Published in final edited form as:

Annu Rev Psychol. 2006 ; 57: 227–254. doi:10.1146/annurev.psych.57.102904.190100.

Explanation and Understanding

Frank C. Keil

Department of Psychology, Yale University, New Haven, Connecticut 06520-8205

Frank C. Keil: frank.keil@yale.edu

Abstract

The study of explanation, while related to intuitive theories, concepts, and mental models, offers important new perspectives on high-level thought. Explanations sort themselves into several distinct types corresponding to patterns of causation, content domains, and explanatory stances, all of which have cognitive consequences. Although explanations are necessarily incomplete—often dramatically so in laypeople—those gaps are difficult to discern. Despite such gaps and the failure to recognize them fully, people do have skeletal explanatory senses, often implicit, of the causal structure of the world. They further leverage those skeletal understandings by knowing how to access additional explanatory knowledge in other minds and by being particularly adept at using situational support to build explanations on the fly in real time. Across development and cultures, there are differences in preferred explanatory schemes, but rarely are any kinds of schemes completely unavailable to a group.

Keywords

concepts; causality; cognition; cognitive development; illusions of knowing; domain specificity; stances

INTRODUCTION

Humans are driven to acquire and provide explanations. Within months of uttering their first words, children ask “why.” Preverbal infants explore phenomena that puzzle them in an attempt to uncover an explanation of why an effect occurred. As adults, we must frequently choose between explanations of why politicians lost, why the economy is failing, or why a war is not winnable. Moreover, explanations are not merely the work of experts. Our friends explain why they have failed to honor a commitment or why a loved one is behaving oddly. Our enemies may offer unflattering explanations of our successes. Explanations are therefore ubiquitous and diverse in nature. This review considers the varieties of explanations, their components and structure, and their uses.

WHAT EXPLANATIONS AND EXPLANATORY UNDERSTANDINGS ARE

Explanations were once thought to be “deductive-nomological” in nature (Hempel & Oppenheim 1948). In this view, explanations are like proofs in logic. A set of basic laws (the nomological part) are stated as axioms and then the deductive consequences of those laws are explored like a logical proof. The complex of laws and the deductive sequence constitute the explanation. Thus, an explanation of the periodicity of pendulums might assume certain laws of classical mechanics and the deductive consequences of those laws when considered in conjunction with additional initial statements about pendulums. This model of explanation has not fared well in the philosophy of science (Salmon 1989). Scientists do not normally proceed in that manner as individuals, and even disciplines as a whole rarely follow so neatly the progression of a proof. Moreover, as one considers explanations in

sciences other than physics, even superficial similarities to deductive chains start to disappear (Salmon 1989). For laypeople as compared with scientists, the deductive-nomological model seems even less plausible.

Everyday explanations depart radically from the image of people methodically considering a set of axioms and running through a deductive chain. For example, people frequently prefer one explanation to another without explicitly being able to say why. They often seem to draw on implicit explanatory understandings that are not easy to put in explicit terms (Kozhevnikov & Hegarty 2001). Even without strictly adhering to the deductive-nomological model of explanations, some explanations in domains such as biology hardly invoke laws at all (Bechtel & Abrahamsen 2005). In many cases, we think of explanations as providing some sense of mechanism (Bechtel & Abrahamsen 2005, Chater & Oaksford 2005, Glennan 2002). Indeed, when explanations of psychological phenomena are stated in law-like ways without mechanism (e.g., if X, then Y), they often instead are called “effects,” which suggests that they are not really explanations (Cummins 2000).

People often recruit causal mechanisms on the basis of one-trial learning (Chater & Oaksford 2005) and not on the basis of a long-term gradual accumulation of statistical data in a manner specified by Hume and since explored in psychology (e.g., Cheng 1997, Dickinson 2001). Although statistical reasoning can certainly guide the formation of causal explanations, in many everyday cases, one or two trial exposures seem to activate a pre-existing schema of a mechanism.

Although the ability to provide verbally explicit explanations may be too strict a criterion to count as having explanatory understanding, not all forms of implicit knowledge that enable prediction will count. One purpose of this review is to convey a better sense of implicit explanatory understanding. More broadly, not only the processes of creating and discovering explanations but also the processes of providing and receiving them are considered, as all these activities are intuitively part of what is meant by “explanatory understanding.”

EXPLANATIONS, THEORIES, AND MENTAL MODELS

Explanations are related to theories; but a focus on explanations brings different issues and bodies of research to the foreground (e.g., Brem & Rips 2000). One difference concerns the transactional nature of explanations. Explanations are often between individuals and reflect an attempt to communicate an understanding. Even within one mind, explanations occur in a manner different from intuitive theories. One can attempt to explain an event to oneself, as sometimes is revealed when people are heard saying to themselves out loud how something works. Even young children playing alone can be observed to explain things to themselves as a verbal strategy to help solve a task (Berk 1994). Explanations create trajectories. Recipients of explanations, if the explanations are at all successful, are expanding their understanding in real time. The transactive process also frequently implies conceptual change. Theories can undergo conceptual change (see, e.g., Carey 1985), but they do not need to and can be highly stable over extended periods, especially when very successful.

Explanations may highlight incompleteness. An explainer will often encounter gaps in understanding that may remain largely invisible when in the form of intuitive theories. Thus, the process of trying to explain explicitly a system to another, or even to oneself, often brings the incompleteness of one’s understanding into much harsher relief (Keil et al. 2004).

Explanations contrast with mental models as well, which can range from formal representations of logical patterns (Johnson-Laird 1983) to image-like representations of the workings of a system (Gentner & Stevens 1983). Mental models are readouts of relations from a mental array and often are understood in spatial terms. Explanations normally are not

seen as mental blueprints or plans that are then read off. They include the interpretations of such plans or blueprints.

Explanations therefore contrast with both intuitive theories and models because of their transactional component, their role in expanding knowledge, and their interpretative role. They are also different from simple, procedural knowledge. Knowing how to operate an automated teller machine or make an international phone call might not entail having any understanding of how either system works. Even a seemingly simple act such as exchanging money at an airport may carry with it no explanatory understanding of how relative currency values are determined. The study of explanations offers insights not obvious from other points of view. Explanations view people less as autodidacts and more as social, interacting agents (Harris 2002).

ARE THERE DIFFERENT KINDS OF EXPLANATIONS?

Explanations occur in different ways (Keil & Wilson 2000a). One can explain why Aunt Edna insulted Uncle Billy at the family holiday dinner, why giraffes have long necks, why salt melts ice, or why seat belts prevent traffic fatalities. There are explanations of individual histories, and of why one thing is functionally “meant” for another. Do all these kinds of explanations work similarly or do they have different properties and perhaps even different developmental trajectories? We can contrast explanations in terms of the causal patterns they employ, the explanatory stances they invoke, the domains of phenomena being explained, and whether they are value or emotion laden.

Causal Patterns

Explanations often refer to causal relations, but in at least four distinct ways: common cause, common effect, linear causal chains, and causal homeostasis. In common-cause explanations, a single cause is seen as having a branching set of consequences (Sober 1984). For example, a virus infects a person and then has a cascade of effects that creates a downward-branching hierarchy differentiating over time. Common-cause explanations are frequently found in diagnoses of problems, such as medical disease, equipment malfunction (e.g., auto repair), or software bugs.

Common-effect explanations involve cases where causes converge to create an event. These sorts of explanations are common in history, where a major event might be attributed to the confluence of several factors. Thus, the defeat of the Spanish Armada in 1588 might be attributed to the converging causes of the Netherlands Revolt, severe weather at sea, personal problems of Phillip II of Spain, and the brilliant marine tactics of Lord Charles Howard.

Explanations as simple linear chains are a special degenerate case of common cause and common effect explanations; namely, there is one unique serial chain from a single initial cause through a series of steps to a single effect. One fanciful case is the child’s verse of how a missing horseshoe nail led to the loss of a kingdom. Simple linear explanations may be quite rare in real life, however. Even if things start with a single chain of effects and causes, at some point those effects start to have multiple effects of their own and the structure starts to branch. Moreover, other causes enter in and influence the reasoning process downstream, thus violating the notion of a single cause.

Causal homeostatic explanations are fundamental to explanatory discourse about natural kinds (Boyd 1999). These explanations seek to account for why sets of things seem to endure as stable sets of properties. A causal homeostatic explanation does not seek to explain how a cause progresses over time to create some effect or effects, but rather how an

interlocking set of causes and effects results in a set of properties enduring together as a stable set over time that then exists as a natural kind. Thus, it seeks to explain why feathers, hollow bones, nest building, flight, and a high metabolic rate might all reinforce the presence of each other in birds (Boyd 1999, Keil 1989).

Several psychological issues arise in the context of these four broad patterns: First, people do seem to find some patterns more natural to think about than other patterns (Ahn et al. 2000, Rehder & Hastie 2004). Second, these differences in naturalness may in turn lead to cognitive biases and distortions in explanations and explanatory understandings. Finally, people may view certain patterns as domain biased, such as seeing causal homeostasis as especially apt for living kinds (Ahn 1998, Keil 1989).

Explanatory Stances

A different way of contrasting types of explanations involves “stances” (Dennett 1987) or “modes of construal” (Keil 1995). People may adopt a stance or mode of construal that frames an explanation. These stances are not in themselves theories; they are far too vague and nonpredictive. However, they do posit certain kinds of relations and properties, and even arguments, as central. Dennett (1987), for example, focused on the mechanical stance, the design stance, and the intentional stance. The mechanical stance considers simple physical objects and their interactions. The design stance considers entities as having purposes and functions that occur above and beyond mechanical interactions (Lombrozo 2004). Some have argued that functional, teleological interpretations come all too readily to mind in evolutionary “storytelling” that infers functional explanations for cases that don’t need them (Gould & Lewontin 1979). Similarly, some see young children as “promiscuously teleological” (Kelemen 1999), whereas still others see them as more selective in their use of functional explanations (Greif et al. 2005, Keil 1992). There are also debates as to whether a design stance or teleological mode of construal requires a notion of an intentional designer (e.g., Bloom 1996, Keil 1995, Kelemen & Carey 2005). Related issues concern the proper scope of teleological arguments in the sciences (Allen et al. 1998). Finally, the intentional stance considers entities as having beliefs, desires, and other mental contents that govern their behavior. Entities are assumed to be governed by a belief-desire calculus and to have mental representations that have causal consequences. Adopting an intentional stance, however, may not require having a full-fledged folk psychology (Bloom & German 2000, Gergely et al. 1995).

