

I Although there are multiple ways in which young children can use language, categorize objects, use tools, and play games, they seem to quickly realize that there is one preferred, or conventional, way to do each of these things. These disparate domains share a common conventional structure.

Conventionality and Cognitive Development: Learning to Think the Right Way

Charles W. Kalish, Mark A. Sabbagh

There are many facts about the world that young children can discover for themselves. For instance, after a certain amount of experience in the world, children might discover that objects are solid, collide with one another, and are subject to the force of gravity. There is a correspondingly large body of facts that we cannot learn from direct observation and instead can be gleaned only from other human agents (Koenig & Harris, 2005). Some of these facts are naturally stipulated (or objective) but difficult for most people to ascertain. For instance, most of us have never seen a spleen and instead have relied on others' reports to tell us where it is in the body, what it looks like, and what it does. Some other facts that we get from other people are overtly subjective. A medical professional might offer an opinion that the spleen is the most elegant organ in the human body. In these cases, the message tells us more about the communicator's beliefs than something that is necessarily correct or true about world.

The focus of this volume is a sort of middle ground between these two kinds of socially gleaned knowledge. We will refer to this middle ground as *conventional knowledge*. Conventional knowledge differs from naturally stipulated, objective knowledge in that there is no way to verify our beliefs independent of other people. Yet conventional knowledge is not subjective in the usual sense; it is supposed to be a truth about the world that can be evaluated as such. In conventional knowledge, the standard by which the truth

of a proposition in a conventionalized domain can be evaluated is whether there is agreement about a particular form among various social agents. Put another way, conventional knowledge is neither strictly objective nor subjective. Rather, it is intersubjective: social sources transmit this knowledge and arbitrate the appropriateness of its forms (see Searle, 1995).

Perhaps the paradigm case of conventional knowledge is language. As Ferdinand de Saussure (1983) first made clear, the relation between words and their referents is essentially arbitrary. For the most part, there is nothing about a given object, event, or concept that indicates how it should be named. It is only through human intention and action that some word comes to stand for one object rather than another. Yet in most situations, we do not consider the relation between words and their referents to be subjective in the usual sense. That is, when someone calls a cup a “cup,” we tend to treat this as a general fact about what the cup is called, as opposed to the person’s idiosyncratic opinion or practice of naming the object. Although it is arbitrary, the fact that others agree on its use means that there is a standard by which use of the word-referent link can be judged. Using another word to refer to a cup is not just different or unusual (the way, say, thinking the spleen is the most elegant organ might be); it is wrong (see Chapters Two and Three, this volume).

Language is not alone as a domain of conventional knowledge (see Searle, 1995). Indeed, many of the conceptual domains that are of great interest to researchers in cognitive development might have conventionalized aspects. A good example is the domain of artifact function and tool use. Although the function and utility of a given artifact or tool are constrained by its physical materials and properties, these constraints fall well short of determining the job the artifact performs (Bloom, 1996; see Chapter Six, this volume) Yet, at least for adults, among the many possible uses of an artifact, one is identified as the proper function, and others might be expressly prohibited. A fork is *supposed* to be used for eating food, and *not supposed* to be used for combing hair. Which function is the proper one is a matter of convention in that it is not determined by physical properties.

More generally, the way an object is correctly categorized can be understood to be a natural discovery or a shared but arbitrary invention (Keil, 1989; Kalish, 2002; see Chapter Four, this volume). Categorization is the process or phenomenon of grouping together similar individuals. Yet as a number of researchers have noted, the process of determining the extent to which two things are similar is not straightforward (Murphy & Medin, 1985). In the same way that physical materials and properties underconstrain artifact function, similarity underconstrains appropriate categorization judgments. Kalish (Chapter Four, this volume) suggests that what makes some categorization judgments appropriate or inappropriate is whether they are going to be shared by other members of the culture or community.

Conventional knowledge also plays a central role in game playing. Certainly when playing formal games but also during impromptu interactions, participants often act within a shared goal structure (Gilbert, 1996; Rakoczy,

Chapter Five, this volume). That goal structure is not objectively in the environment, but neither is it a subjective state of any individual participant. The joint construction of “what we are doing” becomes an intersubjective, conventional fact about how to function in the world.

Language, tool use, categorization, and game playing are all domains of great interest to researchers in cognitive development. However, they are almost always considered separately, with each assumed to have specific learning problems. The hypothesis motivating this volume is that at least at some level, these disparate domains share a common structure and present similar learning challenges. Although there are clearly differences in learning a language, learning to classify, learning about artifacts, and learning to cooperate and play games, successful negotiation of each may require children to marshal whatever cognitive resources are necessary to make sense of conventional knowledge.

