Review TRENDS in Cognitive Sciences Vol.10 No.10

The structure and function of explanations

Tania Lombrozo

Department of Psychology, University of California at Berkeley, Berkeley, CA 94720, USA

Generating and evaluating explanations is spontaneous, ubiquitous and fundamental to our sense of understanding. Recent evidence suggests that in the course of an individual's reasoning, engaging in explanation can have profound effects on the probability assigned to causal claims, on how properties are generalized and on learning. These effects follow from two properties of the structure of explanations: explanations accommodate novel information in the context of prior beliefs, and do so in a way that fosters generalization. The study of explanation thus promises to shed light on core cognitive issues, such as learning, induction and conceptual representation. Moreover, the influence of explanation on learning and inference presents a challenge to theories that neglect the roles of prior knowledge and explanation-based reasoning.

Explanation and cognition

Children, adults and scientists alike confront the world with a common question: why? We wonder why events unfold in particular ways, why objects have specific properties and why people behave as they do. Explanations are more than a human preoccupation – they are central to our sense of understanding, and the currency in which we exchange beliefs. Accordingly, social psychology and philosophy have subfields dedicated to the study of explanation, with social psychology focusing on explanations of behavior [1–3], and philosophy on explanation in science [4]. Yet, only in the past few years has explanation come of age as a topic of study in cognitive psychology [5,6]. As a result, basic empirical questions about the nature of explanation are just beginning to be addressed, including what constitutes an explanation, what makes some explanations better than others, how explanations are generated and when explanations are sought.

Two developments are responsible for the recent interest in explanation within cognitive psychology. First, prominent theories of concepts and conceptual representation accord a central role to explanation [7–10]. Already, explanations have been shown to facilitate category learning [11], influence judgements of the typicality of category members [11,12] and foster conceptual coherence [8,13]. Second, cognitive psychologists have increasingly recognized that prior knowledge, and in particular causal knowledge, exerts a profound influence on learning and inference [14,15]. The intimate relationship between explanation and causation naturally raises the question of the role of explanation in these effects of prior knowledge.

This article reviews recent work on the cognitive psychology of explanation, with a focus on the generation and evaluation of explanations in the course of an individual's reasoning. A variety of interesting and related topics are not addressed, including the phenomenology of explanation [16–18], the pragmatic and communicative aspects of explanation [19], the developmental trajectory of explanations [20,21] and explanation in artificial intelligence [22]. The reader is directed elsewhere for reviews of the literature on explanation in social psychology [2] and philosophy of science [4]. The aim here is to demonstrate that explanations have an important role in learning and inference. In particular, explanations are spontaneously generated and employed as a basis for constraining inference and guiding generalization. The section that follows considers the structure of explanations; I then turn to evidence for the role of explanation in learning and inference.

The structure of explanations

Although the content of explanations can vary wildly, their structure is more constrained. Theories of explanation from the philosophy of science have generally imposed logical or causal constraints on what constitutes an explanation [4]. For example, one early but influential account characterized explanations as arguments demonstrating how what is being explained ('the explanandum') follows deductively from natural laws and empirical conditions [23]. Several contemporary accounts suggest that explanations identify all or a subset of the causes of the explanandum [24,25], or show the explanandum to be an instance of a general pattern [26]. Subsuming the explanandum under patterns or regularities has also been proposed as a mechanism for determining which of the causes of the explanandum are explanatory [27].

Mirroring accounts of explanation from philosophy, psychological evidence supports the predominance of causation in explanation [28] (Box 1), but also the role of pattern subsumption [29,30]. In particular, explanations typically appeal to causes, although knowledge of general patterns constrains which causes are judged probable [28] and relevant [29,31]. For example, in explaining a forest fire by appeal to lightning, one concurrently indicates a cause (i.e. the lightning), presupposes a broader regularity of which the explanandum is an instance (i.e. that

Corresponding author: Lombrozo, T. (lombrozo@berkeley.edu).
Available online 30 August 2006.

www.sciencedirect.com 1364-6613/$ – see front matter doi:10.1016/j.tics.2006.08.004

Box 1. Beyond efficient cause

Aristotle identified four causes, or 'modes of explanation', that pick out different aspects of an answer to a why-question: the efficient, final, formal and material causes (Table I). The majority of empirical research has focused on efficient cause explanations, but two recent lines of work suggest that final and formal explanations correspond to psychologically real modes of understanding.