The same event often can be characterized by each of the distinct stances and thereby yield quite different insights and explanations. For example, a diver going into a tuck position might be explained in terms of the physics of rotating objects (a mechanical stance), in terms of the purpose of pulling in the limbs close to the body (design stance), and in terms of the beliefs the diver has about her actions and the motivations that drive them (the intentional stance). Each of these ways of framing the action will afford different insights and, potentially, different distortions. Similarly, atypical stances sometimes are adopted to shed additional insight into a system or for pedagogical reasons. One can discuss students in a crowded high school hallway between classes as billiard balls careening about, and one can describe magnets as liking certain metals or materials as having memories. There, too, distorted understandings can result (Gentner & Gentner 1983).

Is there a developmental progression of a mechanical stance, a design stance, and an intentional stance? Or does a design stance depend on an earlier emerging intentional stance? Or do they all appear in parallel early in development as foundational core ways of explaining? These questions are some of the most actively investigated in the field today (German & Johnson 2002, Keil et al. 2005, Kelemen 1999, Lombrozo & Carey 2005).

Explanatory Domains

Explanations also can be contrasted in terms of domains roughly corresponding to academic disciplines. Are there distinct explanatory types in intuitive biology, physics, chemistry, psychology, economics, and sociology? Are the ways we explain phenomena pertaining to reproduction of organisms profoundly different from how we explain trajectories of projectiles and the stratifications of social groups? In the philosophy of science, there has been a sea change of opinion in which older reductionist/eliminativist views have been replaced by pluralistic views (Salmon 1989). Thus, biologists, especially evolutionary biologists, might well invoke functional arguments (even as they also caution against reverse causality) far more than physicists or chemists. Psychologists might not consider monotonic relations between inputs and outputs to be as pervasive as are those in the physical sciences, and may expect probabilistic relations usually to obtain between causes and effects (Lehman et al. 1988). A more systematic taxonomy of explanatory domains awaits a more complete formal way of describing the structure of explanations.

Social- and Emotion-Laden Explanations

Explanation may have a different character when it is socially or morally laden. In social attributions, for example, motivational factors seem to color how we construct and accept explanations. Moreover, the presence of explanations can influence how we emotionally experience events (see, e.g., Wilson et al. 2005). The literature is far too vast to consider here in detail (see, e.g., Macrae & Bodenhausen 2000, Malle 2004), but the following point merits mention. Even if motivational factors heavily influence social explanations, such influences in themselves may not change the explanations' structure. The same kinds of causal patterns may still be favored, the same sorts of probabilistic distributions preferred, and the same domains of regularities considered as relevant. Emotions and motivation may simply shift a threshold for acceptance of evidence, or strength of causation. Moral explanations are another special case. A difference may exist between moral thoughts that have the form of explanations and those that defy explanations in phenomena known as "moral dumbfounding" (Greene & Haidt 2002, Haidt 2001). People sometimes explain their moral judgments by reference to a set of principles, rules, or laws; but other times, often with cases of such taboos as incest, sacrilege, and torture, people have "gut" responses more akin to disgust or perception without apparent explanation. The kind of person sought out to provide a moral explanation might also be conceived differently from other kinds of experts (Danovitch & Keil 2005).

WHAT EXPLANATIONS ARE FOR

There are several reasons why we engage in explanations and seek them out. The most commonly offered reason is to be able to predict similar events in the future (Heider 1958). We explain, for example, why a stretch of road is particularly slippery in cold, wet weather so that, on future occasions of similar weather, we might anticipate those conditions more effectively.

Explanations are also used in diagnosis. One might ask why a system failed and then repair a part to bring it back to its normal function (Graesser & Olde 2003). There is little sense of making a prediction in such cases. One is not concerned with predicting but rather with restoring operation of a system. In "broken" mechanical, biological, and social systems, we often seek explanations to "fix" them. Other times, we develop explanations of the causal efficacy or inefficacy of a new action or tool. We might try several different ways of getting a car out of the mud, and when we finally succeed, we try to come up with an explanation for the method that succeeded so that we can reproduce that method in the future. Again,

there is less the feeling of prediction here and more the sense of identifying a causally critical component that one wants to reproduce.

We also often seek explanations for one-time events in which predictions for the future seem completely implausible. The press might search endlessly for explanations of why the car carrying Princess Diana ended up in a fatal crash, but nowhere in that search for explanations is there any motive to predict future crashes. The constellation of factors leading up to that crash was clearly so unusual and unique that looking for an explanation in service of a prediction makes little sense. Why then, do people search for explanations in such cases? The notion of learning from one's mistakes does carry with it the idea of predictive value, but often the motive is more one of affixing blame. One may seek an explanation simply to determine the guilty party that needs to pay for damages or be punished. In the study of history, the search for explanations of events, such as the fall of a dynasty or the start of a war, is often accompanied with severe warning about not using those explanations to predict future events. Thus, there is a great deal of controversy over whether it is appropriate to use counterfactuals in historical analysis (Ferguson 1999).

Other times we engage in explanations to justify or rationalize an action. We explain that we punished a child for his own good, or that we didn't bother voting in an election because our vote would not have mattered anyway, or that we were sickeningly sweet to an enemy because of the strategic value of being nice on that occasion. Our explanations are attempts to represent our actions to others as sensible, well intentioned, pragmatic, or appropriate. Such acts of persuasion can often be self-directed. Thus, it is not uncommon to explain to oneself why one has engaged in a certain act as a way of making sense of one's own actions. I might be surprised at my burst of anger at a meeting and later explain it in terms of situational as opposed to dispositional factors. I might do much worse on a test than expected and explain it in terms of lack of studying, and, in doing so, shift the reason from one of ability to that of effort.

Finally, explanations can be in the service of aesthetic pleasure. One can explain a work of art, a mystery of cosmology, or the intricacies of a poem with the sole goal of increasing appreciation in another, providing that person with a better polished lens through which to view the explanandum. Explanations in their own right can be immensely rewarding things and may be sought out as such, even by the youngest of children (Gopnik 1998). When a young child asks "why," she often simply wants to appreciate more the nature of something that she has observed, with no further agenda.

In short, explanations are used in many different ways beyond mere prediction. It is not yet clear how their content might vary as a function of their use. For example, are self-directed explanations largely ones about ability and individual differences in a social comparison context? Are justifications of actions usually explanations offered in a moral context or are they simply part of a much larger set of ways of describing why we did what we did, just as an accountant might explain an audit procedure? Are explanations of systems fundamentally different from those of individual histories? For example, contrast an explanation of why a particular couple got married with why couples in general get married.

Perhaps explanations have only limited use in day-to-day life because it would take too long to construct an explanation every time one was needed (Fodor 1998). But explanations don't have to work in real time. Instead, they may frequently serve to help people know how to weight information or how to allocate attention when approaching a situation. It used to be doctrine that explanation-based processing of information was slow and later in tasks, whereas early processing was more associative and fast; more recent findings illustrate that

in many cases, presumably where some sort of precompiling takes place, explanation-based effects can occur in the earliest steps of processing (Keil et al. 1998; Luhmann et al. 2002).

INTRINSIC LIMITS OF EXPLANATIONS

Explanations are ubiquitous, come in a variety of forms and formats, and are used for a variety of purposes. Yet, one of the most striking features about most explanations is their limitations. For most natural phenomena and many artificial ones, the full set of relations to be explained is enormous, often indefinitely large and far beyond the grasp of any one individual (Wilson & Keil 1998).

The problem of overwhelming complexity may seem to be unrealistic given how often we seem to be successful with our explanations in everyday life, but that is precisely the point. Somehow, people manage to get by with highly incomplete or partial explanations of how the world around them works. Are there systematic ways in which we extract a gist of the causal structure of the world? The problem is analogous to that confronted by the realistic artist who must attempt to convey, in a two-dimensional plane, the vastly larger amount of information that is embodied in a three-dimensional array (Cavanaugh 1999). Artists use a number of conventions and exploit interpretive “tricks” of the perceptual system. It seems likely that our explanations make similar moves, although we have yet to understand the nature of such compressions of information. To use a different analogy, in categorization, for at least some classes of objects in the artifacts realm, there is a “basic” or optimal level at which to encode information about categories (Murphy 2002, Rosch et al. 1976). Are there comparable optimal levels of causal description in various domains of the physical world? (For a related discussion of the “compression problem” and “global insight,” see Fauconnier & Turner 2002.) As discussed below, the problem of the intrinsic limits of explanation is exacerbated by a tendency to overestimate the depth of our own understandings. It is bad enough that explanations have indefinite depth; it is much worse if we realize too late when we are out of our explanatory depth.

THE CENTRAL ROLE OF CAUSALITY

Not all explanations are about causal relations. One can explain how a mathematical result is achieved, why a design is symmetrical in a subtle manner, how a jazz improvisation resolves itself, or why China is bordered by 14 different countries. Yet, the vast majority of our everyday explanations invoke notions of cause and effect. Moreover, when an explanation contains both causal and noncausal elements, the causal ones tend to occupy center stage and dominate patterns of judgment (Murphy & Medin 1985). Similarly, causal understanding shifts perceptions of the normalcy of properties (Ahn et al. 2003). Causality seems to have a privileged role in most explanations.

