (and therefore cannot be explicitly accessed). Contrary to the second claim, Lin & Dillon's elegant experiments suggest that ordinary language activates representations from core knowledge: When two connected line segments are described as an incomplete object, people connect them; when described as an incomplete path or abstract pattern, they extrapolate the pattern. As they note, it isn't clear whether these verbal descriptions activate core knowledge directly or indirectly, by activating explicit concepts that build on core knowledge. In either case, however, this activation may explain why games that connect images of sets or forms to numerical and geometric language and symbols have stronger, more enduring effects on children's math learning than does play with language and symbols alone (Dean, Dillon, Duflo, Kannan, & Spelke, 2023; Dillon, Kannan, Dean, Spelke, & Duflo, 2017). I look forward to research probing the links connecting core knowledge to language and endorse Lin & Dillon's

conclusions concerning its potential implications for education, economics, and the developmental cognitive sciences.

Finally, language learning may address an important problem raised by **Kaufmann & Clément**, who appropriately distinguish between concepts of individuals, like a particular place, object, plant, or animal, and second-order abstract properties like number and geometry. The latter concepts are radically ambiguous: “How many things are in this room?” is a question with different answers, depending on what things one counts. Solving this problem, they suggest, requires either that different core systems form a multilayered hierarchy, or that the core systems of number and geometry aren’t as abstract as my discussion in *What babies know* implies. I favor the second option: The core systems of number and geometry are limited. Although they represent numerical magnitudes and geometric shapes that can be transformed by operations of arithmetic and geometry, and they support children’s learning of, and adults’ reasoning about, mathematics, they lack the power of explicit mathematical concepts.

The core number system differs from our most intuitive explicit number system – the integers – in at least two ways. First, it fails to support exact enumeration, either because it is inherently approximate or because it interfaces with systems of limited resolution. Second, because core knowledge systems compete for attention, infants can activate representations of objects or numerical magnitudes but not both at once, so the core number system fails to capture the hierarchical structure that allows for representations of *three people*, *three animals*, or *three shoes*. Experiments with navigating animals, and with human adults performing a virtual navigation task, reveal a similar failure to capture, in a unitary representation, the positions of environmental boundaries and of landmark objects. Our core systems of number and geometry therefore do not operate in accord with the hierarchical organization of our mature mathematical concepts. Language, however, is a hierarchically structured, symbolic system, so it might provide a medium for constructing mature representations of number and geometry: topics for my next book.

## R5. Core knowledge, perception, and thought

If my aim, in writing *What babies know*, were reduced to a single sentence, it would be this: To make the case for a level of representation beyond perception and thought, and for the critical contributions of domain-specific cognitive systems at this level to the origins, growth, and nature of human knowledge. Many commentators are not convinced, and for a good reason: The book did not devote nearly enough space to cognitive systems and processes at the borders between core knowledge, perceiving, and thinking.

**Carey** argues that both perceptual and core knowledge representations are iconic, or image-like, whereas conceptual representations are discursive, or propositional and language-like. She concludes that cognitive systems come in two rather than three kinds: perceptual and conceptual systems. I dispute these claims. First, the best theories of perception, dating back to Helmholtz, appeal to representational systems that are both iconic *and* propositional, as they depend on unconscious inferences concerning the sources of sensory experience. Similarly, the contemporary theories of vision described in the book and précis appeal to graphics engines: Generative models, written in an internal programming language, that produce holistic 2D images of the light-reflecting surfaces in a 3D scene. If these theories are correct, then perceptual systems combine iconic and discursive representations (see also Quilty-Dunn, Porot, & Mandelbaum, 2023).