Kelemen *et al.* [68,69] have documented the striking phenomenon that children are 'promiscuous' in their acceptance of final causes: they claim pens are for writing, but also that mountains are for climbing. By contrast, adults accept final cause explanations selectively, typically for artifacts, artifact parts and biological traits. Lombrozo and Carey [29] have shown that this selectivity results from the restriction of final cause explanations to cases for which the function invoked in the explanation had a causal role in bringing about what is being explained: artifacts generally have specific properties *because of* their functions, whereas entities like mountains have properties irrespective of consequences, such as permitting climbing. Causal beliefs constrain the acceptability of final causes, but the sense of understanding that accompanies such explanations is insensitive to the details of the causal mechanisms involved. Insensitivity to mechanistic details is reflected in the fact that radically different processes – such as intentional design and natural selection – can warrant final cause explanations.

Prasada and Dillingham [70] have pioneered the empirical study of formal explanations. They showed that people find it natural to explain some properties but not others by appeal to kind membership. For example, it is reasonable to explain why that (pointing to a carrot) is crunchy by noting that it is a carrot, but it is not reasonable to explain why that (pointing to a car) has a radio by noting that it is a car. This phenomenon holds irrespective of the fact that the same proportion of carrots are judged to be crunchy as cars are judged to have radios, suggesting that the basis for the explanation is not statistical. The phenomenon also holds for objects of all domains, suggesting that the acceptability of formal causes is not tied to specific assumptions about causal mechanisms.

Table I. Aristotle's four 'causes' or modes of explanation

Cause or mode of explanation	Description	Example
Efficient	The proximal mechanisms of change	A carpenter is an efficient cause of a bookshelf
Final	The end, function or goal	Holding books is a final cause of a bookshelf
Formal	The form or properties that make something what it is	Having shelves is a formal cause of a bookshelf
Material	The substance of which something is constituted	Wood is a material cause of a bookshelf

lightning can cause fires under certain conditions) and determines which aspects of the complex causal etiology of the explanandum are explanatorily relevant (i.e. the lightning but not the presence of oxygen or the sound of accompanying thunder). In so doing, explanations embody a great deal of prior knowledge – often more than their surface structure reflects.

Two consequences follow from the structure of explanations. First, a particular explanandum will conform to a subsuming pattern in some respects but not in others. For example, it might be relevant to explaining the forest fire that there was lightning but not that the lightning struck at 22.04 h or that the fire occurred on a Tuesday. Identifying relevant properties can, in turn, provide a principled basis for generalizing from the explanandum to novel cases. Second, when subsuming the explanandum under a general pattern, the general pattern is summoned or created from prior beliefs. Explanations thus privilege a subset of beliefs, excluding possibilities inconsistent with those beliefs. By establishing relevant properties and beliefs, explanations can serve as a source of constraint in reasoning.

The function of explanations

Explanations are accompanied by a sense of understanding [32,33] and satisfaction [16] but the drive to explain might serve a less proximal function. Social psychologists and philosophers have argued that, in revealing the past, explanations help to predict and control the future [1,6,16,34]. The work reviewed below suggests that explanations often support the broader function of guiding reasoning. In particular, engaging in explanation serves as a mechanism through which beliefs are brought to bear on the inference or judgement at hand, providing a source of constraint for underdetermined problems and a basis for generalizing from known to novel cases. This property of explanations is examined in three contexts: inferring the existence of causal relationships, deciding whether a property extends from known to novel cases and learning from text or examples.

Explanation and causal inference

Positing causal relationships from the sparse data of everyday experience requires an inferential leap, whether the inference is about an individual instance (e.g. this fire resulted from lightning, not arson) or a class of phenomena (e.g. lightning can cause fires). Traditional models of causal inference have emphasized covariation or other measures of statistical evidence [35] but a growing body of work suggests that the interpretation and impact of such evidence depends on prior beliefs [36]. In particular, an individual's ability to explain the covariation between a candidate cause and effect can determine whether covariation is taken as evidence for causation [14,37,38].