Philosophers of science have classically turned to the notion of cause to illustrate why some things explain others and not the opposite. Thus, although the presence of a shadow is lawfully related to a flagpole and the time of day, the flagpole and the sun explain the shadow (Bromberger 1966). The shadow cannot explain the flagpole or the sun. Lawful relations and tight correlations do not have the sense of an explanation when the causal relations are absent or inappropriate (Cartwright 2004).

Counterfactual thought also illustrates the central importance of cause in explanations (Lewis 1973). We sense the meaning of a causal relation by stating that, *ceteris paribus*, a particular event B would not have occurred if event A had not occurred first. This “would have” relationship is meant to highlight causation as opposed to mere correlation. Although the ways in which counterfactuals illuminate the nature of causation are philosophically complex (Collins et al. 2001), at a psychological level we often turn to counterfactuals to

explore causal relations or make them more salient (Roese 1997). Counterfactual thought is often triggered by other features, such as the atypicality of an event (Kahneman & Miller 1986), which may then in turn trigger more causal thought. It can also serve to set up a simulation heuristic, in which an event is considered with minor perturbations as a way of better understanding its causal structure (Kahneman & Tversky 1982). The ways in which counterfactual thinking highlight, or perhaps sometimes distort, causal relations are complex and remain an active area of exploration (Spellman & Mandel 1999). One mechanism may involve shifting perceptions of baseline probabilities of the actual event as opposed to alternative events, thereby making a potential cause seem either more or less inevitable (Spellman & Mandel 1999).

Although older views often saw the world as governed by one unified set of laws, i.e., physics, those views do not seem to capture actual science. One alternative argues for a “dappled world” in which there are local causal regularities and fragments, and not a global unified explanation of everything (Cartwright 1999). Even at the same level of reduction, such as the interactions of bounded objects, there might be different domains of regularities. Thus, the movements of a falling leaf may have little in common with those of a falling apple, as the leaf brings in relations of fluid dynamics and the apple those of more classical Newtonian mechanics. If the world is a heterogeneous collection of different causal patterns, it is reasonable to ask if those patterns result in different explanatory styles. Indeed, evolutionary psychologists have argued that there might be a selective advantage to develop different explanatory schemes to resonate with these different kinds of causal regularities (Cosmides & Tooby 1987). Just as organisms have developed different sense organs to handle the heterogeneous forms of information carried in sound, light, and chemicals, they may also have developed mental organs to grasp different kinds of causal regularities that have stably existed in the world for millions of years.

Probabilistic and deterministic causation are also importantly different. Evolutionary theory, economics, and much of psychology are deeply dependent on notions of probabilistic causation. Yet, in the simple mechanics of macroscopic bounded objects, we often think of discrete causes of events. The philosophy of science has long recognized that both types of causation have explanatory value in the different sciences (Humphreys 1989, Salmon 1980). Higher levels of education in some areas of the social sciences may foster a sensitivity to certain kinds of probabilistic causal reasoning (see, e.g., Lehman et al. 1988), which suggests we tend to weigh an interpretative causal schema more or less strongly as a function of the perceived domain.

Adults and children therefore tend to associate particular kinds of causal schemata with specific domains. Beyond such domain differences, however, there are also consistent biases towards specific causal patterns. Features that are earlier in causal chains are often deemed more important in category learning (Ahn et al. 2000), and features that are more causally interdependent on others are similarly more critical in category learning and induction (Sloman et al. 1998). Thus, people prefer explanations in which the most emphasized features and properties either are early in a causal chain or are the most causally interdependent on others.

In some studies, common-cause, common-effect, and simple cause-chain schemas facilitated induction to the same extent for biological kinds, artifacts, and nonliving natural kinds (Rehder & Hastie 2004). Similarly, a preference for the first element in a causal chain, the causal status effect, has been found to be influential in both artifact and biological domains (Ahn 1998). Yet, even children as young as three, when spontaneously asking questions about novel artifacts and novel animals, seem to seek out different kinds of causal relations, asking about the functions of artifacts as a whole but not of animals. With animals, children

focus more on the locations in which they are found (Greif et al. 2005). Thus, although a wide array of causal patterns can be applied to each kind, there are different preference orderings across domains.

A related influence of causal structure can be found in a “causal diversity effect” (Kim & Keil 2003). Single causes usually have downwardly branching sets of effects and causes. Thus, a bacterial infection might initially cause a white blood cell elevation and fever. Those effects in turn might cause anemia and headache, which in turn can cause other effects. The further apart two final effects are in the branching tree of causes the more they are judged good evidence for the initial cause. Explanations are more appealing when they use diverse forms of evidence for initial causes. Indeed, in 1847, William Whewell, the first philosopher of science, argued for the notion of “consilience,” in which the most appealing scientific explanations were ones for which evidence of effects came from maximally diverse sources (Whewell 1847).

Causality has figured prominently in explanatory thinking in recent years in the context of Bayesian learning and related paradigms. Causal patterns may be discerned from merely correlational ones by construction of “causal Bayes nets” (Glymour 2001), which track patterns of nested contingencies in a manner that strongly rules out noncausal relations. For example, if one screens off a particular variable *X* and the effect *Y* no longer occurs, evidence accumulates for a causal relationship between *X* and *Y*. Not only adults but also young children use such strategies to infer causal relations (Gopnik & Glymour 2004, Sobel et al. 2004). More recently, Bayesian techniques have explored how particular causal and relational patterns might be learned and linked to domains in a manner that guides explanatory interpretation. For example, in one study, a learning system that was fed data about biological systems constructed a taxonomic interpretative structure and then used that structure to guide future learning (Kemp et al. 2004). The system was quickly able to detect the taxonomic relations and then use them to constrain further learning. In a different domain, such as political party preferences, a different relational structure would be learned. The most relevant case involves a learning system that learned to distinguish relational patterns quite similar to those of causal chains, causal hierarchies, and causal homeostasis (Kemp et al. 2004). Not all causal knowledge may arise from Bayesian learning; indeed, prior causal explanatory schemes may distort Bayesian learning in systematic ways (Krynski & Tenenbaum 2003). In short, there is a recent convergence between behavioral studies, which show sensitivity at all ages to a wide range of causal and relational structures, and newer techniques of modeling learning, which show the detection of such structures in specific domains and then the use of them to further constrain learning in those domains. Those simulation accounts are also revealing how prior beliefs about certain causal relational patterns can constrain further Bayesian learning in a domain and, therefore, generation of explanations.

Analogies often are critical to successful causal explanations, especially those of systems as opposed to singular event sequences. Analogies have been used in scientific explanations ever since the first attempts at understanding nature have been recorded (Holyoak & Thagard 1995). Analogies can be used both in the process of scientific discovery and in the process of giving an explanation, but they may work in different ways. In day-to-day discovery in the actual laboratory, analogies may be plentiful but may be relatively mundane as they make linkages between closely related areas. In a molecular biology laboratory, scientists might use analogies from one gene of the HIV virus to another gene of the same virus (Dunbar 1995, Dunbar & Blanchette 2001). In contrast, when scientists explain systems to outsiders or teach students, the analogies may draw on much more distant domains, such as the familiar solar-system-as-atom case (Gentner 1983) or electricity as a flowing liquid (Home 1981). Thus, when people are seeking an explanatory understanding

in an area, they might use close analogies; but when they are actually trying to formulate explanations to others, analogies between more distant domains may become more common. Different facets of explanation therefore seem to draw on different forms of analogies enabling access to different levels of causal analysis. Because analogies draw on intrinsically relational patterns and because relational patterns may be at the core of many explanations, analogies also help both adults and children to go beyond surface patterns to see the causal patterns that lie beneath (Gentner et al. 1995).

RECOGNIZING BAD EXPLANATIONS AND CHOOSING THE BEST EXPLANATION

Three important dimensions guide our evaluations of explanations: circularity, relevance, and coherence. These criteria can be surprisingly elusive in real-world contexts. For example, although some cases of circularity are blatantly empty, e.g., “This diet pill works because it helps people lose weight,” more complex and lengthy circularities can be much harder to detect (Rips 2002). People tend to confuse pragmatic factors, such as repetition of facts, with true circularities in evaluating the reasonableness of explanations. They recognize circularity in its simplest and most straightforward forms, but can be easily sidetracked in more complex bodies of discourse. Thus, our ability to detect faulty explanations, even ones that are logically circular, is not fail-safe.

Relevance would also seem to be straightforward. Yet, levels of abstraction, analogies, and surprising connections to other domains can complicate the assessment of relevance. In one view, an input is relevant to the extent that, at any time, the cognitive effects of processing that input are positive (Wilson & Sperber 2004). Explanations are good if they are high on this dimension of relevance. Relevance can also be understood as a principle of speech act theory: Speakers should be informative (Grice 1975). Explanations therefore suffer if presented at the wrong level of detail. Thus, if asked why John got on the train from New Haven to New York, a good explanation might be that he had tickets for a Broadway show. An accurate but poor explanation at too low a level might say that he got on the train because he moved his right foot from the platform to the train and then followed with his left foot. An accurate but poor explanation at too high a level might say that he got on the train because he believed that the train would take him to New York from New Haven. Finding the appropriate level of analysis has long been a problem in analyses of event structures (e.g., Schank & Abelson 1977), so it should be no surprise that it is an important factor in judging the appropriateness of an explanation.

Explanations can also go awry for another speech act–related reason concerning egocentrism. Although originally discovered by Piaget (1926) in children and later expanded on a variety of communication tasks (e.g., Kraus & Glucksberg 1969), egocentrism is a common problem for people of all ages (Keysar et al. 1998, Nickerson 2001). An explanation of how something works will fail either if it provides too much detail or if it presupposes too much and skips over essential details. One version of this deficit sees adults and children alike as laboring under a “curse of knowledge,” in which having knowledge biases one to think that others have the same knowledge (Birch & Bloom 2003). More broadly, we may try to understand what others know by looking at our own knowledge and using that as a model to project what others know (Nickerson 2001). We adjust our explanation to take into account the other, but by using our own knowledge as a point of departure, we may egocentrically distort it. This process influences explanations by making us miscalculate the informational common ground between explainer and explainee (Clark 1996). Although speakers attempt to negotiate a common ground (Clark 1996), the egocentrism bias intrudes.