Following Ullman, Spelke, Battaglia, and Tenenbaum (2017), I suggested that core knowledge systems also function as generative models of the entities in their domains: Models that support learning when they are run in the forward direction and that support inference both when run forward and when inverted. If that is true, then the representations formed by core knowledge systems also are both iconic and discursive, contrary to the dichotomy urged by Carey and by Block (2022). Finally, investigators from diverse perspectives have long argued that explicit thought, memory, and prospection involve both mental simulation and language-like reasoning (see Kosslyn, 1980; Shepard, 1984). This argument has been made by investigators with diverse views on the nature of the representations that guide reasoning and simulation, but one family of accounts appeals to the same sorts of generative models as those proposed to underlie innate capacities for visual perception and core knowledge (e.g., Tenenbaum, Kemp, Griffiths, & Goodman, 2011; Ullman, 2015; Ullman & Tenenbaum, 2020).

In sum, perception, thought, and core knowledge all may depend on image-like simulations produced by language-like representations and computations. What, then, distinguishes core knowledge from perception and thought? First, our perceptions and beliefs change with experience: Older children are better at distinguishing faces and recognizing known others than younger ones, and chess masters see relationships on a chess board and plan moves that novices miss, as **Krøjgaard, Sonne, & Kingo** (Krøjgaard et al.) notes. Core knowledge, by contrast, shows the same signature limits in adults and infants. Second, perception and belief depend in part on unconscious processes but give rise to conscious experience, whereas core knowledge is wholly unconscious: We can become aware of its existence and properties by doing experiments but not by introspecting. Third, as **Vallortigara**’s commentary argues, the plastic neocortex is the seat of our conscious, malleable perceptions, memories, planning, concepts, and beliefs, whereas core knowledge in animals and infants likely depends on subcortical brain systems that are hard-wired and impenetrable.

**Scholl** raises crucial questions concerning the relationship between core knowledge and visual perception of objects, scenes, causal relations, and animacy. For decades, he and his associates have built on classic experiments – such as Michotte’s studies of the perception of causality and Heider and Simmel’s study of perception of animacy – to probe the processes by which we perceive these phenomena. Although perception itself is conscious, his studies reveal that many perceptual phenomena depend on processes that are unconscious and automatic. Moreover, the representations that they deliver are strikingly similar to those found in studies of infants, even though the experiments themselves differ. Scholl’s commentary provides the simplest, most natural explanation for this consilience: Core knowledge systems *are* perceptual systems.

Although **Scholl** notes that the consilience is not complete, it may be even stronger than his commentary argues, as he focuses on wonderful findings from foraging and predator-detection tasks but no findings bearing on perceptual knowledge of places or people. The occipital place area (OPA) of the visual cortex is consistently activated by visual scenes, it is specifically attuned to the navigational affordances of these scenes, such as an open doorway in a room (Bonner & Epstein, 2017), and it is causally involved in human reorientation (Julian, Ryan, Hamilton, & Epstein, 2016). Studies of other regions in the occipital lobe connect mid- or high-level visual perception to core social cognition: Not only



to representations of agents, as Scholl described, but also to representations of individual people and their relationships: crucial information for navigating the social world. Although Scholl notes that no visual area specializes in representing acts such as helping or gift-giving, earlier-emerging aspects of social knowledge may well have counterparts in the multiple visual areas that are sensitive to faces, as **Carey** describes.

There are two ways to think about these findings. First, my book may be misnamed: All the abilities that I find in infants may depend on perceptual systems, not systems of knowledge. It is hard to construe navigation or social reasoning as perceptual abilities, however, for several reasons. First, navigation builds on enduring representations of unseen destinations and of the paths that led from one place to another in the past, and social reasoning builds on enduring representations of the imperceptible bonds that connect one person to another. We may learn how two people are related by observing their interactions, but we store this information in a way that respects their previously experienced social relationships. Second, both these tasks are associated with activity in the hippocampus: A structure that is commonly associated with memory and action planning and is rarely considered to be a perceptual system. For these reasons, I favor the other possibility that **Scholl's** discussion invites: Core knowledge lies between or beneath our perceptions, memories, beliefs, and plans. If this view is correct, then perceptual signatures of core concepts, like the retinotopic adaptation effects of repeated viewing of causal interactions, would not be explained by adaptation of the core knowledge system but by adaptation of the cortical, perceptual representations that it activates.