Explanations constrain causal inference by reducing the range of possibilities considered to those consistent with prior beliefs – in particular, beliefs about causal mechanisms. For example, if provided with evidence that cars of a particular color and size have better gas mileage, children and adults will disregard the confounding factor of color to conclude that car size causes the mileage difference [37]. However, if participants are provided with an explanation for how color might influence mileage (e.g. by affecting the mood of the driver), the potential causal contribution of the confounding factor is more likely to be acknowledged. Such results suggest that mechanistic explanations are the favored basis for causal inference. In fact, if asked to uncover the cause of an event, people overwhelmingly request information that sheds light on mechanisms that could explain the event, not information about factors that covary with the event [39]. When asked

Box 2. Explanation and generalization

A recent study by Rehder [48] suggests that even when participants lack the requisite knowledge to generate explanations, they readily use provided explanations as a basis for generalizing properties. By using artificial categories, Rehder was able to exert precise control over participants' knowledge, and in particular their access to explanations for the properties being generalized. In one experiment, participants judged whether a property true of a source item generalized to a target item that shared several features (high similarity) or few features (low similarity). Crucially, participants sometimes received an explanation stating that the property was caused by one of the features of the source item (Table I). When no explanation was provided, participants were more likely to generalize a property from the source to a target if the items were highly similar (Figure Ia). However, when an explanation was provided, the effect of similarity was almost completely eliminated. Instead, participants extended the property from the source to the target if the target shared the feature cited in the explanation. In a second experiment, the benefits of diverse evidence were likewise eliminated when participants generalized properties equipped with explanations but not when the properties were unexplained (Figure Ib).

Table I. Sample stimuli involving the artificial category 'Kehoe ant'. Features are indicated with variables but involved descriptions in the actual experiment

Condition	Question	Low similarity	High similarity
No explanation	A Kehoe ant with properties F_1, F_2, F_3, but not F_4, has choroidal parasites attached to its eyes. How likely is it that the following ant also has choroidal parasites attached to its eyes?	A Kehoe ant with F_1 and F_4, but not F_2 or F_3	A Kehoe ant with F_1 and F_2, but not F_3 or F_4
Explanation (cause present)	A Kehoe ant with properties F_1, F_2, F_3, but not F_4, has a venom that gives it a stinging bite. The stinging sensation is caused by F1. How likely is it that the following ant also has a venom that gives it a stinging bite?	A Kehoe ant with F_1 and F_4, but not F_2 or F_3	A Kehoe ant with F_1 and F_2, but not F_3 or F_4
Explanation (cause absent)		A Kehoe ant with F2 and F4, but not F1 or F3	A Kehoe ant with F_2 and F_3, but not F_1 or F_4

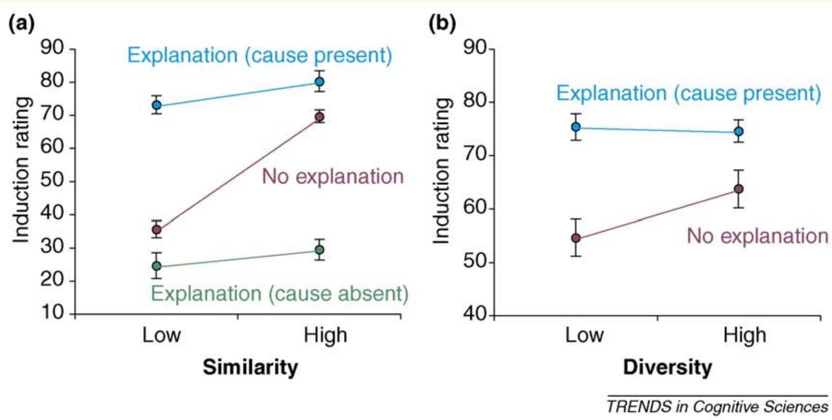

Figure I. Induction rating as a function of similarity **(a)** and diversity **(b)**. Reproduced, with permission, from Ref. [48].

to justify or argue for a claim, people likewise offer explanations over evidence [40,41], especially when evidence is sparse [42].

Generating explanations for why a claim might be true provides a way to assess the probability of that claim in light of prior beliefs. The following section demonstrates that explaining a claim can also affect belief in related claims. Whether these effects of explanation reflect a naïve epistemology [40] or a legitimate and necessary use of prior knowledge [37] is controversial. But because explanations embody prior beliefs, they have an undisputed danger: when generated from true beliefs, explanations provide an invaluable source of constraint; when generated from false beliefs, explanations can perpetuate inaccuracy.