Theories of Conventional Knowledge

To begin to assess this hypothesis, we need to characterize what cognitive resources are necessary to make sense of conventional knowledge. We take as our starting point three major conceptual perspectives on conventions, each of which is represented to varying degrees in the following chapters. Each theory implies a set of cognitive abilities and processes that would support and characterize the acquisition of conventional knowledge.

Conventions Emerge from Successful Coordination. Arguably the dominant account of conventions derives from David Lewis’s work (1969) on coordination problems. Coordination problems occur when an actor has multiple attractive options, but each is attractive only if another actor makes a specific choice. For example, consider students choosing seats in a class. The first day of school, each student makes a complex calculation to select the optimal seat given the others’ positions. With enough shuffling, some satisfactory distribution is achieved. Lewis notes that the students can reliably achieve a good distribution if each adopts the practice of sitting in the same spot throughout the term. In characterizing this equilibrium, or convention, it is not the case that a student simply concludes, “The best thing for me is to sit in the same spot.” Rather, the preference is conditional on expectations about the others’ actions: “The best thing for me is to sit in the same spot as long as I expect others to do the same.” On this analysis, the cognitive processes underlying coordination are relatively simple: agents need only to be able to keep track of outcomes and form expectations about others’ actions.

An important feature of Lewis’s model is that it can be applied to either social or nonsocial regularities. For instance, the processes that might shape the students’ seating behavior might also be at play for an angler deciding where to fish. That is, the angler selfishly goes to where the fish are likely to be, and the fish selfishly (though perhaps stupidly) go to where they might find some food. What is important about this parallelism is that it

highlights the fact that the behavior patterns established are not necessarily normative or prescriptive. The angler who arrives at the regular spot but finds no fish may feel surprised and upset but not cheated. Similarly, the students do not feel that each has a proper place (“my seat”), only a likely one.

Collective Intentions and Conventions. Accounting for this more normative or prescriptive characteristic of conventions is the central focus of theorists who argue that conventional systems necessarily involve collective action or shared perspective (Gilbert, 1996; Searle, 1995, 2001). Although the angler shows up expecting to find fish and the fish show up expecting to find food, there is no shared, intersubjective construal of the interaction, and thus there is no prescriptive entailment of the convention. Only when there is shared, intersubjective understanding of how to act (or think) in a given situation is there a standard by which actions (or propositions) can be evaluated for appropriateness. For example, the students in the classroom come to see the room as containing not merely tables and chairs but “your place” and “my seat.” What is established is a way something should be treated or an appropriate, correct way to act. The established conventions do not necessarily depend on past successes or provide better outcomes in the future. A student need not believe that a particular chair is the best or most effective place to sit in order to believe it is the place he or she is supposed to sit.

On some level, it might seem that participating in conventional systems according to these criteria may require relatively high-level inferences about what myriad others think about a particular situation. However, the extent to which such cognitive sophistication is required for using and acquiring conventions is very much an open question (Barr, 2005; Kalish, 2002). Theories of collective intentions were developed as conceptual analyses of conventional systems and social facts (Gilbert, 1996; Searle, 1995). What makes these analyses compelling at some level is that they highlight the fact that apparently objective phenomena actually rely on social construction. For example, people expect there is some natural basis for the value of money, such as a gold standard (Searle, 1995). The apparent objectivity of a given system may make some aspects of learning conventional systems much like learning other objective facts. Of course, there may be some situations in which more explicit awareness of the conventionality of a given domain might bring complex judgments about others’ knowledge into play. Many of the chapters in this volume address the question of how children’s developing appreciation of the conventional characteristics of a given domain might change their acquisition strategies.

Communities and Conventions. Sociocultural theories provide an interesting third perspective on the acquisition of conventions. In general, sociocultural theories of cognitive development start from the perspective that cognitive and conceptual development is best construed not as the accumulation of ideas in the head but rather as emerging abilities to participate in social practices involving behaviors, habits, and attitudes (Rogoff, 2003). A central tenet of such theories is that a complete or mature under-

standing in any given domain is a result of engaging in practices, not a prerequisite. From this sociocultural perspective, then, there is a sense in which conventionality is integrally woven into cognitive development. After all, the social practices, behaviors, habits, and attitudes are all likely best described as conventions.

Although sociocultural theories have not always clearly identified the cognitive processes and mechanisms underlying acquisition of practices, we think that the approach has two important implications within the current context. First, because of the focus on participation in cultural activities, sociocultural researchers have identified the importance of understanding the means by which experts guide novices to successful negotiation of various challenges. This approach may be fruitful for understanding how children come to act in ways that are consistent with an appreciation of conventionality (see Chapter Seven, this volume). Second, sociocultural theories provide a rich account of motivation. Acquisition of cultural practices is an important part of identity; it is by doing things “our way” that an individual comes to see himself or herself as part of a community. Moreover, successful and willful engagement in well-defined conventional practices, such as ways of speaking, serves to mark group solidarity and social distinctions. One of the goals for research in this area is understanding how the cognitive and emotional mechanisms that allow young children to be sensitive to expert guidance and motivated to be part of a group provide for the acquisition of conventional systems (see Chapter Five, this volume).