Explanations can also be seen as bad if they fail to cohere, or “hang together.” The different elements of an explanation must work in concert to achieve an internally consistent package. The doctrine of coherentism is seen as an important alternative to foundationalist views of scientific explanations that attempt to reduce all phenomena to the bedrock of physics (Amini 2003). Coherentism argues that a set of statements at a particular level of explanation, such as the psychological, can cohere as a tightly organized and interrelated unit that offers insights and explanations in its own right, without having to depend on lower levels of explanation (Amini 2003).

Coherence has been defined in terms of constraint satisfaction (Thagard 2000, Thagard & Verbeurgt 1998). A set of elements is coherent to the extent that each element in a set positively constrains other ones, often causally. Elements can also be negatively constraining, that is, they contradict or causally block other elements. One can quantitatively derive coherence values from sets of elements and implement those analyses in computational systems such as neural networks (see, e.g., Thagard 2000, Thagard & Verbeurgt 1998). These formal analyses are difficult to apply to everyday explanations, but the degree to which elements hang together seems to be an important intuitive component in evaluating the quality of explanations. Coherence is also related to a notion of systematicity, the extent to which elements form a tightly interconnected, mutually supporting relational structure. For example, Gentner & Toupin (1986) suggest that a focus on relational predicates can reveal cases in which elements are tightly linked as opposed to merely being properties associated with a category.

There are, however, both conceptual and empirical problems with the idea that people are highly sensitive to coherence. Conceptually, there is the problem of “holism” (Fodor 1998): Virtually all elements eventually are causally connected to all others. Thus, if one were trying to explain how a multispeed bicycle works, much of the coherence would seem to rest in how the mechanical elements of a bicycle interact with each other. Yet, those metal mechanical elements are also constrained by elements having to do with human anatomy, physiology, and goals. Similarly, they are constrained by riding surfaces, economics of construction, and so on in an ever-widening set of relations. This holism problem may be surmountable by heuristics that prune out links to other elements that are below a certain level of density; but such heuristics are difficult to implement in practice and are highly sensitive to context.

Empirically, a full commitment to coherence seems to evaporate when people are probed concerning their actual beliefs about how the world works. They seem to violate coherence in two ways. First, they often seem to know only fragments of a full system. Those fragments may consist of so few elements that they hardly qualify as cohering. In the extreme, people may have little more than a set of phenomenological primitives, or p-prims, that are very small pieces of understanding with few relations of consistency or coherence to other elements (Di Sessa 1993, Di Sessa et al. 2004). Second, people often seem to have sets of beliefs that directly contradict each other (Chin & Brewer 1993, Winer et al. 2002). The conflicts may be ignored until they are made explicit; and even then, any changes in beliefs may be resisted. Coherence may be most powerful when all relevant elements are explicitly considered at the same time. However, given the limits of working memory, the number of such elements actually considered is likely to be quite small. In short, coherence and systematicity influence our sense of the quality of explanations, but they also are limited in ways that are not yet fully understood.

We frequently have to choose between competing explanations. Sometimes both explanations might be possible as converging sources of influence, but other times they may directly conflict such that only one can be correct. Explanation choice between reasonable

explanations may be quite distinct from rejecting truly inappropriate explanations. One bias governing explanation choice may involve the order in which we receive explanations. Thus, if we are presented with parts of one causal explanation of how a system works, we may then discount subsequently presented fragments from an alternative explanation, often on the basis of a single trial of learning. Causal discounting may not require an explicit evaluation and comparison of two explanations, and it often may occur in a more automatic manner outside of awareness (Oppenheimer 2005).

Choices between explanations have been described as operating through a process known as inference to the best explanation (IBE) (Harman 1965). In IBE, an individual generates alternative hypotheses and then considers, as a result of those hypotheses, tradeoffs between external verification of the hypotheses and some metric of internal parsimony. Thus, if an explanation seems to predict and account for a huge array of natural phenomena but is a large set of seemingly unrelated ad hoc elements, it might be preferred less than one with a bit less predictive power but with a much simpler internal structure. IBE approaches, however, require clarifications of notions of parsimony, inference, and empirical success. Simplicity and complexity are notoriously slippery notions (e.g., Chater & Vitanyi 2003, Sober 1975), and as noted earlier, logical inference as a deductive chain rarely characterizes real life cognition either in scientists or in laypeople.

We may also choose between explanations by asking which of two mechanisms seems more plausible, or how well a mechanism conforms to early emerging causal schemata. Plausibility can also be driven by familiarity with similar classes of mechanisms. Explanation choice may also be guided by the extent to which a visual animation of mechanism easily comes to mind (Barsalou 1999, Hegarty 1992). Appeals to mechanism, however, have been criticized for relying on preexisting causal schemata that ultimately may have to be stipulated as innate (Glymour 2001). In summary, we use a wide array of heuristics to evaluate explanations and may use additional ones when choosing between two reasonable alternatives.

ILLUSIONS OF EXPLANATORY UNDERSTANDING

The pursuit of explanations and the desire to offer them are driven by our intuitions concerning the quality of explanations we already have. In trying to gain an explanatory understanding of a system, one stops upon reaching a “working understanding.” People often claim to have rushes of insight or flashes of understanding, yet these intuitions often may not be accurate. People do have strong senses of making cognitive advances in understanding, ranging from a visceral “rush” of understanding (Gopnik 1998) to being in a cognitive “flow” (Csikszent-mihalyi 1990). The “aha” sense has been discussed ever since Archimedes was described as euphorically shouting, “Eureka!” To some extent, these intuitions must be tracking real progress in understanding; but recent evidence also suggests that these intuitions are only crude indicators that can mislead, sometimes quite dramatically.

People of all ages tend to be miscalibrated with respect to their explanatory understandings; that is, they think they understand in far more detail than they really do how some aspect of the world works or why some pattern in the world exists. This bias, the “illusion of explanatory depth” or IOED (Rozenblit & Keil 2002), can be demonstrated by asking people to rate how well they understand devices such as helicopters, cylinder locks, and zippers. Having given those ratings, they are asked to describe in detail everything they know about how some of the devices work. They are then asked to re-rate their understanding (*a*) in light of having given an explanation, (*b*) having attempted to answer a critical diagnostic question (e.g., How do you pick a lock?), and finally, (*c*) having been given a concise expert

explanation of how the device works. (The last measure helps assess whether they had a more implicit understanding they couldn't verbalize in the earlier trials but that fosters a surge of recognition when the expert explanation is presented.)

There is a consistent strong IOED, with later ratings being substantially lower than earlier ones (Rozenblit & Keil 2002). This drop in ratings, however, does not occur for self-assessments of other kinds of knowledge such as facts (e.g., capitals of countries), procedures (e.g., how to make an international phone call), or narratives (e.g., plots of familiar movies). The IOED is therefore not merely another case of a general overconfidence effect (e.g., Krueger & Dunning 1999, Yates et al. 1997). There is something about partial explanatory understanding that provides particularly compelling senses of knowing more than one really does. The same specificity is found in quite young children as well (Mills & Keil 2004).

Why should explanatory understanding mislead us into thinking we know more than we really do, whereas other forms of knowledge are much more accurately assessed? Several distinctive properties of explanations, many of them described above, seem to be involved (Keil et al. 2004). Explanations are especially vulnerable to function-for-mechanism confusions because of their hierarchical structure. Explanations also have relatively unclear end states, making it more difficult to assess success. Moreover, explanations rarely are offered or heard in much detail, which makes it difficult to have much practice in assessing one's own explanatory understanding. Finally, we may confuse our ability to represent or to know fully an explanation in our heads with the ability to fuse more fragmentary understandings with situational cues that fill out critical explanatory detail when an object is directly in front of us in real time and is available for inspection (Clark 1997). These factors are not as dominant for procedures, facts, or narratives.

Consider in a bit more detail the function-for-mechanism confusions. Many complex systems both in nature and in the manufactured world can be understood as hierarchically organized sets of relatively autonomous units or "stable subassemblies" (Simon 1996, 2000). With most artifacts and with many systems in biology, there are salient functional relations between the highest levels of these systems. Thus, a computer mouse controls the movement of a cursor on a screen, a carburetor mixes fuel and air, and a kidney filters blood. When we learn one of these causal-functional relations, we get an appropriate surge of insight into an explanatory relation that we did not have before. The problem occurs when we attach that surge of insight to an inappropriately low level of analysis. That is, we might assume that we have gained some insight into the workings of the mouse, the ways in which fuel and air are mixed, or how the kidney actually processes blood. Moreover, if we have achieved some success at deciphering a set of mechanistic relations for a system in real time when it is in front of us, we may mistake that success and the genuine understanding we have at higher functional levels with having a stored mental representation of lower level, more mechanistic relations.

People also seem to use misleading heuristics to assess how well they understand a system. Most notably, if they can see or easily visualize several components of a system, they are more convinced they know how it works. Thus, the more easily visible are parts in a system, relative to hidden ones, the stronger the IOED (Rozenblit & Keil 2002). Visual influences of this sort may be related to the appeal of visual "mental animations" in constructing and evaluating explanations of devices (Hegarty 1992).

HOW DO WE DEAL WITH GAPS IN EXPLANATORY UNDERSTANDING?

Even though we often think we understand a system in far more detail than we really do, we rarely think we have a fully exhaustive understanding of any complex system (Keil

2003a,b). How, then, do we deal with recognized gaps in explanatory understanding? Sometimes we suppress our ignorance through a memory distortion of the explanatory failure or by a reconstruction of the event itself into a more explanatorily tractable form (see, e.g., Schacter et al. 1995). We might also attempt to fill in the gap by acquiring new details of function and/or mechanism. Finally, we may recognize the gap and decide to “outsource” it to another mind, believing that the outsourced area of understanding supports our own explanations in a safe and reliable manner.