Explanation and the generalization of properties

Suppose chocolate contains compound X. Is it more likely that the compound is also found in coffee or in tea? The extension of properties from known to novel cases has generated a substantial literature and serves as a model for developing theories of induction [43]. Among documented phenomena are the effects of similarity and diversity. Briefly, the probability that a property will be generalized from a known to a novel case is proportional to the similarity of the entities involved. If coffee is judged to be more similar to chocolate than is tea, coffee is more likely to contain compound X than is tea. Diversity manifests when generalizing to superordinate categories. A property true of lemurs and chimpanzees is judged more likely to hold for all primates than is a property true of bonobos and chimpanzees because the former provides a broader evidential base.

The effects of similarity and diversity are robust in the absence of prior knowledge about the entities and properties involved. But when participants can generate an explanation for why the property holds in given cases, the impact of similarity and diversity is eliminated or reduced (Box 2). For example, people assign a higher probability to a target statement given the truth of a base statement when the statements share an explanation, as in (A) (below), than when they involve different explanations, as in (B), irrespective of the fact that the similarity

Box 3. Explaining theory of mind

Explanation facilitates learning in a variety of contexts, including science instruction [50], mathematical problem-solving [57] and strategic game-playing [55]. Recent work by Amsterlaw and Wellman [56] demonstrated that explanation can likewise facilitate an understanding of false beliefs, a milestone in children's developing theory of mind. In variations on classic false-belief tasks, children were asked questions such as where a bear would search for an object moved from one location to another while the bear napped. To respond correctly, children must appreciate that the behavior of the bear is governed by a belief about the location of the object, in this case a false belief. On a pretest, Amsterlaw and Wellman's predominantly 3-year-old participants exhibited the typical response for their age, incorrectly predicting that the object would be sought in its true location. Over a period of several weeks, these same children participated in one of two training regimens. In an explanation condition, children received 24 false-belief problems over 12 sessions. They not only responded to the problems (e.g. predicting where the bear would look), but were also prompted to explain the correct response once it was revealed. For example, children who responded incorrectly would be asked why the bear, in fact, searched in the original location, and if they failed to provide a response were asked what the bear thought. Children in a comparison condition likewise received 24 problems and solutions but did so over fewer sessions and were prompted to explain only half of the problems. A control group participated in the pre- and post-tests.

In the post-test, children in the explanation condition significantly outperformed the comparison and control groups on false-belief tasks such as those encountered in training (Figure I). Moreover, only the explanation group succeeded on a transfer problem that had not appeared in training. The improved learning and transfer of the explanation group can be attributed to frequent explanation as the comparison group received an equivalent amount of corrective feedback but learned no more than the control group. Explanation studies in other domains find that explaining the correct response fosters learning more effectively than only receiving feedback [52]. Amsterlaw and Wellman's findings suggest that explanation not only facilitates knowledge acquisition, but can also promote conceptual change.

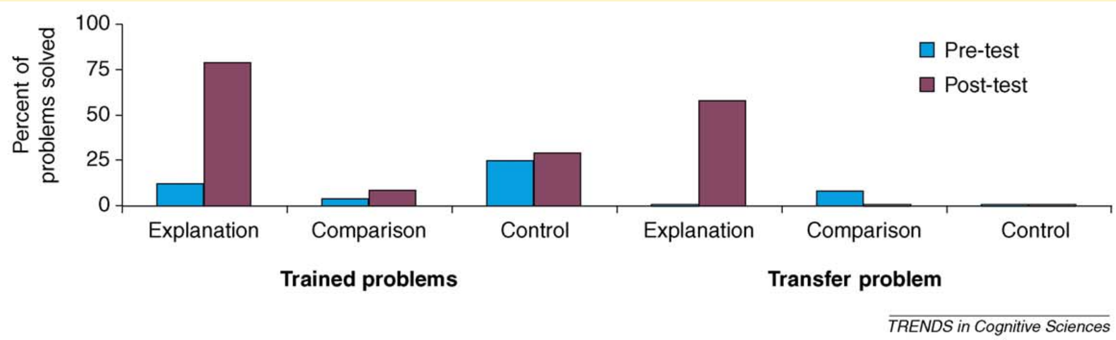

Figure I. Pre- and post-test scores as a function of condition. Reproduced, with permission, from Ref. [56].

between premise and conclusion categories is constant [44,45]:

(A) Given that furniture movers have a hard time financing a house, how likely is it that secretaries have a hard time financing a house?

(B) Given that furniture movers have bad backs, how likely is it that secretaries have bad backs?