If that's true, then the processes occurring at the interfaces of core knowledge with perception and thought beg for further study, and the commentary by **Kaicher, Conti, Dedhe, Aulet, & Cantlon (Kaicher et al.)** suggests ways to study them. Kaicher et al. argue for cognitive systems with all the properties of core knowledge systems except two: The systems crosscut the core domains on which the book focused, and they contribute to the efficiency and adeptness with which we perceive and reason, rather than to the content of knowledge. For example, the system that gives rise to categorical perception applies to entities in diverse domains, but it operates automatically and unconsciously, as do the systems of core knowledge.

I think there are good reasons to distinguish domain-specific systems of knowledge from domain-general systems for optimizing cognitive processing, but I agree that both types of systems are needed to explain how minds work, at all ages. Categorical perception reflects a fundamental property of perceptual systems and perceptual learning; it therefore serves to foster both the rapid identification of the entities in each core domain as well as abilities to distinguish one such entity from another and to learn about each entity's properties and behavior: Crucial skills for navigating the geographical world, as **Duval** argued, the social world of people and their relationships, as **Kaufmann & Clément** argue, and the world of object kinds, with distinctive forms and functions, as **Liu & Xu's** research reveals. By increasing the efficiency of perceptual processing, categorical perception enhances learning at the interface of perception with core knowledge.

I agree that the development of knowledge requires more cognitive mechanisms at the interface between core knowledge and thought. **Krøjgaard et al.** usefully propose that mental rotation, episodic memory, and chess playing depend on such systems. In section R2, I argued that cognitive mechanisms for simulating events (including rotations of objects), and for distinguishing

simulations from reality, are needed to account for core knowledge, perception, memory, and reasoning. Episodic memory depends on the hippocampus and likely is activated and strengthened when humans or animals review past events. Chess playing requires both perceptual mechanisms permitting rapid analysis of relationships between the pieces on the board, and mechanisms of action planning allowing comparisons among sequences of possible moves. These processes operate at the interfaces of core knowledge with perception and with memory, prediction, and planning.

These considerations suggest an alternative account of the findings with which **Jenkin & Markson** and **Liu & Xu** challenge two of my central claims for core knowledge: The claim that core knowledge is constant over development and unaffected by later-emerging beliefs or attitudes, and the claim that the core knowledge systems form a natural kind, with common functions and modes of operation. Against both claims, Jenkin & Markson argue that the core social system is revisable, and Liu & Xu argue that all the core systems are revisable, based on their findings that preschool children develop beliefs at odds with core knowledge of objects and agents when given verbal questions about events that are similar to those providing evidence for core knowledge in infants. Liu & Xu also argue that the object and number systems are harder to revise, as are perceptual systems, whereas the agent and social systems are more readily revised, as are belief systems. These arguments cut to the heart of the claims made in *What babies know*.

Adults' and children's explicit beliefs about the entities singled out by core knowledge, and the actions that we perform on the entities in core domains, are various, changeable, and sometimes at odds with core knowledge. Nevertheless, the hypothesis of constant core systems, interfacing with malleable systems for perceiving and thinking, seems far more plausible to me than the hypothesis that core knowledge itself is revisable. For example, mathematicians develop concepts of complex numbers and high-dimensional spaces. Where research has been conducted, however, their reasoning about difficult problems in mathematics has activated the same systems of core knowledge as are activated by numerical or geometric tasks in children and ordinary adults (e.g., Amalric & Dehaene, 2016).

Thus, I differ from **Liu & Xu** over the sources of our malleable, flexible reasoning. Infants reason flexibly, I submit, because the properties captured by each impenetrable core knowledge system apply to all habitable natural environments. For example, the core agent system applies to all actions that agents can undertake: Not just reaching for objects and locomoting to places, but reaching to places and locomoting to objects. That is why the 3-month-old infants in the study cited by Liu & Xu learned with equal facility that the goal of a reach was an object or a place, depending on the evidence that they received. The flexibility of children's learning about agents and their goals is consistent, moreover, with the impenetrability of core knowledge: Impenetrable core knowledge systems can promote flexible learning throughout life, because the abstract properties they capture apply to all the entities in their domain, for people of all ages, in all environments. Humans go beyond core knowledge, and even contradict it, when we develop explicit, learned beliefs, because beliefs are malleable. Core knowledge is not malleable, however, and so it continues to function, despite these beliefs.

