Published in final edited form as:

Crit Rev (N Y). 2010 January 1; 22(4): 415–426. doi:10.1080/08913811.2010.541695.

WHEN AND WHY DO HEDGEHOGS AND FOXES DIFFER?

Frank C. Keil

Abstract

Philip E. Tetlock's finding that "hedgehog" experts (those with one big theory) are worse predictors than "foxes" (those with multiple, less comprehensive theories) offers fertile ground for future research. Are experts as likely to exhibit hedgehog- or fox-like tendencies in areas that call for explanatory, diagnostic, and skill-based expertise—as they did when Tetlock called on experts to make predictions? Do particular domains of expertise curtail or encourage different styles of expertise? Can we trace these different styles to childhood? Finally, can we nudge hedgehogs to be more like foxes? Current research can only grope at the answers to these questions, but they are essential to gauging the health of expert political judgment.

In *Expert Political Judgment* (2005), Philip E. Tetlock identified two very different ways of making political predictions. Across a wide swath of political prediction tasks, he showed that in many cases "foxes"—those who don't base their predictions on one principle, but instead use several pieces of information even when they are somewhat contradictory—tend to make more accurate predictions than do "hedgehogs"—who favor one big idea and use it with gusto and usually great confidence.

This fascinating dichotomy raises several provocative questions for anyone concerned about the reliability of experts. Among these questions: Is the fox/hedgehog dichotomy also applicable to non-predictive forms of expertise, such as expertise about how something works? Is the contrast also evident in nonpolitical domains of expertise, or is it a manifestation of the properties of only some domains of expertise? Can we locate the source of the different cognitive approaches taken by foxes and hedgehogs in some aspect of cognitive development? Is the tendency to be a hedgehog inextricably linked to illusions of insight and understanding? And given that hedgehogs usually are less likely to make successful predictions than foxes, are there ways to guide them to more foxlike behaviors?

Beyond Prediction: Feedback as an Aid to Diagnostic Expertise

Experts aren't always valued for their alleged ability to make good predictions. Sometimes they are asked to diagnose a problem in a system. Other times they are asked to perform or demonstrate a skill to the untutored. Finally, they are often asked to explain a system or phenomenon, where the measure of success lies not in their predictive powers but in their

© 2010 Critical Review Foundation

Frank C. Keil, Department of Psychology, Yale University, New Haven, CT 06520, is the author, *inter alia*, of *Concepts, Kinds, and Cognitive Development* (MIT, 1989).

Publisher's Disclaimer: Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

ability to weave together all the components in a causally coherent and consistent manner (Thagard 2007). Do foxes and hedgehogs populate these non-predictive areas of expertise?

Those who are expert in diagnosing a problem, such as doctors, car mechanics, and many others, can be said to fall into two groups, similar to Tetlock's hedgehogs and foxes. It is possible to imagine, for instance, a car mechanic who always tries to shoehorn a diagnosis into one kind of causal chain, one key component, or one kind of system, such as trying to explain most cases of rough car performance in terms of poor-quality fuel, the fuel-injection system, or the valves. Such hedgehog theorizing may likely be more common among novices, who simply don't know enough about alternative potential causes of a problem, than among more experienced experts. But even an experienced expert may not be not reflective enough to realize when he is overgeneralizing, or he may tend to make snap decisions (Frederick 2005), which in turn biases him toward focusing in on the same explanation time and time again. A critical facet of most diagnostic situations, however, is that the expert can find out within a relatively short period whether his diagnosis was correct or not. In such situations the feedback is often frequent, relatively rapid, and relatively unambiguous, which may perhaps reduce hedgehog tendencies, whereas in the political arena feedback is often not immediate and is also murky enough to be prone to alternative interpretations.

There are, of course, nonpolitical diagnostic situations in which the feedback, while being clear and determinate, takes a long time to obtain and may often not be worth the effort to verify. For example, diagnosing that an aging patient is losing cognitive abilities because of Alzheimer's as opposed to other causes can take a long time to determine and may often not be seen as worth pursuing. By contrast, diagnosing whether a toothache is a result of an infected root can be linked to much quicker and more accurate feedback.