Moreover, explanations can lead reasoners to override the influence of similarity. If told that herring and tuna have a disease, naïve participants are more likely to extend the property to wolffish, the more similar item, than to dolphins [46]. However, among fishing experts, who can generate an explanation for why the property might hold (e.g. tuna contract the disease by eating infected herring), similarity is less predictive of property extensions. Instead, properties are extended if the explanation generalizes (e.g. to dolphins, who also eat herring). Explanations can similarly override the benefits of diverse evidence [47]. Most participants judge (C) to be a stronger argument than (D), irrespective of the fact that sparrows and seeds are a more diverse sample of living things than are sparrows and dogs:

(C) Sparrows have property X; dogs have property X. Therefore, all living things have property X.

(D) Sparrows have property Y; seeds have property Y. Therefore, all living things have property Y.

Because participants can generate an explanation for why sparrows and seeds might share a property not common to all living things – namely, that sparrows eat seeds – the greater diversity fails to provide a more representative basis for generalizing to all living things. Explanations attenuate the influence of similarity and diversity by providing a more restrictive basis for generalizing from known to novel cases. Put differently, explanations isolate the sense of similarity relevant for the induction being considered, reducing reliance on global similarity [43]. Explanations are generated spontaneously when participants possess sufficient prior knowledge [44–47], and are readily used as a basis for judgement once they are made available [48]. As the next section documents, explanations can also serve as a vehicle for generalizing entire systems of knowledge, rather than isolated properties.

Learning by explaining

Explaining novel information to oneself can facilitate learning from text or examples, and foster generalization (Box 3). This is known as the 'self-explanation' effect, and has been documented in the acquisition of both procedural [49] and declarative knowledge [49,50]. Self-explanation is a more effective learning strategy than thinking out loud [51], reading study materials twice [50] or merely receiving feedback [52,53]. Although self-explaining can improve memory for procedures [54] and facts [53], the greatest and most reliable benefit is in generalization as assessed by solving transfer problems that go beyond the study materials [50,51,53,55–57]. For example, third- and fifth-grade school students who received practice on mathematical equivalence problems (e.g. '$7 + 3 + 4 = 7 + \_\_$') were more likely to succeed in solving transfer problems that involved subtraction if they were prompted to self-explain the study materials [57].

The benefits of explanation on learning are well established, but the mechanisms underlying the effect are only partially understood. Most broadly, explaining promotes learning by requiring the integration of novel information with prior beliefs, as reflected in the finding that the self-explanations most effective for learning and transfer are those that relate the information being studied to general principles [49,50,58]. In fact, 'explanations' that merely identify relevant principles improve learning [53], whereas explanations that do not relate novel information to prior beliefs are less effective [59].

The structure of explanations renders self-explaining a particularly effective strategy for learning in a way that fosters generalization. By subsuming study materials under general patterns, explanations highlight the aspects of the explanandum relevant for particular conclusions to follow or solution strategies to be effective. For example, in self-explaining the solution to a physics problem, it might become apparent that a solution strategy is appropriate by virtue of the relationship between forces (which figures in the explanation) and not the more superficial characteristics of the problem (which do not). Explaining to oneself thus facilitates generalization to transfer problems by isolating relevant senses of similarity, helping learners to overcome 'the frailties of induction' [53].

Explanation versus causal reasoning

Although explanations clearly have a role in inference and learning, one might legitimately wonder whether explanations exert an influence by virtue of something about explanations, or simply by virtue of the causal knowledge they happen to express. There are three reasons to attribute the effects of explanation to explanation *per se*. First, the structure of explanations is such that some beliefs are privileged at the expense of others. I have suggested that in summoning general patterns and isolating relevant aspects of causal factors, explanations are uniquely equipped to constrain underdetermined judgements and identify bases for generalization. Explanations thus provide an important filter on the causal beliefs brought to bear on a given inference.

A second reason to credit explanations comes from findings that prior knowledge might not be deployed through other means. Experiments that directly manipulate whether participants explain a hypothetical claim find that explaining why a claim might be true or false changes the perceived probability of that claim. For example, explaining why a psychiatric patient might commit suicide [60] or why risk takers are better firefighters [61] increases the probability assigned to that claim. Subsequently explaining the opposite attenuates the effect, suggesting that the beliefs of participants are consistent with possibilities only considered through explanation [62].