This third way of dealing with gaps is relatively unexplored in psychology, yet may represent the dominant way in which we handle incomplete understandings. We are not isolated learners each on our own desert island of intellectual exploration. We rely heavily on expertise in other minds. The study of distributed cognition has long recognized this fact (e.g., Crowley et al. 2001, Hutchins 1995), but largely has involved demonstrations of how people in groups share cognitive burdens and achieve results not possible as individual problem solvers (Cole & Engestrom 1993, Crowley et al. 2001, Hutchins 1995). With incomplete explanations, however, we often are not in direct interaction with others. Instead, we rely on others who are removed from us in both space and time.

To use others as sources of information, we have to know if they are speaking from an area of expertise or are posturing, bluffing, or otherwise unreliable. Sources can be deficient on several distinct grounds: motivational states, general competence, or competence in the area that is being explained. Motivational states matter when we believe that explainers have conflicts of interest (Walster et al. 1966). The drug company spokesperson who touts the scientific reasons for taking the company’s new drug seems less credible than the independent scientists or consumer advocates making the same claims, but who are normally critical of the drug companies. People do discount explanations when such conflicts of interest are made salient (Miller 1999). Moreover, such discounting occurs early in a child’s development (Mills & Keil 2005). We also discount the quality of a person’s explanation if we feel that the person is incompetent as revealed by being intoxicated, uneducated, or excessively emotional. Finally, it may be important to know when people do not care if what they are saying is true or false and are simply attempting to make another person think that they know what they are talking about (Frankfurt 2005).

Even when sources are not contaminated by self-interest or obvious global incompetence, we may discount them as outside their area of expertise. Such intuitions arise from an ability to track not only the divisions of labor noted by Smith and Durkheim (e.g., Durkheim 1893, Smith 1776), but also the corresponding divisions of cognitive and linguistic labor (Kitcher 2001, Putnam 1975). To surmount gaps in our own knowledge, we need to know what groups of experts are relevant.

Adults and surprisingly young children can often use their coarse, fragmentary explanatory understandings to leverage access to much deeper explanations when needed (Danovitch & Keil 2004, Keil 2005, Lutz & Keil 2002). Thus, if a person has a better-than-average understanding of one phenomenon, we can guess about what else they are likely to know on the basis of our own more limited knowledge of how causal regularities in the world cluster. For example, if a person knows a lot about how spinning tops stay up, adults and children alike will infer that the same person is likely to know more about why basketballs bounce than why animals have to breathe (Keil 2005). They assume that there is a pattern of causal regularities concerning bounded objects in motion such that an expert on one phenomenon arising from those regularities is likely to understand other phenomena arising from the same regularities. One may not know much about those regularities, but instead may have just a shallow impression of distinctive causal patterns corresponding roughly to physical mechanics. In this way, great gaps in understanding can rest on firmer ground. A related line

of work examines how adults and children evaluate testimony, and suggests a similar early emerging sense of how to decide what messages are more reliable than others (Clément et al. 2004, Harris 2002).

HOW DO EXPLANATORY SKILLS DEVELOP?

Explanatory skills develop along several dimensions. Most obvious are the explanations offered by children. A second dimension concerns children's understanding of what makes a good explanation. Finally, there are the developing abilities to evaluate one's own explanatory understanding and the epistemological status of explanations.

Signs of explanatory ability may first be seen when children ask questions about the workings of the world around them. Although some "why" questions can simply be attempts to get parents to change their minds (e.g., "Why can't I have dessert before dinner?"), in many other cases young children ask questions that try to uncover mechanisms or explain inconsistencies. Young preschoolers will ask "why" questions about the causal structure of the world, questions that are shaped by the causal explanations they receive from parents and others (Hickling & Wellman 2001). In their answers, parents tend to focus more on prior causal factors than on consequences, and such answers are indeed the ones being sought (Callanan & Oakes 1992). Children often give off information about their level of understanding in their gestures when they ask questions, information that is not always fully utilized by parents. Thus, when parents are taught how to attend to children's gestures in asking questions, adult explanations will show an improvement in quality and effectiveness (Kelly et al. 2002). The subtlety of children's questions can be quite striking, e.g., preschoolers ask about the functions of artifacts as wholes but only about the parts of animals (Greif et al. 2005).

Almost as early as the first "why" questions are the first attempts at explanations. Although certainly the explanations of three-year-olds differ in complexity and clarity from those of an older child or adult, they will offer causal accounts that seem to be roughly appropriate for the domain involved. Psychological phenomena are explained with psychological relations, physical events with physical causal relations, and biological patterns with biological mechanisms (Gelman 2002, Hickling & Wellman 2001, Inagaki & Hatano 2002, Wellman & Schult 1997). Thus, within a year of having any facility with language, children map causal patterns in the world with kinds of explanations roughly corresponding to those patterns. They do so systematically and consistently for at least the domains of biology, psychology, and physical mechanics and social conventions (Gelman & Kremer 1991, Hickling & Wellman 2001). In one pioneering set of studies in this area, preschoolers reliably chose different kinds of explanatory schemas to account for departures from physical, social conventional, and moral regularities (invoking magic, defective mental states, and ill will, respectively) (Lockhart 1981).

The ability to explain events after the fact may emerge earlier than the ability to predict those events. For example, in explaining the mental behaviors of others, preschoolers sometimes find it easier to explain why a person was deceived or mistaken than that they will be mistaken in the same task (Bartsch 1998, Wimmer & Mayringer 1998). This pattern may be related to the hindsight bias in adults in which people who are given an explanation of an outcome falsely assume they could have predicted that outcome.

In many cases, implicit explanatory understanding seems to precede explicit versions in development, often by several years. This progression can sometimes be uncovered by eye-tracking studies in which children anticipate with their eyes more advanced explanations while explicitly offering more primitive ones (e.g., Clements et al. 2000). Similarly, gestures can reveal an explanatory understanding that hasn't yet appeared in explicit terms (Garber &

Goldin-Meadow 2002). What sort of implicit knowledge, however, should count as explanatory understanding as opposed to a simpler tracking of some causal or relational pattern? Implicit precursors of explicit forms are better candidates than are patterns of causal tracking buried in cognitively impenetrable modules.

As described earlier, children's explanations do suffer, sometimes dramatically, from an inability to gauge correctly the level of understanding in another. Not until quite late in development do children have a mature sense of the "epistemology" of explanations (Kuhn & Franklin 2005). A mature epistemology of explanations sees explanations as falsifiable, recognizes evidence as the usual means for falsification, and realizes that theory and evidence are different and play complementary roles (Kuhn & Franklin 2005). During adolescence, teenagers shift from viewing explanations as absolutist to viewing knowledge as more pluralistic or relativistic (Hofer & Pintrich 2002, Kuhn & Franklin 2005). Younger children assume that there must be one correct explanation and that science or other areas of inquiry proceed by adding more true facts and principles to that explanation. Only later do they appreciate that two different communities of equally reputable scientists and scholars might disagree on which of two competing explanations is correct. Moreover, the transition out of absolutist ways of understanding occurs in a domain-specific manner, for example, occurring earlier for aesthetic judgments than for more fact-based reasoning (Kuhn et al. 2000).

EXPLANATION ACROSS CULTURES

Although cultures may differ in their dominant explanatory styles, they rarely, if ever, have a particular explanatory style completely unavailable to them. That more dated view of cultural differences, often aided and abetted by implicit assumptions of western cultures as being more scientific, analytic, and abstract in explanatory styles, no longer seems plausible (see, e.g., Cole & Means 1981, Sperber & Hirschfeld 1999). Instead, cultural differences are now more commonly thought of as different default hierarchies. A particular explanatory style may be the first option in one culture and a less-prominent option in another, but rarely is one form of explanation completely unavailable. One well-known example of explanatory differences occurs in cultures that are considered collectivist or concerned more with the rights of the group relative to those of the individual, versus cultures that are individualist and that have the opposite emphasis (Triandis 1995). Collectivist cultures tend to provide more situational explanations about behaviors, whereas individualist cultures tend to provide more dispositional explanations. Extending beyond the social realm, even inanimate objects, such as the movement of a stick in a stream, may be explained in more situational terms in a collectivist culture and more dispositionally in an individualist culture (Nisbett 2003). Do these different orderings of explanatory styles hold across all domains, e.g., a situational bias for phenomena ranging from physical mechanics to social interactions, or do the orderings themselves differ across domains and cultures in a more complex interactional pattern? Based on analogous results for moral reasoning (e.g., Turiel 1998), the more complex interactional result seems likely.

Different explanatory styles may also emerge as a function of the predominant practices of a culture. Cultures using forms of agriculture that are sensitive to soil conservation might adopt more ecologically oriented forms of biological explanation that see more commonalities in explanations for both plants and animals than other groups using slash and burn agriculture. Similarly, people may reason about the same groups of fish in goal-centered versus ecological ways because of different cultural lenses (Medin & Atran 2004, Medin et al. 2005). More subtly, some groups in a domain may make inductions based on reference to ideals in a category, while others may make similar inductions based on category-central prototypes (Lynch et al. 2000).

A different pattern of cultural variation in explanatory form happens when a core explanatory idea is shared, but it becomes elaborated on in different ways across cultures. Consider, for example, vitalism, the tendency to assume that living things have vital forces inside them that are responsible for growth and, in animals, for activity (Inagaki & Hatano 2002). Many cultures may share this explanatory framework but will fill it out differently. One may intentionalize agents within living things, whereas another might use a notion of fluid energy (Inagaki & Hatano 2002, 2004). In a related vein, it appears that all cultures adopt a kind of Cartesian dualism in explanations about minds and bodies, but they fill out the details in very different ways (Bloom 2004). Many turn to religious notions of the soul, others invoke animistic spirits, while still others may rely on a shared being in several bodies at once. These examples illustrate a differentiation model in which a core explanatory schema is shared by all cultures but can become manifested in quite different concrete ways.