**Jenkin & Markson's** commentary, focused on the core social system, raises a further question about the interface of this system with action and thought: How and why do children come to organize their social knowledge in accord with social categories of

people, based on attributes like race, gender, or social class, developing biases toward or against individuals in these categories, if the core social system applies to all potential social beings and supports learning about individuals of all races (Kinzler & Spelke, 2011) and, indeed of many species (e.g., Pascalis, de Haan, & Nelson, 2002)? I believe these effects stem from the mechanism of categorical perception discussed by **Kaicher et al.**: As children gain increasing exposure to social beings of a familiar race, they come to perceive faces of that race more clearly, and information conveyed by the same-race face comes to them more vividly and rapidly: They quickly see, for example, that a pictured face of a familiar race looks happy or scared, whereas the emotion expressed by a pictured face of a less familiar race will be seen more slowly and less vividly. Later in development, responses to individuals of differing races may be modulated by explicit beliefs about the characteristic attributes and behaviors of people in different groups. At no time, I believe, will core knowledge of agents or social beings change. Instead, changes in racial attitudes likely depend on the plasticity of perception and thought. These are largely untested predictions, however. In the final section, I ask how future research might serve to shed more light on the development of knowledge in infancy.

### R6. Beyond *What babies know*

*What babies know* omitted important questions. Throughout my research on infants' knowledge of number, for example, I wondered whether the core number system was exact or approximate, but I concluded that the question was not answerable by current methods: An inability to distinguish *six* from *seven* could stem from limits either to the core number system or to the perceptual, memory, and action systems with which it interfaces. The lack of evidence bearing on this question was not a reason to avoid discussion of the question, however: On the contrary, such discussion is a needed prelude to research. A second omission occurred in chapter 10, where I proposed (a) that infants' language learning brings them a new concept of people as *social agents*, based on evidence for the development of new concepts of *social actions* like helping and gift-giving, and (b) that in the absence of language learning, domain-general learning processes support a weaker understanding of social agents as beings whose specific actions on objects may predictably precede or follow specific social gestures. I stand by these claims but regret my failure to consider how language learning, and language-independent predictive learning, might combine to account for developmental changes in social cognition, through processes occurring at the border between core knowledge and thought.

In general, the commentaries have prompted me to think more about the interfaces between perception, thought, and core knowledge. At many points in the book, I suggested that neuroimaging studies of human infants could shed light on open questions about the origins of knowledge in infancy, but I never mentioned experiments using the methods of multivariate pattern analysis (MVPA), including representational similarity analysis (RSA). Although other ways of analyzing brain data have shed light on diverse aspects of numerical representations, some questions have proved exceedingly hard to answer, despite a rich body of research using the methods of psychophysics and cognitive neuroscience on adults, children, and infants. **Ellis** focused on one such question: Do infants, children, and adults represent number *per se*, or do they represent continuous variables that correlate with number, such as continuous spatial extent or temporal duration?

**Ellis** proposed that neuroimaging experiments using MVPA in general, and RSA in particular, can address this question. Indeed, a recently published study, conducted by Dehaene-Lambertz and her collaborators, embraced this challenge. Gennari, Dehaene, Valera, and Dehaene-Lambertz (2023) used MVPA to decode for number in 3-month-old infants, using high-density electroencephalography (EEG). As infants rested or slept, they heard sequences of tones varying in number (4 or 12), sequence duration, tone duration, and also tone frequency and timbre. Gennari et al. trained a decoder to identify and distinguish between sequences of 4 and 12 tones, using input from 256 sensors on the baby's scalp during brief intervals that followed the end of each sequence. During training, the decoder was presented with a critical subset of these intervals, chosen such that successful discrimination of the two numbers required that it ignores differences in sequence duration, tone duration, and the other variables. After training, the decoder reliably distinguished 4- from 12-tone sequences in the remaining data, providing evidence for representations of number in the infant brain.