Finally, properties of explanations, such as their generality or simplicity, can influence probabilistic judgements [63–65]. Explanations that are simpler in the sense of invoking fewer causes are assigned a higher prior probability [66], and claims that can explain multiple observations are judged more likely to be believed and more valuable [67]. Because properties such as simplicity and generality apply to explanations and not to causal statements, these findings bolster the claim that the effects of explanation are not simply artifacts of causal reasoning.

Box 4. Outstanding questions

- Once an explanation is generated or provided, it influences reasoning in reliable ways. However, little is known about the mechanisms employed in the generation and evaluation of explanations. How is the content of explanations generated and evaluated? How does this ability develop?
- Most research has focused on causal explanations. Are there genuinely noncausal explanations, and, if so, are they generated and evaluated by the same mechanisms responsible for causal explanations?
- Although explanations are spontaneously generated in a variety of laboratory tasks, much less is known about the role of explanations in real-world reasoning. In particular, under what conditions do people spontaneously generate or seek explanations?
- In general, engaging in explanation leads to the preservation of existing beliefs. Can explanations lead to the rejection of prior beliefs? If so, how and when? How and when do people recognize that an existing explanation is inadequate?
- Explanations promote understanding, whether they are self-generated or received from others. Why is explanation so fundamental to understanding?

Concluding remarks

Explanations mediate a great deal of everyday reasoning. In particular, the generation and evaluation of explanations can constrain inferences by appropriately summoning prior beliefs. In evaluating claims, the existence of explanations can constitute evidence, and serve as a basis for eliminating possibilities to assess probability. In generalizing from facts or examples, explanations subsume provided information under a general pattern, thereby highlighting the senses of similarity that warrant induction.

The predominance of explanation presents a challenge for approaches to reasoning and inference that focus exclusively on decontextualized statistical evidence or similarity. When people can generate or are provided with explanations, judgements are less likely to be based only on covariation (in the case of causal attribution) or on similarity-based metrics (in the case of category-based induction). Even when dealing with complex domains of knowledge such as science or mathematics, explaining problems to oneself can foster learning and generalization more effectively than reading text or receiving feedback. Although a great deal remains to be learned (Box 4), the findings reviewed here suggest that explanation provides a unique window onto the mechanisms of learning and inference.

Acknowledgements

The author wishes to acknowledge Elizabeth Allen, Susan Carey, Tom Griffiths, Steven Sloman and Michael Weisberg for helpful comments on earlier drafts.