The influence of culture, context, and even social class on explanatory styles is clearly a growth area for research. Already, however, it seems likely that the predominant pattern will be one in which no culture is completely blocked from certain forms of explanatory styles or schemes. Instead, it seems more plausible that cultures will vary in the ordering preferences of such explanatory forms in each domain and in terms of how they fill out the more mechanistic details of each scheme.

CONCLUSIONS

The processes of constructing and understanding explanations are intrinsic to our mental lives from an early age, with some sense of explanatory insight present before children are even able to speak. A focus on explanation and understanding brings to the fore issues and perspectives that are not as prominent in discussions of related topics such as intuitive theories and mental models. Explanations are sought after and provided not only in interpersonal interactions but also within the mind of a single individual. There is considerable diversity in kinds of explanations in terms of the causal patterns they invoke, the broad stances they employ, and the more local domains in which they occur. Qualitatively different patterns of explanation seem to be used in talking about domains such as physical mechanics, biological function, and social interactions. A major challenge remains in specifying in more detail the structures of these explanatory types and demarcating the sorts of domains that are responsible for such structural differences. Explanatory types may also vary as a function of the different uses of explanations beyond those of prediction, such as for justification and rationalization, culpability determination, and aesthetic insight.

The causal and relational complexity inherent in much of the world makes many explanations necessarily incomplete or flawed. We therefore must rely on coarser gists that provide effective explanatory frameworks while nonetheless missing many details. We are adept at supplementing these gists by “outsourcing” knowledge to other minds and relying on the divisions of cognitive labor that occur in all cultures. People also use a wide variety of heuristics to recognize bad explanations and additional ones to choose among competing explanations that may be reasonable. These skills, however, do not prevent us from frequently overestimating the depth of our explanatory understanding, often dramatically, in ways that are systematically related to the entity or system being explained. All of us throughout the world may share the same drive for explanation, the same assortment of explanatory styles and strategies for dealing with gaps, and similar developmental patterns. Our differences may lie in which explanatory styles come to mind first in specific contexts, not in terms of fundamental explanatory abilities.

Acknowledgments

Preparation of this review and some of the research described herein were supported by NIH grant R-37-HD023922 to Frank Keil. Thanks to Esther Schlegel for help in manuscript preparation.

LITERATURE CITED

- Ahn W. Why are different features central for natural kinds and artifacts? The role of causal status in determining feature centrality. *Cognition* 1998;69:135–78. [PubMed: 9894403]
- Ahn, W.; Kalish, CW. The role of mechanism beliefs in causal reasoning. 2000. p. 199-226. See Keil & Wilson 2000b
- Ahn W, Kim NS, Lassaline ME, Dennis MJ. Causal status as a determinant of feature centrality. *Cogn Psychol* 2000;41:361–416. [PubMed: 11121260]
- Ahn W, Novick L, Kim NS. “Understanding it makes it more normal”: Causal explanations influence person perception. *Psychon Bull Rev* 2003;10:746–52. [PubMed: 14620373]
- Allen, C.; Bekoff, M.; Lauder, G., editors. *Nature’s Purposes: Analyses of Function and Design in Biology*. Cambridge, MA: MIT Press; 1998.
- Amini M. Has foundationalism failed? A critical review of *Coherence in Thought and Action* by Paul Thagard. *Hum Nat Rev* 2003;3:119–23.
- Barsalou L. Perceptual symbol systems. *Behav Brain Sci* 1999;22:577–609. [PubMed: 11301525]
- Bartsch K. False belief prediction and explanation: which develops first and why it matters. *Int J Behav Dev* 1998;22(2):423–28.
- Bechtel W. Levels of description and explanation in cognitive science. *Minds Mach* 1994;4:1–25.
- Bechtel W, Abrahamsen A. Explanation: a mechanistic alternative. *Stud Hist Philos Sci C Stud Hist Philos Biol Biomed Sci* 2005;36:421–41.
- Berk LE. Why children talk to themselves. *Sci Am* 1994;271(5):78–83. [PubMed: 7997866]
- Birch SA, Bloom P. Children are cursed: an asymmetric bias in mental state attribution. *Psychol Sci* 2003;14:283–86. [PubMed: 12741755]
- Bloom P. Intention, history, and artifact concepts. *Cognition* 1996;60:1–29. [PubMed: 8766388]
- Bloom, P. *Descartes’ Baby: How the Science of Child Development Explains What Makes Us Human*. New York: Basic Books; 2004.
- Bloom P, German T. Two reasons to abandon the false belief task as a test of theory of mind. *Cognition* 2000;77:B25–31. [PubMed: 10980256]
- Boyd, R. Homeostasis, species, and higher taxa. In: Wilson, R., editor. *Species: New Interdisciplinary Studies*. Cambridge, MA: MIT Press; 1999. p. 141-85.
- Brem SK, Rips LJ. Explanation and evidence in informal argument. *Cogn Sci* 2000;24:573–604.
- Bromberger, S. Why-questions. In: Brody, BA., editor. *Readings in the Philosophy of Science*. Englewood Cliffs, NJ: Prentice Hall; 1966. p. 66-84.
- Callanan M, Oakes L. Preschoolers’ questions and parents’ explanations: causal thinking in everyday activity. *Cogn Dev* 1992;7:213–33.
- Carey, S. *Conceptual Change in Childhood*. Cambridge, MA: Bradford Books/MIT Press; 1985.
- Cartwright, N. *The Dappled World: A Study of the Boundaries of Science*. London: Cambridge Univ. Press; 1999.
- Cartwright, N. From causation to explanation and back. In: Leiter, B., editor. *The Future for Philosophy*. London: Oxford Univ. Press; 2004. p. 230-45.
- Cavanaugh, P. Pictorial art and vision. In: Wilson, RA.; Keil, FC., editors. *MIT Encyclopedia of Cognitive Sciences*. Cambridge, MA: MIT Press; 1999. p. 648-51.
- Chater, N.; Oaksford, M. Mental mechanisms: speculations on human causal learning and reasoning. In: Fiedler, K.; Juslin, P., editors. *Information Sampling and Adaptive Cognition*. London: Cambridge Univ. Press; 2005. In press
- Chater N, Vitanyi P. Simplicity: a unifying principle in cognitive science? *Trends Cogn Sci* 2003;7(1): 19–22. [PubMed: 12517354]
- Cheng PW. From covariation to causation: a causal power theory. *Psychol Rev* 1997;104:367–405.

- Chinn CA, Brewer WF. The role of anomalous data in knowledge acquisition: a theoretical framework and implications for science instruction. *Rev Educ Res* 1993;63:1–49.
- Church RB, Goldin-Meadow S. The mismatch between gesture and speech as an index of transitional knowledge. *Cognition* 1986;23:43–71. [PubMed: 3742990]
- Clark, A. *Being There: Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press; 1997.
- Clark, HH. *Using Language*. London: Cambridge Univ. Press; 1996.
- Clément F, Koenig M, Harris PL. The ontogenesis of trust. *Mind Lang* 2004;19:360–79.
- Clements W, Rustin CL, McCallum S. Promoting the transition from implicit to explicit understanding: a training study of false belief. *Dev Sci* 2000;3(1):81–92.
- Cole, M.; Engstrom, Y. A cultural approach to distributed cognition. In: Salomon, G., editor. *Distributed Cognitions: Psychological and Educational Considerations*. New York: Cambridge Univ. Press; 1993. p. 1-46.
- Cole, M.; Means, B. *Comparative Studies of How People Think: An Introduction*. Cambridge, MA: Harvard Univ. Press; 1981.
- Collins, J.; Hall, E.; Paul, L. *Causation and Counterfactuals*. Cambridge, MA: MIT Press; 2001.
- Cosmides, L.; Tooby, J. From evolution to behavior: evolutionary psychology as the missing link. In: Dupre, J., editor. *The Latest on the Best: Essays on Evolution and Optimality*. Cambridge, MA: MIT Press; 1987. p. 277-306.
- Crowley K, Callanan MA, Jipson J, Galco J, Topping K, Shrager J. Shared scientific thinking in everyday parent-child activity. *Sci Educ* 2001;85:712–32.
- Cummins, R. “How does it work?” versus “What are the laws?” Two conceptions of psychological explanation. 2000. p. 117-45. See Keil & Wilson 2000b
- Cziksentmihalyi, M. *Flow: The Psychology of Optimal Experience*. New York: HarperCollins; 1990.
- Danovitch JH, Keil FC. Should you ask a fisherman or a biologist? Developmental shifts in ways of clustering knowledge. *Child Dev* 2004;5:918–31. [PubMed: 15144494]
- Danovitch, JH.; Keil, FC. Scientists or saints: the emergence of an understanding of the mental characteristics needed to solve moral or scientific problems. Poster presented at 2005 Meet. Soc. Child Dev; Atlanta, GA. 2005.
- Dennett, D. *The Intentional Stance*. Cambridge, MA: MIT Press; 1987.
- Dickinson A. Causal learning: an associative analysis. *Q J Exp Psychol* 2001;54B:3–25.
- di Sessa A. Towards an epistemology of physics. *Cogn Instruct* 1993;10:105–225.
- di Sessa A, Gillespie NM, Esterly JB. Coherence versus fragmentation in the development of the concept of force. *Cogn Sci* 2004;28:843–900.
- Dunbar, K. How scientists really reason: scientific reasoning in real-world laboratories. In: Sternberg, R.; Davidson, J., editors. *The Nature of Insight*. Cambridge, MA: MIT Press; 1995. p. 365-96.
- Dunbar K, Blanchette I. The in vivo/in vitro approach to cognition: the case of analogy. *Trends Cogn Sci* 2001;5:334–39. [PubMed: 11477002]
- Durkheim, E. *The Division of Labor in Society*. Simpson, G., translator. New York: Free Press; 1893/1997. (From French)
- Fauconnier, G.; Turner, M. *The Way We Think: Conceptual Blending and the Mind’s Hidden Complexities*. New York: Basic Books; 2002.
- Ferguson, N., editor. *Virtual History: Alternatives and Counterfactuals*. New York: Basic Books; 1999. First publ. 1997 by Picador, London
- Fodor, JA. *The Language of Thought*. New York: Thomas Crowell; 1975.
- Fodor, JA. *Concepts: Where Cognitive Science Went Wrong*. New York: Oxford Univ. Press; 1998.
- Frankfurt, H. *On Bullshit*. Princeton, NJ: Princeton Univ. Press; 2005.
- Garber P, Goldin-Meadow S. Gesture offers insight into problem-solving in adults and children. *Cogn Sci* 2002;26:817–31.
- Gelman, R. Cognitive development. In: Pashler, H.; Medin, DL.; Gallistel, R.; Wixted, J., editors. *Stevens’ Handbook of Experimental Psychology*. 3. Vol. 3. New York: Wiley; 2002. p. 396-443.
- Gelman, SA. *The Essential Child*. New York: Oxford Univ. Press; 2003.