Expert Performances

It is easy to imagine hedgehogs and foxes making an appearance, too, in another non-predictive realm of expertise, that of skilled performance. One person might try to always approach a problem in the same way, whereas another might use a diverse array of strategies. One kind of expert chef, for example, might make a complex dish in exactly the same way and in the same order each time, whereas another chef might think it critical to subtly modify the recipe and even the sequence of operations as ingredients vary. More broadly, one can imagine someone who uses a bag of tricks to solve a problem as opposed to another who, to paraphrase Maslow, has only a hammer and sees everything as a nail. For example, a hedgehog might be a wizard at using duct tape to solve an enormous range of problems, from leaky boats to broken hula hoops to torn luggage. In contrast, a fox might employ a different method for fixing each of these problems. I suspect that, in the sphere of skilled expertise, the two styles of expert—the expert who is more flexible and multi-skilled, and the expert who has more rigid, fixed procedures—are analogous to, respectively, the foxes and hedgehogs that Tetlock describes. And there might a comparable outcome, namely that the foxes tend to get better results.

In a similar vein, one can imagine two very different ways an expert might demonstrate to a novice how to do a task. A hedgehog expert might stress always approaching the task in the same way, while a fox may stress being flexible and adjusting the technique one uses according to context or as the task evolves over time. The critical question is whether hedgehogs would show the same biases in the area of skilled performance as they do in political predictions.

Expert Explanations

Experts are often called upon to provide explanations of how something works, or why it is the way it is. In contrast to diagnosing a problem, the expert may be asked to focus on the system in its normal recurring state, such as explaining the causes of the tides, how a television works, or why most sports teams win more games at home. In such cases there may be no prediction.

Whether or not one finds a dichotomy between foxes and hedgehogs here may well depend on the kinds of explanations involved. Explanations have a wide range of forms and functions (Keil 2006), and only some types of explanation may tend to encourage the fox/hedgehog divide. For example, where there is only one plausible mechanism in a given domain, such that expertise consists of cumulative and detailed knowledge about that mechanism (such as how the tides work), there may not be any fox/hedgehog differences among different experts. However, where the phenomenon to be explained—such as why sports teams win more games at home—has several plausible explanations, differences between foxes and hedgehogs may flourish. This may especially be true when experts are called upon to explain unique individuals or events, such as why a particular sports team lost a championship series, as opposed to explaining whole systems. These types of explanation, in which the data are open to numerous interpretations, seem similar to the situation facing experts in the political arena.

Different Types of Prediction

Just as there are different types of expertise, not all of which involve prediction, so, too, are there different types of prediction.

Consider for example, weather forecasting. One can imagine an expert who focuses heavily on one factor, such as barometric pressure or typical weather at that time, location, and date. In contrast, another expert might try to integrate as many different sources of information as possible. Conceivably, a hedgehog strategy that focused on the wrong variable would be pruned out by negative feedback (Nickerson 1998), leaving only successful hedgehogs with little need for any foxes. On the other hand, in many forms of health care that do not lack for negative feedback, foxes might thrive due to the large number of biological differences among patients, leading to many different interpretations of which factors are important. Similarly, those who predict horse races, stock prices, or football-games might be pushed toward fox-like styles of expertise. In all of these cases, unlike the political case, it be more difficult to massage the feedback and convince oneself that one was partly right (although sports fans often talk about near misses as a sign of their success).

If so, then do situations that produce clear and dichotomous outcomes diminish the presence of foxes and hedgehogs, as compared to situations in which a wide variety of potential factors might all be plausibly used in the prediction? Similarly, to what extent does the size and presence of fox/hedgehog differences depend upon the frequency of, as well as the length of the interval between, the expert's prediction and the feedback that confirms or disproves it? Perhaps in domains where the frequency of tests is high and the feedback interval is low, the contrast between foxes and hedgehogs will play a minor role.