References

1 Heider, F. (1958) *The Psychology of Interpersonal Relations.* John Wiley & Sons, Inc

2 Gilbert, D.T. (1998) Ordinary personology. In *The Handbook of Social Psychology* (4th edn) (Gilbert, D.T. *et al.*, eds), McGraw Hill
3 Malle, B.F. (2004) *How the Mind Explains Behavior: Folk Explanations, Meaning, and Social Interaction.* MIT Press
4 Salmon, W. (1989) *Four Decades of Scientific Explanation.* University of Minnesota Press
5 Keil, F.C. and Wilson, R.A. (2000) *Explanation and Cognition.* MIT Press
6 Keil, F.C. (2006) Explanation and understanding. *Annu. Rev. Psychol.* 57, 227–254
7 Carey, S. (1985) *Conceptual Change in Childhood.* Bradford Books, MIT Press
8 Murphy, G.L. and Medin, D.L. (1985) The role of theories in conceptual coherence. *Psychol. Rev.* 92, 289–316
9 Murphy, G.L. (2002) *The Big Book of Concepts.* MIT Press
10 Keil, F.C. (2003) Folkscience: coarse interpretations of a complex reality. *Trends Cogn. Sci.* 7, 368–373
11 Murphy, G.L. and Allopenna, P.D. (1994) The locus of knowledge effects in concept learning. *J. Exp. Psychol. Learn. Mem. Cogn.* 20, 904–919
12 Ahn, W. *et al.* (2002) Effect of theory-based feature correlations on typicality judgments. *Mem. Cognit.* 30, 107–118
13 Patalano, A.L. *et al.* (2006) The importance of being coherent: category coherence, cross-classification, and reasoning. *J. Mem. Lang.* 54, 407–424
14 Koslowski, B. and Thompson, S. (2002) Theorizing is important, and collateral information constrains how well it is done. In *The Cognitive Basis of Science* (Carruthers, P. *et al.*, eds), Cambridge University Press
15 Tenenbaum, J.B. *et al.* (2006) Theory-based Bayesian models of inductive learning and reasoning. *Trends Cogn. Sci.* 10, 309–318
16 Gopnik, A. (2000) Explanation as orgasm and the drive for causal knowledge: the function, evolution, and phenomenology of the theory-formation system. In *Explanation and Cognition* (Keil, F. and Wilson, R.A., eds), pp. 299–324, MIT Press
17 Rozenblit, L.R. and Keil, F.C. (2002) The misunderstood limits of folk science: an illusion of explanatory depth. *Cogn. Sci.* 26, 521–562
18 Trout, J.D. (2002) Scientific explanation and the sense of understanding. *Philos. Sci.* 69, 212–233
19 Hilton, D.J. (1990) Conversational processes and causal explanation. *Psychol. Bull.* 107, 65–81
20 Hickling, A.K. and Wellman, H.M. (2001) The emergence of children's causal explanations and theories: evidence from everyday conversation. *Dev. Psychol.* 37, 668–683
21 Wellman, H.M. and Lagattuta, K.H. (2004) Theory of mind for learning and teaching: the nature and role of explanation. *Cogn. Dev.* 19, 479–497
22 Lewis, C. (1988) Why and how to learn why: analysis-based generalization of procedures. *Cogn. Sci.* 12, 211–256
23 Hempel, C. and Oppenheim, P. (1948) Studies in the logic of explanation. *Philos. Sci.* 15, 135–175
24 Salmon, W. (1984) *Scientific Explanation and the Causal Structure of the World.* Princeton University Press
25 Woodward, J. (2003) *Making Things Happen: A Theory of Causal Explanation.* Oxford University Press
26 Kitcher, P. (1989) Explanatory unification and the causal structure of the world. In *Scientific Explanation* (Kitcher, P. and Salmon, W., eds), pp. 410–505, University of Minnesota Press
27 Strevens, M. (2004) The causal and unification accounts of explanation unified – causally. *Nous* 38, 154–179
28 Einhorn, H.J. and Hogarth, R.M. (1986) Judging probable cause. *Psychol. Bull.* 99, 1–19
29 Lombrozo, T. and Carey, S. (2006) Functional explanation and the function of explanation. *Cognition* 99, 167–204
30 Wellman, H.M. and Liu, D. Causal reasoning as informed by the early development of explanations. In *Causal Learning: Psychology, Philosophy, and Computation* (Gopnik, A. and Schulz, L.E., eds), Oxford University Press (in press)
31 Hilton, D.J. and Erb, H. (1996) Mental models and causal explanation: judgments of probable cause and explanatory relevance. *Thinking & Reasoning* 2, 273–308
32 Dawes, R. (1999) A message from psychologists to economists: mere predictability doesn't matter like it should (without a good story appended to it). *J. Econ. Behav. Organ.* 39, 29–40
33 Ahn, W. *et al.* (2003) Understanding it makes it more normal. *Psychon. Bull. Rev.* 10, 746–752
34 Craik, K. (1943) *The Nature of Explanation.* Cambridge University Press
35 Shanks, D.R. (1995) *The Psychology of Associative Learning.* Cambridge University Press
36 Ahn, W. *et al.* Dynamic interpretations of covariation data. In *Causal Learning: Psychology, Philosophy, and Computation* (Gopnik, A. and Schulz, L.E., eds), Oxford University Press (in press)