- Gelman SA, Kremer KA. Understanding natural cause: children's explanations of how objects and their properties originate. *Child Dev* 1991;62(2):396–414. [PubMed: 2055130]
- Gentner D. Structure-mapping: a theoretical framework for analogy. *Cogn Sci* 1983;7(2):155–70.
- Gentner, D.; Gentner, D. Flowing waters or teeming crowds: mental models of electricity. In: Gentner, D.; Stevens, AL., editors. *Mental Models*. Hillsdale, NJ: Erlbaum; 1983. p. 99-129.
- Gentner, D.; Rattermann, MJ.; Markman, AB.; Kotovsky, L. Two forces in the development of relational structure. In: Simon, T.; Halford, G., editors. *Developing Cognitive Competence: New Approaches to Process Modeling*. Hillsdale, NJ: Erlbaum; 1995. p. 263-313.
- Gentner, D.; Stevens, A. *Mental Models*. Hillsdale, NJ: Erlbaum; 1983.
- Gentner D, Toupin C. Systematicity and surface similarity in the development of analogy. *Cogn Sci* 1986;10:277–300.
- Gergely G, Nadasdy Z, Csibra G, Biro S. Taking the intentional stance at 12 months of age. *Cognition* 1995;56(2):165–93. [PubMed: 7554793]
- German TP, Johnson SA. Function and the origins of the design stance. *J Cogn Dev* 2002;3:279–300.
- Glennan S. Rethinking mechanistic explanation. *Philos Sci* 2002;69:S342–53.
- Glymour C. Learning causes: psychological explanations of causal explanations. *Minds Machines* 1998;8(1):39–60.
- Glymour, C. *The Mind's Arrows: Bayes Nets and Graphical Causal Models in Psychology*. Cambridge, MA: MIT Press; 2001.
- Gopnik A. Explanation as orgasm. *Minds Machines* 1998;8(1):101–18.
- Gopnik A, Glymour C, Sobel D. Causal maps and Bayes nets: a theory of causal inference in young children. *Psychol Rev* 2004;111(1):1–30.
- Gould SJ, Lewontin RC. The spandrels of San Marco and the Panglossian paradigm: a critique of the adaptationist programme. *Proc R Soc Lond B* 1979;205:581–98. [PubMed: 42062]
- Graesser AC, Olde BA. How does one know whether a person understands a device? The quality of the questions the person asks when the device breaks down. *J Educ Psychol* 2003;95:524–36.
- Greene J, Haidt J. How (and where) does moral judgment work? *Trends Cogn Sci* 2002;6(12):517–23. [PubMed: 12475712]
- Greif, M.; Gutierrez, F.; Keil, F.; Kemler-Nelson, D. What do children want to know about animals and artifacts? Domain-specific requests for information. Poster presented at 2005 Meet. Soc. Child Dev; Atlanta, GA. 2005.
- Grice, HP. Logic and conversation. In: Cole, P.; Morgan, J., editors. *Syntax and Semantics*. New York: Academic; 1975. p. 41-58.
- Griffiths, TL.; Baraff, ER.; Tenenbaum, JB. Using physical theories to infer hidden causal structure. *Proc. 26th Annu. Conf. Cogn. Sci. Soc*; Chicago. Mahwah, NJ: Erlbaum; 2004. p. 446-51.
- Haidt J. The emotional dog and its rational tail. *Psychol Rev* 2001;108:814–34. [PubMed: 11699120]
- Harman G. The inference to the best explanation. *Philos Rev* 1965;74:88–95.
- Harris, PL. What do children learn from testimony?. In: Carruthers, P.; Stich, SP.; Siegal, M., editors. *The Cognitive Basis of Science*. London: Cambridge Univ. Press; 2002. p. 316-34.
- Hegarty M. Mental animation: inferring motion from static displays of mechanical systems. *J Exp Psychol Learn Mem Cogn* 1992;18:1084–102. [PubMed: 1402712]
- Heider, F. *The Psychology of Interpersonal Relations*. New York: Wiley; 1958.
- Hempel, CG. *Aspects of Scientific Explanation*. New York: Free Press; 1965.
- Hempel CG, Oppenheim P. Studies in the logic of explanation. *Philos Sci* 1948;15:135–75.
- Hickling AK, Wellman HM. The emergence of children's causal explanations and theories: evidence from everyday conversation. *Dev Psychol* 2001;37:668–83. [PubMed: 11552762]
- Hofer, B.; Pintrich, P. *Epistemology: The Psychology of Beliefs About Knowledge and Knowing*. Mahwah, NJ: Erlbaum; 2002.
- Holyoak, KJ.; Thagard, P. *Mental Leaps: Analogy in Creative Thought*. Cambridge, MA: MIT Press; 1995.
- Home, RW. *The Effluvial Theory of Electricity*. New York: Arno; 1981.

- Humphreys, P. *The Chances of Explanation: Causal Explanations in the Social, Medical, and Physical Sciences*. Princeton, NJ: Princeton Univ. Press; 1989.
- Hutchins, E. *Cognition in the Wild*. Cambridge, MA: MIT Press; 1995.
- Inagaki, K.; Hatano, G. *Young Children's Naïve Thinking About the Biological World*. New York: Psychol. Press; 2002.
- Inagaki K, Hatano G. Vitalistic causality in young children's naive biology. *Trends Cogn Sci* 2004;8:356–62. [PubMed: 15335462]
- Johnson-Laird, PN. *Mental Models*. Cambridge, MA: Harvard Univ. Press; 1983.
- Kahneman D, Miller DT. Norm theory: comparing reality to its alternatives. *Psychol Rev* 1986;93:136–53.
- Kahneman, D.; Tversky, A. The simulation heuristic. In: Kahneman, D.; Slovic, P.; Tversky, A., editors. *Judgment Under Uncertainty: Heuristics and Biases*. New York: Cambridge Univ. Press; 1982. p. 201–8.
- Kayoko I, Giyoo H. Vitalistic causality in young children's naive biology. *Trends Cogn Sci* 2004;8(8): 356–62. [PubMed: 15335462]
- Keil, FC. *Concepts, Kinds and Cognitive Development*. Cambridge, MA: MIT Press; 1989.
- Keil, FC. The emergence of an autonomous biology. In: Gunnar, M.; Maratsos, M., editors. *Modularity and Constraints in Language and Cognition: The Minnesota Symposia*. Hillsdale, NJ: Erlbaum; 1992. p. 103–38.
- Keil, FC. The growth of causal understandings of natural kinds: modes of construal and the emergence of biological thought. In: Premack, A.; Sperber, D., editors. *Causal Cognition*. New York: Oxford Univ. Press; 1995. p. 234–62.
- Keil FC. Categorization, causation and the limits of understanding. *Lang Cogn Process* 2003a;18:663–92.
- Keil FC. Folkscience: coarse interpretations of a complex reality. *Trends Cogn Sci* 2003b;7:368–73. [PubMed: 12907233]
- Keil FC. That's life: coming to understand biology. *Hum Dev* 2003c;46:369–77.
- Keil FC. Doubt, deference and deliberation. *Oxf Stud Epistemol*. 2005 In press.
- Keil, FC.; Greif, M.; Kerner, R. A world apart: artifacts. In: Margolis, E., editor. *Creations of the Mind: Essays on Artifacts and Their Representation*. Oxford Univ. Press; 2005. In press
- Keil, FC.; Rozenblit, LR.; Mills, C. What lies beneath? Understanding the limits of understanding. In: Levin, DT., editor. *Thinking and Seeing: Visual Metacognition in Adults and Children*. Cambridge, MA: MIT Press; 2004. p. 227–49.
- Keil FC, Smith C, Simons DJ, Levin DT. Two dogmas of conceptual empiricism: implications for hybrid models of the structure of knowledge. *Cognition* 1998;65:103–35. [PubMed: 9557380]
- Keil, FC.; Wilson, RA. Explaining explanation. 2000a. p. 1–18. See Keil & Wilson 2000b
- Keil, FC.; Wilson, RA. *Explanation and Cognition*. Cambridge, MA: MIT Press; 2000b.
- Kelemen D. Function, goals, and intention: children's teleological reasoning about objects. *Trends Cogn Sci* 1999;3:461–68. [PubMed: 10562725]
- Kelemen, D.; Carey, S. The essence of artifacts: developing the design stance. In: Margolis, E.; Lawrence, S., editors. *Creation of the Mind: Essays on Artifacts and Their Representation*. New York: Oxford Univ. Press; 2005. In press
- Kelly SD, Singer M, Hicks J, Goldin-Meadow S. A helping hand in assessing children's knowledge: instructing adults to attend to gesture. *Cogn Instruct* 2002;20(1):1–26.
- Kemp, CS.; Perfors, A.; Tenenbaum, JB. Learning domain structures. *Proc. 26th Annu. Conf. Cogn. Sci. Soc*; Chicago. Mahwah, NJ: Erlbaum; 2004. p. 663–68.
- Keysar B, Barr DJ, Horton WS. The egocentric basis of language use: insights from a processing approach. *Curr Dir Psychol Sci* 1998;7:46–50.
- Kim NS, Keil FC. From symptoms to causes: diversity effects in diagnostic reasoning. *Mem Cogn* 2003;31:155–65.
- Kitcher, P. *Science, Truth and Democracy*. London: Oxford Univ. Press; 2001.