It is worth emphasizing that domains vary not only in the frequency of feedback and the interval of time before receiving it, but also in terms of the number of possible variables that might cause the outcomes being predicted—with potential variables ranging from a small set of discrete elements, such as four to six, to hundreds of less-well defined ones. However, it is not always true that a larger number of factors makes predictions harder. For example, in attempting to predict whether a particular bridge will collapse under the load of a twenty-ton

truck, an expert might consider a small set of variables: the truss design, the fastener specifications, and the materials used. Solving this problem might still be very complex because of the ways variables can interact. The classical 3 and n-body problems in physics is an even more powerful example. Here, trying to predict where just three (let alone more) bodies in motion will be after a short interval can make for a computationally intractable problem, even when the only factors influencing movement are the masses of the bodies, their initial velocities, and the gravitational forces between them.

Separate from the *difficulty* of making accurate predictions is the question of which types of predictive situation might cause experts to use hedgehog mode more often—regardless of whether it is wise or not: complex situations, even with a small number of variables; or situations that are complex *because of* the sheer number of possible factors. Examples of the latter type range from reasons for the stock market rising or falling to reasons for a company's success or failure to predictions about an individual's performance in a new job. I am inclined to think that the sheer number of factors may matter more in determining whether an expert will typically incline toward being a fox or a hedgehog, but that is just a hunch.

Other aspects of a given domain might also increase or diminish the relative prevalence of foxes and hedgehogs. These might include, first, the relevance of various components of a domain in predicting or explaining a phenomenon. In some domains, all experts will agree that all components are clearly relevant even if they differ on exactly how they all interact with each other, whereas in other domains the controversy is about which factors matter at all. Next, domains may differ in whether certain factors of that domain are dichotomous or graded in effects. Another aspect to consider is whether the factors in a given domain are important independently of one another, or whether they heavily interact with each other in determining a given outcome. Finally, some domains may have clearly discrete subcomponents, each with their own rich causal processes, while other domains have simple unitary causes. In this last case it may well be that the richer and more complex a subcomponent is in its own right, the easier it is for a hedgehog to latch onto it without being clearly contradicted. These and other possible differences among domains raise intriguing research questions about the various ways in which some domains might spawn hedgehogs while others would strongly discourage such tendencies. Similarly, they raise the question of whether aspects of various domains allow the same person to be a strong hedgehog in one domain but constrain him from being a hedgehog in a different domain.

How Foxes May Turn into Hedgehogs

Are hedgehogs and foxes born or bred? That is, do hedgehog and fox tendencies emerge as cognitive styles that are apparent in young children, or do they only show up in adulthood? Apart from intra-individual variations, is there a general developmental tendency towards either a more hedgehoggy or foxy style of thinking in younger children?

The answers are not clear, in part because they depend on how these two styles of thought might coincide with other cognitive differences.

Consider the provocative idea that young children might be more fox-like than hedgehog-like. There is considerable evidence that young children can be more holistic than analytic (Werner 1940), and that, at least when categorizing and naming objects and phenomena, they will often assign equal weight to a large number of variables rather than zeroing in on just a few features that they see as defining (Keil 1989). The shift from hedging one's bets by embracing all typically occurring features to a focus on just a few features that are seen as defining a category represents the way that children naturally progress from relatively little knowledge to expertise. Moreover, this shift is not acute but gradual, and occurs at

different points in development as a function of particular domains. Would young children therefore be more likely to holistically assess a large number of features to make predictions or offer explanations?

We also know that young children can be more tolerant than adults of contradictions and inconsistencies and have more difficulty seeing circularities (Osherson and Markman 1975; Baum, Danovitch, and Keil 2008). If a child hasn't thought enough about a given domain to have a well-worked-out explanatory or predictive theory, will she simply weigh everything she knows holistically? In addition, young children have often been observed to treat different dimensions of a problem as integral, not separable (Smith and Kemler 1978). So it could be argued that they would have more difficulty teasing out one single facet of a problem as crucial.