Using RSA, moreover, Gennari et al. found that the response to number that was trained on the auditory sequences, heard during sleep, generalized with no further training to stationary visual arrays that the infants viewed while awake, before or after the sleep session. The latter finding accorded with earlier behavioral research on newborn infants (Izard, Sann, Spelke, & Streri, 2009) but went beyond it, because in this study, tones of different numbers were randomly intermixed, engendering no expectations that either number would be repeated, and infants were tested while drowsy or sleeping, allowing tests for generalization not only over changes in modality but also over changes in the infant's state.

Further analyses by Gennari et al. speak on **Ellis'** discussion of the vexed question of whether infants respond to number, or to continuous variables that correlate with number. They showed that the infants' brain signals allowed not only for successful decoding of number in the tone sequences but also for successful decoding of sequence duration and tone duration. Moreover, their analyses showed that infants' brain responses to the two duration variables were independent of their brain responses to number. Thus, the experiment provided evidence that 3-month-old infants represent number *and* two aspects of duration. It is hard to imagine how the independence and robustness of infants' numerical representations could have been tested without these methods and analyses. Using MVPA and RSA, Gennari et al. discovered a signal in the infant brain that is specific to number, as well as signals specific to other quantities.

Might similar experiments resolve the question of whether the core number system delivers an approximate or exact representation of number? Further RSA analyses by Gennari et al. showed that a decoder, trained on the interval that followed the third and the seventh tone in the 12-tone sequences, came to distinguish between them, even though the sleeping infants never heard sequences consisting of three or seven tones. This finding suggests that infants increment their numerical representations after each tone in a sequence, raising the possibility of exact numerical representations at each step in the sequence. Against this conclusion, the decoder failed to discriminate the third from the fifth tone, but this failure could have occurred because the 4- and 12-tone sequences that the infants heard all differed by a large ratio, encouraging a focus on approximate numerical magnitudes. I hope future studies will use these methods to test for exact number representations in infants. If the core number system is exact and error is introduced by perceptual and memory

systems, then the brains of infants, presented in sleep with sequences of five versus six tones, for example, might produce a signal that a coder could be trained to detect, independently not only of continuous variables but also of the signals generated by other brain systems.

If decoders can use data from functional brain imaging to decode for numerical representations, then they might also serve to decode for representations of social actions and of the people who engage in them. Using these methods, investigators could focus on the changes in brain activity that occur with changes in children's understanding of social and communicative actions, and of the mental states of the people who perform such actions. As babies begin to learn that people's social gestures and object-directed actions occur in regular patterns, does this learning change their conceptions of people as social agents, or are concepts of people invariant over development, with changes only to infants' understanding of the actions they perform? Neural recordings, analyzed by MVPA and RSA, might address these questions.

Another potential avenue for advancing understanding of cognition in infancy is suggested by **Vallortigara's** commentary. He noted that newly hatched chicks, which respond to objects, places, and agents similarly to human infants, represent objects primarily via a midbrain structure, the optic tectum, that is a homologue of the human superior colliculus, which is also a subcortical brain structure. In the book, I speculated that all the core knowledge systems reside in subcortical structures, from which their activity (prenatal and postnatal) propagates to the plastic cortical regions underlying perception, memory, and learning. To my knowledge, no one has imaged subcortical activity in chicks or human infants during performance of tasks providing evidence for core knowledge. Such studies might shed light on the operation of the core knowledge systems and their constancy versus malleability by experience. Further studies of the activity that propagates from subcortical systems to the cortex then could address open questions concerning the interface between core knowledge, perception, and thought.