The Challenges of Acquiring Conventional Knowledge

Our aim in this volume is to consider the possibility that there might be important commonalities in learning across domains of conventional knowledge. As cognitive developmentalists, then, our first question might be something like: To what extent is the task of learning a language like the task of learning to use tools or play games? Although there is no one chapter here that will answer this particular question, the set of chapters encourages addressing a more general question regarding just what is and is not common to the structure of a set of conventional systems. Our hope is that the chapters taken as a whole may provide a starting taxonomy of conventional domains that will then allow more refined hypotheses about common and distinctive cognitive processes underlying acquisition. As a step in this direction, we conclude this introduction with a discussion of six dimensions that we believe are integral to mapping conventional domains.

Prescriptive Force. As noted in our discussion of conceptual analyses of conventions, the conventions established in a domain can be construed as either prescriptive norms governing proper behavior or expectations regarding successful behavior. You might grab a hammer rather than a wrench because you expect that a hammer will do a better job pounding a nail or because you know what hammers are supposed to do. You might

use the word *hammer* to refer to the tool because you expect that label will best convey your intent or because that is the correct or prescribed way to do so. Clearly these motives are not exclusive. However, the extent to which a given conventional practice has prescriptive qualities might differ across domains. Games typically involve prescriptive rules with sanctions for violations (such as penalties). In contrast, the norms governing tool use may be less prescriptive. Although the main purpose of screwdrivers is to drive screws, there is nothing forbidding their use as paint can openers. More important is the question of whether variations in prescriptive force affect children's strategies for acquisition. What are the psychological differences? Interestingly, one way or the other, all of the chapters in this volume consider how people react to nonstandard actions: Are they errors to be corrected (high prescriptive force) or instead imperfect but perhaps good enough strategies (low prescriptive force)? We suspect that exploring the development of intuitions about error, about being wrong as opposed to different or ineffective, will be one of the keys for psychological accounts of conventions.

Invention. At some point in history, conventional forms had to be invented. However, the degree to which invention is obvious may vary across conventional domains. For example, the conventions that govern games may seem more obviously invented than the conventions that govern the use of common tools, which may be less obvious because they have clear physical affordances that also shape their everyday use (see Chapter Six, this volume). A critical question is whether learners' beliefs about the role of invention in the origins of a given convention have an impact on the mechanisms recruited for acquisition. That is, does the child who appreciates the constructed nature of language rules learn differently from the child who sees the rules as facts of nature? Might learning in more obviously invented domains facilitate an appreciation of the less obviously invented domains?

Scope Restriction. A key feature of conventions is that they can account for both flexibility and stability in a particular representation or activity. This is partly because conventions vary with respect to their effective scope. Some conventions are held by a relatively large number of individuals (for example, all English speakers will agree on appropriate referents for the word *chair*), whereas other conventions are held by a smaller number of people or in more restricted contexts. As an example of the latter, we know of one family whose members regularly referred to soy milk as "kitty" when talking with one another (for example, "Would you like some kitty and chai?"). Recognizing the scope of a given convention involves understanding the contextual and institutional features that define a receptive population. The important question this raises for acquisition concerns how and when children come to identify conventions as restricted in scope and what the principles are that individuate different conventional contexts.

Codification. Certain conventions, such as the rules of a game or the referent of *chair*, are well established, and perhaps even codified in some way. Transmission of these conventions from one party to another is expected to be reasonably precise and faithful to the well-established standard. Other conventions, however, may be more ad hoc and transitory. For instance, a dyad might for purposes of convenience decide that bottle caps will serve as pawns in a chess game and thereby establish an ad hoc transitory convention. Nevertheless, how those bottle caps will be allowed to move in the game will remain governed by the codified rules of chess. An important point is that children are not full-fledged members of conventional communities: they may not experience, and may not have, the same degree of flexibility in conventions as adults. At the same time, recognizing instances in which conventions may be more or less codified may play some role in shaping acquisition. That is, do children encode ad hoc conventions the same way they encode the more well-established ones?

Importance of Teachers. Part of the appeal of conventionality is that in each of its domains, children's active participation and use of forms can play some role in shaping the form that is ultimately acquired. Nonetheless, it is important to note that there may also be instances in which conventional forms have to be acquired from "teachers." For instance, in some cases, social status may give someone special power in unilaterally establishing a particular convention. For example, a schoolteacher who states, "First we read, then we do math," is not merely reporting a convention the students may be unaware of; she is using her unique authority to set the standard. In Chapter Six, German, Truxaw, and Defeyter argue that the person who makes a tool has a special role in determining its usual function. Importantly, in this case, experience with the tool has little to do with assigning its conventional function.