Children's holistic and integral forms of thinking do not necessarily mean that children are always fox-like. In some domains, the simplicity of focusing on a single feature might be so preferable to having to keep in mind several dimensions at once that it wins out. We know, for example, that when children attempt to predict when a balance beam with different weights and different distances from the fulcrum will balance, that they may first focus on a single dimension such as weight or height while ignoring other ones, and only later bring in other dimensions (Siegler 1976).

Children are also prone to be more overconfident than adults about what they know, or think they know (Mills and Keil 2004). They overestimate their explanatory and predictive abilities as well as their abilities to do simple things, such as recalling a set of items (Schneider 1999). Their overconfidence may arise from an inability to reflect metacognitively on a problem and what the solution to it entails (Schneider 2008), as well as from a general overoptimism about their abilities (Lockhart, Chang, and Story 2002).

Of course, most children aren't experts on much of anything, and it may be that expertise itself results in the fox-hedgehog split. This hypothesis can be tested, however, because in certain restricted domains, children can become quite expert, such as having surprisingly sophisticated knowledge of dinosaurs, their habitats, and their diets (Chi and Koeske 1983). In such areas of expertise, are young children, compared to adult experts, nonetheless more hedgehog or fox-like?

Studies on the developmental emergence of the split between foxes and hedgehogs would be difficult to arrange because of the problems of children's limited expertise, but such studies might be invaluable to understanding the foundational biases that set the ground for adults' biases. One possible research question is whether children have an increasing need for closure on problems, and, if so, whether this need leads to hedgehog tendencies over time. One can further ask if, and to what degree, early learning experiences influence the development of children's problem-solving strategies.

Illusions of Explanation and Understanding

All of us fall sway to illusions of explanatory depth and of understanding. Illusions of explanatory depth (IOED) occur when people think they can explain the workings of the world in far more detail than they really can (Rozenblit and Keil 2002).

We are more likely to overestimate our explanatory abilities than our abilities to perform procedures or to recall simple facts and narratives. The realm of explanations is particularly conducive to overconfidence for several reasons. First, most people are rarely asked to explain things in detail, so they don't have much experience with the challenges involved. Second, when explaining how things work, we often have the object or situation in front of

us, giving us real-time explanatory cues that we conflate with more explanatory knowledge than we actually possess, independent of the situation. Finally, it can be harder to have one's explanation disproven than to be corrected about a factual mistake or to fail to execute a procedure properly.

Perhaps being a hedgehog is made possible by the IOED: Hedgehogs may be more influenced than foxes by the seduction of explanations. But it would be useful to know the extent to which hedgehogs *do* labor under a large IOED. In addition to being more confident about their predictions than are foxes, is the disconnect between what they really know and what they think they know larger than a fox's disconnect? Does a hedgehog think that his big idea has great depth and texture when in fact, if probed, his big idea is shown to be little more than a simple, potentially unsupportable rule? Is there a relationship between the size of the IOED and a tendency to think one has a straightforward big idea that can explain everything about the phenomenon of interest?

Another illusion related to explanations is the illusion of insight (Weisberg, Keil, Goodstein, Rawson, and Gray 2008). Sometimes a concrete but irrelevant feature of an explanation can make one think one has understood a phenomenon much better than one really has. Thus, in one series of studies Weisberg et al. (ibid.) showed that while college students could easily distinguish empty, circular explanations of psychological phenomena from non-circular ones, they largely lost that ability when irrelevant neuro-imaging data were added. In that case, they regarded both the circular and non-circular explanations as insightful. It seemed that the mere presence of irrelevant neuro-imaging findings gave people the illusion of insight. This illusion may arise from a feeling that an explanation involving a simple, concrete piece of information (e.g., an image of a physical region of the brain) is somehow much more insightful than more abstract, relational psychological information. Perhaps hedgehogs are especially vulnerable to such features of explanations, even in domains where they have no expertise. It would be helpful to examine whether hedgehogs are more likely than foxes to seize on irrelevant but simple, clear, and concrete bits of information in explaining phenomena outside their area of expertise. Future researchers could conduct a series of studies analyzing what sort of information people see as most insightful in hearing novel explanations, and then seeing if they can map their findings onto Tetlock's.