Beyond research in cognitive neuroscience, advances promise to come from field research, conducted in diverse countries and cultures, and following infants over extended timespans, as **Lin & Dillon** discussed. Still more advances may come from studies of infants whose everyday experiences differ from those of the infants who most often are studied by developmental cognitive scientists: For example, infants who learn their language from overheard speech, because adults in their culture do not speak to infants; infants with limited vision; infants with exceptional abilities like absolute pitch; or deaf children who invent their own language, as **Goldin-Meadow** argued. Finally, I look forward to insights from computational cognitive science, leveraging the data from experiments testing large samples of infants recruited through online platforms, large-scale field experiments, or collaborative replications of classic findings. Analyses of data from these sources could serve to evaluate diverse computational models of infant cognition and learning (e.g., Gandhi, Stojnik, Lake, & Dillon, 2021), including the probabilistic generative models discussed in the book. *What babies know* featured few experiments on infants using any of these methods, leaving rich territory for future books to explore.

**Acknowledgments.** The errors in this response are mine alone, but I thank Ghislaine Dehaene-Lambertz, Stanislas Dehaene, Sam Gershman, and Rachel Keen for advice and helpful information.

## References

- Amalric, M., & Dehaene, S. (2016). Origins of the brain networks for advanced mathematics in expert mathematicians. *Proceedings of the National Academy of Sciences*, 113(18), 4909–4917.
- Bergelson, E. (2020). Why do older infants understand words better? *Child Development Perspectives*, 14(3), 142–149.
- Block, N. (2022). *The border between seeing and thinking*. Oxford University Press.
- Bonner, M. F., & Epstein, R. A. (2017). Coding of navigational affordances in the human visual system. *Proceedings of the National Academy of Sciences of the United States of America*, 114, 4793–4798.
- Clifton, R. K., Rochat, P., Litovsky, R. Y., & Perris, E. E. (1991). Object representation guides infants' reaching in the dark. *Journal of Experimental Psychology: Human Perception and Performance*, 17(2), 323–329.
- Dean, J. T., Dillon, M. R., Duflo, E., Kannan, H., & Spelke, E. S. (2023). Number and geometry games combining symbols with intuitive material durably enhance poor children's learning of first-grade mathematics. *Unpublished manuscript, Abdul Latif Jameel Poverty Action Lab (J-PAL) South Asia, New Delhi, India*.
- DeLoache, J. S. (2004). Becoming symbol-minded. *Trends in Cognitive Sciences*, 8(2), 66–70.
- Dillon, M. R., Kannan, H., Dean, J. T., Spelke, E. S., & Duflo, E. (2017). Cognitive science in the field: A preschool intervention durably enhances intuitive but not formal mathematics. *Science*, 357(6346), 47–55.
- Doeller, C. F., & Burgess, N. (2008). Distinct error-correcting and incidental learning of location relative to landmarks and boundaries. *Proceedings of the National Academy of Sciences*, 105(15), 5909–5914.
- Farooq, U., & Dragoi, G. (2019). Emergence of preconfigured and plastic time-compressed sequences in early postnatal development. *Science*, 363(6423), 168–173.
- Gandhi, K., Stojnik, G., Lake, B. M., & Dillon, M. R. (2021). Baby intuitions benchmark (BIB): Discerning the goals, preferences and actions of others. *Advances in Neural Information Processing Systems 34*. Conference Proceedings (pdf, Supplement) (arXiv:2102.11938).
- Gennari, G., Dehaene, S., Valera, C., & Dehaene-Lambertz, G. (2023). Spontaneous supra-modal encoding of number in the infant brain. *Current Biology*, 33(10), 1906–1915.
- Hamlin, J. K. (2023). Core morality? Or merely core agents and social beings? A response to Spelke's *What babies know*. *Mind & Language*, 38(5), 1323–1335.
- Izard, V., Sann, C., Spelke, E. S., & Streri, A. (2009). Newborn infants perceive abstract numbers. *Proceedings of the National Academy of Sciences*, 106(25), 10382–10385.
- Julian, J. B., Ryan, J., Hamilton, R. H., & Epstein, R. A. (2016). The occipital place area is causally involved in representing environmental boundaries during navigation. *Current Biology*, 26(8), 1104–1109.
- Kinzler, K., & Spelke, E. S. (2011). Do infants show social preferences for people differing in race? *Cognition*, 119(1), 1–9.
- Kosslyn, S. M. (1980). *Image and mind*. Harvard University Press.
- Muller, M., & Wehner, W. (1988). Path integration in desert ants. *Proceedings of the National Academy of Sciences of the United States of America*, 85, 5287–5290.
- Neff, M. B., & Martin, A. (2023). Do face-to-face interactions support 6-month-olds' understanding of the communicative function of speech? *Infancy*, 28, 240–256.
- Norman, Y., Yeagle, E. M., Khuvis, S., Harel, M., Mehta, A. D., & Malach, R. (2019). Hippocampal short-wave ripples linked to visual episodic recollection in humans. *Science (New York, N.Y.)*, 365, eaax1030.
- Pascalis, O., de Haan, N., & Nelson, C. A. (2002). Is face processing species-specific during the first year of life? *Science*, 295(5571), 1321–1323.
- Powell, L. J., & Spelke, E. S. (2013). Preverbal infants expect members of social groups to act alike. *Proceedings of the National Academy of Sciences of the United States of America*, 110(41), 3965–3952.
- Pyers, J. E., & Senghas, A. (2009). Language promotes false belief understanding: Evidence from learners of a new sign language. *Psychological Science*, 20(7), 805–812.
- Pyers, J. E., Shusterman, A., Senghas, A., Spelke, E. S., & Emmory, K. (2010). Evidence from an emerging sign language reveals that language supports spatial cognition. *Proceedings of the National Academy of Sciences of the United States of America*, 107(27), 12116–12120.
- Quilty-Dunn, J., Porot, N., & Mandelbaum, E. (2023). The best game in town: The emergence of the language-of-thought hypothesis across the cognitive sciences. *Behavioral and Brain Sciences*, 46, e261.
- Rochat, P. (2018). The ontogeny of human self-consciousness. *Current Directions in Psychological Sciences*, 27(5), 345–350.
- Shepard, R. M. (1984). Ecological constraints on mental representation: Resonant kinematics of perceiving, imagining, thinking and dreaming. *Psychological Review*, 91, 417–447.
- Spaepen, E., Coppola, M., Spelke, E. S., Carey, S., & Goldin-Meadow, S. (2011). Number without a language model. *Proceedings of the National Academy of Sciences of the United States of America*, 108(8), 3163–3168.
- Spaepen, E., Flaherty, M., Coppola, M., Spelke, E. S., & Goldin-Meadow, S. (2013). Generating a lexicon without a language model: Do words for number count? *Journal of Memory and Language*, 69(4), 496–505.