37 Koslowski, B. (1996) *Theory and Evidence: The Development of Scientific Reasoning.* MIT Press
38 Koslowski, B. and Masnick, A. (2002) The development of causal reasoning. In *Blackwell Handbook of Child Cognitive Development* (Goswami, U., ed.), Blackwell Publishers
39 Ahn, W. *et al.* (1995) The role of covariation versus mechanism information in causal attribution. *Cognition* 54, 299–352
40 Kuhn, D. (2001) How do people know? *Psychol. Sci.* 12, 1–8
41 Glassner, A. *et al.* (2005) Pupils' evaluation and generation of evidence and explanation in argumentation. *Br. J. Educ. Psychol.* 75, 105–118
42 Brem, S.K. and Rips, L.J. (2000) Explanation and evidence in informal argument. *Cogn. Sci.* 24, 573–604
43 Heit, E. (2000) Properties of inductive reasoning. *Psychon. Bull. Rev.* 7, 569–592
44 Sloman, S.A. (1994) When explanations compete: the role of explanatory coherence on judgments of likelihood. *Cognition* 52, 1–21
45 Sloman, S.A. (1997) Explanatory coherence and the induction of properties. *Thinking & Reasoning* 3, 81–110
46 Shafto, P. and Coley, J.D. (2003) Development of categorization and reasoning in the natural world: novices to experts, naïve similarity to ecological knowledge. *J. Exp. Psychol. Learn. Mem. Cogn.* 29, 641–649
47 Medin, D.L. *et al.* (2003) A relevance theory of induction. *Psychon. Bull. Rev.* 10, 517–532
48 Rehder, B. (2006) When similarity and causality compete in category-based property generalization. *Mem. Cognit.* 34, 3–16
49 Chi, M.T.H. *et al.* (1989) Self-explanations: how students study and use examples in learning to solve problems. *Cogn. Sci.* 13, 145–182
50 Chi, M.T.H. *et al.* (1994) Eliciting self-explanations improves understanding. *Cogn. Sci.* 18, 439–477
51 Wong, R.M.F. *et al.* (2002) The effects of self-explanation training on students' problem solving in high-school mathematics. *Learning & Instruction* 12, 233–262
52 Siegler, R.S. (2002) Microgenetic studies of self-explanations. In *Microdevelopment: Transition Processes in Development and Learning* (Granott, N. and Parziale, J., eds), pp. 31–58, Cambridge University Press
53 Aleven, V. and Koedinger, K.R. (2002) An effective metacognitive strategy: learning by doing and explaining with a computer-based cognitive tutor. *Cogn. Sci.* 26, 147–179
54 Calin-Jageman, R.J. and Ratner, H.H. (2005) The role of encoding in the self-explanation effect. *Cognition & Instruction* 23, 523–543
55 Crowley, K. and Siegler, R.S. (1999) Explanation and generalization in young children's strategy learning. *Child Dev.* 70, 304–316
56 Amsterlaw, J. and Wellman, H.M. (2006) Theories of mind in transition: a microgenetic study of the development of false belief understanding. *J. Cogn. Dev.* 7, 139–172
57 Rittle-Johnson, B. (2006) Promoting transfer: the effects of direct instruction and self-explanation. *Child Dev.* 77, 1–15
58 Renkl, A. (1997) Learning from worked-out examples: a study on individual differences. *Cogn. Sci.* 21, 1–29
59 O'Reilly, T. *et al.* (1998) A comparison of self-explanation and elaborative interrogation. *Contemp. Educ. Psychol.* 23, 434–445
60 Ross, L.D. *et al.* (1977) Social explanation and social expectation: effects of real and hypothetical explanations on subjective likelihood. *J. Pers. Soc. Psychol.* 35, 817–829
61 Anderson, C.A. *et al.* (1980) Perseverance of social theories: the role of explanation in the persistence of discredited information. *J. Pers. Soc. Psychol.* 39, 1037–1049
62 Koehler, D.J. (1991) Explanation, imagination, and confidence in judgment. *Psychol. Bull.* 110, 499–519
63 Thagard, P. (1989) Explanatory coherence. *Behav. Brain Sci.* 12, 435–467
64 Pennington, N. and Hastie, R. (1993) Reasoning in explanation-based decision making. *Cognition* 49, 123–163

65 Read, S.J. and Marcus-Newhall, A. (1993) Explanatory coherence in social explanations: a parallel distributed processing account. *J. Pers. Soc. Psychol.* 65, 429–447
66 Lombrozo, T. and Rutstein, J.J. (2004) Simplicity in explanation. *Proceedings of the 26th Annual Conference of the Cognitive Science Society,* 837–842, Erlbaum
67 Preston, J. and Epley, N. (2005) Explanations versus applications: the explanatory power of valuable beliefs. *Psychol. Sci.* 16, 826–832
68 Kelemen, D. (1999) Function, goals, and intention: children's teleological reasoning about objects. *Trends Cogn. Sci.* 3, 461–468
69 Kelemen, D. and DiYanni, C. (2005) Intuitions about origins: purpose and intelligent design in children's reasoning about nature. *J. Cogn. Dev.* 6, 3–31
70 Prasada, S. and Dillingham, E.M. (2006) Principled and statistical conceptions in common sense conception. *Cognition* 99, 73–112