Hedgehog Hunting

Although neither foxes nor hedgehogs do very well in Tetlock's studies, there is still a clear overall advantage to being a fox, as a fox is more likely to produce accurate predictions and to know the likelihood of being wrong. The question naturally arises as to what it would take to turn hedgehogs into foxes.

There are clearly many external incentives for hedgehogs to exist. They are better masters of the sound bites that the media craves, since it is much easier for a hedgehog to neatly encapsulate an explanation with one big memorable idea. Hedgehogs also radiate more confidence in their opinions, making them more attractive commentators than those who emphasize their uncertainty or insist on caveats. Despite all these pro-hedgehog pressures, there may still be ways to shift some hedgehogs to more fox-like views.

One technique may simply be to increase metacognitive awareness of the task at hand and to encourage people to think about whether they are incorporating all the relevant information. Standing back and asking why one is excluding large swaths of information might make one rethink the wisdom of one's big idea. Running through a checklist that fostered such metacognitions might gently nudge people away from hedgehog styles before they press the last key to send their opinion out into the world. It is striking how successful simple checklists can be (Gawande 2010).

A second intervention might be to ask hedgehogs to offer more and more details supporting their big idea. In doing so, they may quickly realize that their big idea often is more sound bite than substance. In studies on the IOED, people become more modest about what they really know when forced to supply detailed explanations (Rozenblit and Keil 2002). One might also expose hedgehogs to experts with sharply contradictory views in a domain where the hedgehogs have no great intellectual investment. If both groups of experts are equally successful in their predictions, the tie may suggest to hedgehogs that other views must somehow be accommodated.

It might seem that there are certain areas where hedgehogs have great virtues. For example, consider the determined scientist who, in a drive for parsimony and elegance, breaks through a complicated system of many small factors to find a simple, elegant explanation that pivots on one key concept. The switch from Ptolemy's epicycles to Kepler's and Newton's much simpler models of planetary motion is an example. Similarly the fundamental dogma of molecular biology—that information flows in one direction, from nucleic acids to proteins—could be seen as a hedgehog-like triumph in the biological sciences, where a massively cluttered set of accounts was swept away by one elegant idea.

It is not clear, however, how often such big ideas really do work in the sciences as opposed to collections of less clear and partially contradictory hunches (Dunbar 1995, 365–96). Even some of the biggest ideas in the history of science, such as the central dogma of molecular biology, have turned out to be quite a bit more complicated and messy than they first appeared (e.g. Barry 2007). Thus, even in the natural sciences, hedgehog-like tendencies may end up distorting insights as much as clarifying them.

In the end, the kinds of experts we choose to rely on and defer to may depend on us as much as they depend on the alleged experts. All cultures depend profoundly on divisions of labor and on the divisions of expertise associated with those labors (Durkheim [1893] 1997). To benefit from such expertise, members of a society should not only know which experts to approach for which problems, but should also to know how to choose among experts when they disagree. Tetlock's work suggests that confident, adamant experts should be regarded with some caution and that those who are more tentative and equivocal might well be more valuable. To the extent that the public is more informed about the real complexity of problems, the more it might resist explanations that act as if everything can be explained in one way. To the extent that the public reads not just the sources that confirm their own biases, but a wider range of opinions, it is likely to encounter equally successful but contradictory positions and start to realize how elements of both might be at work. When laypeople engage in such practices, they might actually start to shape the behaviors of experts in ways that guide them to more considered and nuanced opinions.

Acknowledgments

This paper, and some of the research it describes, was prepared with the support of NIH Grant R-37-HD023922.