- Spelke, E. S. (2023). Core knowledge, language learning and the origins of morality and pedagogy: Reply to reviews of *What Babies Know*. *Mind and Language*, 38(5), 1336–1350.
- Tenenbaum, J. T., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure and abstraction. *Science (New York, N.Y.)*, 331(6022), 1279–1285.
- Thomas, A. J., Saxe, R., & Spelke, E. S. (2022). Infants infer potential social partners by observing the interactions of their parents with unknown others. *Proceedings of the National Academy of Sciences*, 119(32), e2121390119.
- Ullman, T. D. (2015). *On the nature and origin of intuitive theories: Learning physics and psychology* [Doctoral dissertation]. MIT.
- Ullman, T. D., Spelke, E. S., Battaglia, P., & Tenenbaum, J. B. (2017). Mind games: Game engines as an architecture for intuitive physics. *Trends in Cognitive Sciences*, 21(9), 649–665.
- Ullman, T. D., & Tenenbaum, J. B. (2020). Bayesian models of conceptual development: Learning as building models of the world. *Annual Review of Developmental Psychology*, 2(1), 533–558.
- von Hofsten, C., & Spelke, E. S. (1985). Object perception and object-directed reaching in infancy. *Journal of Experimental Psychology: General*, 114(2), 198–212.
- Winkler-Rhoades, N., Carey, S., & Spelke, E. S. (2013). Two-year-old children interpret abstract, purely geometric maps. *Developmental Science*, 16(3), 365–376.