- Kozhevnikov M, Hegarty M. Impetus beliefs as default heuristics: dissociation between explicit and implicit knowledge about motion. *Psychon Bull Rev* 2001;8(3):439–53. [PubMed: 11700895]
- Krauss RM, Glucksberg S. The development of communication: competence as a function of age. *Child Dev* 1969;40:255–66.
- Krueger J, Dunning D. Unskilled and unaware of it: how difficulties in recognizing one's own incompetence lead to inflated self-assessments. *J Personal Soc Psychol* 1999;77:1121–34.
- Krynski, TR.; Tenenbaum, JB. The role of causal models in reasoning under uncertainty. *Proc. 25th Annu. Conf. Cogn. Sci. Soc.*; London: LEA. 2003. p. 692-97.
- Kuhn D, Black J, Keselman A, Kaplan D. The development of cognitive skills to support inquiry learning. *Cogn Instruct* 2000;18:495–523.
- Kuhn, D.; Franklin, S. The second decade: What develops (and how)? In: Damon, W.; Lerner, R.; Kuhn, D.; Siegler, RS., editors. *Handbook of Child Psychology*. 6. Vol. 2. 2005. In press
- Lehman DR, Lempert RO, Nisbett RE. The effects of graduate training on reasoning: formal discipline and thinking about everyday life events. *Am Psychol* 1988;43:431–43.
- Lewis, D. *Counterfactuals*. Cambridge, MA: Harvard Univ. Press; 1973.
- Lockhart KL. The development of knowledge about uniformities in the environment: a comparative analysis of the child's understanding of social, moral, and physical rules. *Dissert Abstr Int* 1981;41:2793.
- Lombrozo, T. Teleological explanation: causal constraints and regularities. Paper presented at 1st joint meet. Soc. Philos. Psychol. Eur. Soc. Philos. Psychol; Barcelona, Spain. 2004.
- Lombrozo T, Carey S. Functional explanation and the function of explanation. *Cognition*. 2005 In press.
- Luhmann CC, Ahn WK, Palmeri TJ. Theories and similarity: categorization under speeded conditions. *Proc 24th Annu Conf Cogn Sci Soc* 2002;590–95.
- Mahwah NJ, Erlbaum Lutz DR, Keil FC. Early understanding of the division of cognitive labor. *Child Dev* 2002;73:1073–84. [PubMed: 12146734]
- Lynch EB, Coley JD, Medin DL. Tall is typical: central tendency, ideal dimensions and graded category structure among tree experts and novices. *Mem Cogn* 2000;28(1):41–50.
- Machamer P, Darden L, Craver C. Thinking about mechanisms. *Philos Sci* 2000;67:1–25.
- Macrae CN, Bodenhausen GV. Social cognition: thinking categorically about others. *Annu Rev Psychol* 2000;51:93–120. [PubMed: 10751966]
- Malle, BF. *How the Mind Explains Behavior: Folk Explanations, Meaning, and Social Interaction*. Cambridge, MA: MIT Press; 2004.
- Markman EM. Realizing that you don't understand: a preliminary investigation. *Child Dev* 1977;48(3):986–92.
- Medin DL, Atran S. The native mind: biological categorization, reasoning and decision making in development across cultures. *Psychol Rev* 2004;111:960–83. [PubMed: 15482069]
- Medin DL, Ross N, Atran S, Cox D, Wakaua HJ, et al. The role of culture in the folk-biology of freshwater fish. *Cogn Psychol*. 2005 In press.
- Miller DT. The norm of self-interest. *Am Psychol* 1999;54:1053–60. [PubMed: 15332526]
- Mills C, Keil FC. Knowing the limits of one's understanding: the development of an awareness of an illusion of explanatory depth. *J Exp Child Psychol* 2004;87:1–32. [PubMed: 14698687]
- Mills C, Keil FC. The development of cynicism. *Psychol Sci* 2005;16:385–90. [PubMed: 15869698]
- Murphy, GL. *The Big Book of Concepts*. Cambridge: MIT Press; 2002.
- Murphy GL, Medin DL. The role of theories in conceptual coherence. *Psychol Rev* 1985;92:289–316. [PubMed: 4023146]
- Nickerson RS. The projective way of knowing. *Curr Dir Psychol Sci* 2001;10:168–72.
- Nisbett, R. *The Geography of Thought*. New York: Free Press; 2003.
- Oppenheimer DM. Spontaneous discounting of availability in frequency judgment tasks. *Psychol Sci* 2004;15:100–5. [PubMed: 14738516]
- Piaget, J. *The Language and Thought of the Child*. New York: Routledge & Kegan Paul; 1926.

- Putnam, H. The meaning of meaning. In: Gunderson, K., editor. *Language, Mind and Knowledge*. Minneapolis: Univ. Minn. Press; 1975. p. 131-93.
- Rehder B, Hastie R. Category coherence and category-based property induction. *Cognition* 2004;91:113–53. [PubMed: 14738770]
- Rips LJ. Circular reasoning. *Cogn Sci* 2002;26:767–95.
- Roese NJ. Counterfactual thinking. *Psychol Bull* 1997;121:133–48. [PubMed: 9000895]
- Rosch E, Mervis C, Gray W, Johnson D, Boyes-Braem P. Basic objects in natural categories. *Cogn Psychol* 1976;8:382–439.
- Rozenblit LR, Keil FC. The misunderstood limits of folk science: an illusion of explanatory depth. *Cogn Sci* 2002;26:521–62.
- Salmon W. Probabilistic causality. *Pacific Philos Q* 1980;61:50–74.
- Salmon, W. *Scientific Explanation and the Causal Structure of the World*. Princeton, NJ: Princeton Univ. Press; 1984.
- Salmon, W. Four decades of scientific explanation. In: Kitcher, P.; Salmon, W., editors. *Scientific Explanation*. Minneapolis: Univ. Minn. Press; 1989. p. 3-219.
- Schacter, DL.; Coyle, JT.; Fischbach, GD.; Mesulam, MM.; Sullivan, LE., editors. *Memory Distortion: How Minds, Brains, and Societies Reconstruct the Past*. Cambridge, MA: Harvard Univ. Press; 1995.
- Schank, RC.; Abelson, RP. *Scripts, Plans, Goals and Understanding*. Hillsdale, NJ: Erlbaum; 1977.
- Simon, HA. *Science of the Artificial*. 3. Cambridge, MA: MIT Press; 1996.
- Simon, HA. Discovering explanations. In: Wilson, R.; Keil, F., editors. *Cognition and Explanation*. Cambridge, MA: MIT Press; 2000. p. 21-59.
- Sloman SA, Love BC, Ahn W. Feature centrality and conceptual coherence. *Cogn Sci* 1998;22:189–228.
- Smith, A. *An Inquiry into the Nature and Causes of the Wealth of Nations*. London: Methuen; 1776/1904.
- Sobel D, Tenenbaum JB, Gopnik A. Children's causal inferences from indirect evidence: backwards blocking and Bayesian reasoning in preschoolers. *Cogn Sci* 2004;28:303–33.
- Sober, E. *Simplicity*. Oxford, UK: Clarendon; 1975.
- Sober E. Common cause explanation. *Philos Sci* 1984;51:212–41.
- Spellman BA, Mandel DR. When possibility informs reality: counterfactual thinking as a cue to causality. *Curr Dir Psychol Sci* 1999;8:120–23.
- Sperber, D.; Hirschfeld, LR. Culture, cognition and evolution. In: Wilson, R.; Keil, F., editors. *MIT Encyclopedia of the Cognitive Sciences*. Cambridge, MA: MIT Press; 1999. p. cxi-cxxxii.
- Thagard, P. *Coherence in Thought and Action*. Cambridge, MA: MIT Press; 2000.
- Thagard P, Verbeurgt K. Coherence as constraint satisfaction. *Cogn Sci* 1998;22:1–24.
- Triandis, HC. *Individualism and Collectivism*. Boulder CO: Westview; 1995.
- Turiel, E. *Handbook of Child Psychology*. 5. Vol. 3. New York: Wiley; 1998. Moral development; p. 863-932.
- Walster E, Aronson E, Abrahams D. On increasing the persuasiveness of a low-prestige communicator. *J Exp Soc Psychol* 1966;2:325–42.
- Wellman HM, Schult CA. Explaining human movements and actions: children's understanding of the limits of psychological explanation. *Cognition* 1997;62:291–324. [PubMed: 9187061]
- Whewell, W. *The Philosophy of the Inductive Sciences*. London: Cass; 1847.
- Wilson, D.; Sperber, D. Relevance theory. In: Ward, G.; Horn, L., editors. *Handbook of Pragmatics*. Oxford, UK: Blackwell Sci; 2004. p. 607-32.
- Wilson RA, Keil FC. The shadows and shallows of explanation. *Minds Machines* 1998;8:137–59.
- Wilson TD, Centerbar DB, Kermer DA, Gilbert DT. The pleasures of uncertainty: prolonging positive moods in ways people do not anticipate. *J Personal Soc Psychol* 2005;88:5–21.
- Wimmer H, Mayringer H. False belief understanding in young children: Explanations do not develop before predictions. *Int J Behav Dev* 1998;22:403–22.

- Winer GA, Cottrell JE, Gregg V, Fournier JS, Bica LA. Fundamentally misunderstanding visual perception: adults' beliefs in visual emissions. *Am Psychol* 2002;57:417–24. [PubMed: 12094435]
- Yates JF, Lee JW, Bush JG. General knowledge overconfidence: cross-national variations, response style, and “reality”. *Organ Behav Hum Decis Process* 1997;70(2):87–94 .*Annual Review of Psychology* 2006;57