REFERENCES

- Barry, Patrick. Genome 2.0: Mountains of New Data are Challenging Old Views. *Science News*. 2007; 172(10):154.
- Baum, Laura A.; Danovitch, Judith H.; Keil, Frank C. Children's Sensitivity to Circular Explanations. *Journal of Experimental Child Psychology*. 2008; 100:146–155. [PubMed: 18078950]
- Chi MTH, Koeske RD. Network Representation of Child's Dinosaur Knowledge. *Developmental Psychology*. 1983; 19:29–39.

- Dunbar, K. How Scientists Formally Reason: Scientific Reasoning in Real-World Laboratories. In: Sternberg, Robert J.; Davidson, Janet E., editors. *The Nature of Insight*. Cambridge, Mass: MIT Press; 1995.
- Durkheim, Emile. *The Division of Labor in Society*. Simpson, George, translator. New York: The Free Press; [1893] 1997.
- Frederick, Shane. Cognitive Reflection and Decision Making. *Journal of Economic Perspectives*. 2005; 19:25–42.
- Gawande, Atul. *The CHECKLIST Manifesto*. London: Profile Books; 2010.
- Keil, Frank C. Concepts, Kinds, and Cognitive Development. Cambridge, Mass.: MIT Press; 1989.
- Keil, Frank C. Explanation and Understanding. *Annual Review of Psychology*. 2006; 57:227–254.
- Lockhart, Kristi L.; Chang, Bernard; Story, Tyler. Young Children's Beliefs about the Stability of Traits: Protective Optimism? *Child Development*. 2002; 73:1408–1430. [PubMed: 12361309]
- Mills, Candice M.; Keil, Frank C. Knowing the Limits of One's Understanding: The Development of an Awareness of an Illusion of Explanatory Depth. *Journal of Experimental Child Psychology*. 2004; 87:1–32. [PubMed: 14698687]
- Nickerson, Raymond S. Confirmation Bias: A Ubiquitous Phenomenon in Many Guises. *Review of General Psychology*. 1998; 2:175–220.
- Osherson, Daniel; Markman, Ellen. Language and the Ability to Evaluate Contradictions and Tautologies. *Cognition*. 1975; 3:213–226.
- Rozenblit, Leonid R.; Keil, Frank C. The Misunderstood Limits of Folk Science: An Illusion of Explanatory Depth. *Cognitive Science*. 2002; 26:521–562. [PubMed: 21442007]
- Siegler, Robert S. Three Aspects of Cognitive Development. *Cognitive Psychology*. 1976; 8:481–520.
- Smith, Linda B.; Kemler, Deborah G. Levels of Experienced Dimensionality in Children and Adults. *Cognitive Psychology*. 1978; 10:502–532. [PubMed: 699515]
- Schneider, Wolfgang. The Development of Metamemory in Children. In: Gopher, Daniel; Koriati, Asher, editors. *Attention and Performance, vol. XVII. Cognitive Regulation of Performance: Interaction of Theory and Application*. Cambridge, Mass.: MIT Press; 1999.
- Schneider, Wolfgang. The Development of Metacognitive Knowledge in Children and Adolescents: Major Trends and Implications for Education. *Mind, Brain and Education*. 2008; 2:114–121.
- Thagard, Paul. Coherence, Truth, and the Development of Scientific Knowledge. *Philosophy of Science*. 2007; 74:28–47.
- Tetlock, Philip E. *Expert Political Judgment: How Good Is It? How Can We Know?*. Princeton: Princeton University Press; 2005.
- Weisberg, Deena Skolnick; Keil, Frank C.; Goodstein, Joshua; Rawson, Elizabeth; Gray, Jeremy R. The Seductive Allure of Neuroscience Explanations. *Journal of Cognitive Neuroscience*. 2008; 20:470–477. [PubMed: 18004955]
- Werner, Heinz. *Comparative Psychology of Mental Development*. New York: International Universities Press, Inc.; 1940.