A second case in which teachers might be particularly important is when some members of the community are relatively more experienced and knowledgeable with respect to the shared forms. Giving greater credence to the actions and statements of experts is a general principle of learning—valid in both conventional and natural domains (for example, the division of linguistic labor; Putnam, 1975). Yet there may be crucial differences between what constitutes expertise in conventional versus nonconventional domains. In nonconventional domains, experience with a particular phenomenon would seem to be sufficient to confer knowledge. But what kinds of experiences make an informant knowledgeable about what everyone else in the community does? What can actors do to signal that they are experts in word meanings or artifact functions? It remains an open question how children identify the best informants about conventional knowledge.

Flexibility. Theoretically if action in a given domain were specified solely by conventions, we might expect highly stereotyped behavior in that domain. However, no domain of behavior is completely specified by conventions (and perhaps none is completely free). Instead, there is substantial

room for innovation and idiosyncratic performance in any of the domains that we have discussed up to now. For instance, rules of games specify prescribed actions, but the manner and style of how these actions are carried out can vary across players. In golf, the only rule that governs putting is that the player cannot straddle the imaginary line running from the golf ball to the hole. All golfers abide by this rule, but it obviously underconstrains all of the other aspects of the mechanics of putting. This raises a very important problem for acquisition: How do children figure out which actions within a general domain are governed by convention and which are not (see Kalish, Chapter Four, this volume)? What cues are available in the environment that might help distinguish conventionalized aspects from idiosyncratic ones?

Relations Among Dimensions. Although we see these dimensions of conventionality as conceptually distinct, in reality there may be important relations among them. For instance, conventions that have a wide scope (such as the possible referents for the word *chair*) might also be more codified, more likely to carry prescriptive force, and require an appropriate teacher. In contrast, conventions that have a more restricted scope might be less codified, less likely to carry prescriptive force, and be within the purview of anyone to establish. We elected to separate these dimensions because it seems possible to come up with many examples of instances in which these general relations did not seem to hold (this journal includes entries in the References section in the required word count for manuscripts, a convention that has narrow scope but is codified and carries prescriptive force). Yet if there is some way in which features of conventions reliably coalesce, then it raises some interesting questions regarding development. In particular, early in development, do children have a more rudimentary understanding of conventions that conflates dimensions (for example, do they mistakenly believe that all stable conventions will have a wide scope)? If so, how might these dimensions ultimately get separated?

Conclusion

Many of the domains that have interested researchers in children's cognitive and conceptual development have characteristics that are conventionalized: the judged appropriateness of the forms is arbitrated by social consensus. Our goal in this introduction has been to suggest that these domains, although disparate with respect to their content, might have a common structure and thus might pose similar problems to young children in the course of acquisition. To delineate what these problems might be, and provide a framework for positioning each of the following chapters, we reviewed theoretical and conceptual analyses of conventionality to develop some hypotheses about the cognitive processes that might be implicated in either acquiring or participating in conventionalized domains. We also highlighted six dimensions on which conventions can vary among themselves and the special problems that each might pose for acquisition.

References

- Bloom, P. (1996). Intention, history, and artifact concepts. *Cognition*, 60, 1–29.
- de Saussure, F. (1983). *Course in general linguistics*. London: Duckworth.
- Gilbert, M. (1996). *Living together: Rationality, sociality, and obligation*. Lanham, MD: Rowman and Littlefield.
- Kalish, C. W. (2002). Essentialist to some degree: Beliefs about the structure of natural kind categories. *Memory and Cognition*, 30, 340–352.
- Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.
- Koenig, M. A., & Harris, P. L. (2005). The role of social cognition in early trust. *Trends in Cognitive Sciences*, 9, 457–459.
- Lewis, D. K. (1969). *Convention: A philosophical study*. Cambridge, MA: Harvard University Press.
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289–316.
- Putnam, H. (1975). *Mind, language and reality*. Cambridge: Cambridge University Press.
- Rogoff, B. (2003). *The cultural nature of human development*. New York: Oxford University Press.
- Searle, J. R. (1995). *The construction of social reality*. New York: Free Press.
- Searle, J. R. (2001). *Rationality in action*. Cambridge, MA: MIT Press.

CHARLES W. KALISH is professor of educational psychology at the University of Wisconsin-Madison. His work addresses the development of inductive inference. He is especially interested in the role of norms in social explanation and prediction.

MARK A. SABBAGH is associate professor of psychology at Queen's University in Kingston, Ontario, Canada. He is interested in the social, cognitive, and neural bases of children's language and social-cognitive development.