Scientific Explanation and the Sense of Understanding*

J. D. Trout†‡
Philosophy Department and the Parmly Hearing Institute
Loyola University Chicago

Scientists and laypeople alike use the sense of understanding that an explanation conveys as a cue to good or correct explanation. Although the occurrence of this sense or feeling of understanding is neither necessary nor sufficient for good explanation, it does drive judgments of the plausibility and, ultimately, the acceptability, of an explanation. This paper presents evidence that the sense of understanding is in part the routine consequence of two well-documented biases in cognitive psychology: overconfidence and hindsight. In light of the prevalence of counterfeit understanding in the history of science, I argue that many forms of cognitive achievement do not involve a sense of understanding, and that only the truth or accuracy of an explanation make the sense of understanding a valid cue to genuine understanding.

1. Introduction. Few products of intellectual life are more exhilarating, more pleasing to give and receive, than a good explanation. While *theories* of explanation can be quite technical, and the content of particular *scientific* explanations quite arcane, the cue for acceptable explanation remains cheerfully informal. A good explanation "feels right." In all cases, the cue seems the same: A familiar sense of understanding that the explanatory story, causal or otherwise, delivers to us.[1] It is not news, of

*Received October 2001; revised February 2002.

†Send requests for reprints to the author, Philosophy Department and the Parmly Hearing Institute, Loyola University Chicago, 6525 North Sheridan Road, Chicago, IL 60626; e-mail: jtrout@orion.it.luc.edu.

‡Thanks are due to Paul Abela, Michael Bishop, Richard Farr, Robert Klee, and Joseph Mendola for very helpful comments on this paper. I especially want to thank Dwayne Mulder for early conversations and arguments with me about the nature and function of explanation.

1. I believe that models play a central role in the most promising objectivist account of explanation. Models stand as durable repositories of explanatory information. They

Philosophy of Science, 69 (June 2002) pp. 212–233. 0031-8248/2002/6902-0003$10.00

course, that this sense of understanding alone is not necessarily a reliable guide to truth, nor is it a necessary condition for good explanation. Still less is it sufficient for good explanation. As one philosopher of explanation puts this observation: "It is no explanation to provide a distorted representation of the world, and the 'understanding' induced by such incorrect models is illusory at best" (Humphreys 1989, 103).

What causes us to accept an explanation? Is it the sense of a coherent story it conveys? Undeniably, there is a special kind of intellectual satisfaction—an affective component—that occasions the acceptance of an explanation, a sense that we have achieved understanding of the phenomena. Peirce identifies the distinctive cognitive experience of explanatory understanding by isolating the moment of final acceptance; the good explanation "is turned back and forth like a key in a lock" (1908, 100). This description alone should supply little solace to those holding that good explanations are epistemically reliable. After all, alchemists surely felt the key turn, but once inside we find only false descriptions of causal mechanisms. And when Galen arrived at a diagnosis of melancholy due to black bile, his sense of understanding was so gratifying it must have balanced his humors. Ptolemy claimed that "if the earth did not lie in the middle [of the universe], the whole order of things which we observe in the increase and decrease in the length of daylight would be fundamentally upset" (1984, 42/H20).

Some of our favorite stories in the history of science, such as Kekule's famous "Eureka" episode, feature dramatic journeys to compelling explanations.[2] In these scenarios, the explanation "felt right." This sense of understanding is causally potent. It can be so comforting that explanation stops when this feeling is experienced. I will argue that the psychological sense of understanding is just a kind of confidence, abetted by hindsight, of intellectual satisfaction that a question has been adequately answered.[3]

are not fleeting impressions in the scientist's mind; they remain long after the often fanciful occasions of their production. Because we can actually represent their parts, we have a permanent record of how our theoretical postulates fare under test. But I will not defend this claim here.

2. One night, the 19th century chemist August Kekule, labored over the structure of compounds until he

". . . fell into a reverie, and lo, the atoms were gamboling before my eyes. Whenever, hitherto, these diminutive beings had appeared to me, they had always been in motion; but up to that time, I had never been able to discern the nature of their motion. Now, however, I saw how, frequently, two smaller atoms united to form a pair; how a larger one embraced two smaller ones; how still larger ones kept hold of three or even four of the smaller; whilst the whole kept whirling in a giddy dance. I saw how the larger ones formed a chain. . . . I spent part of the night putting on paper at least sketches of these dream forms." (Quoted in Rothenberg 1993, 291).

3. This sense of understanding is most likely at the bottom of many standards used as

Thus this sense of satisfaction is confidence that one enjoys an accurate description of the underlying causal factors sufficient (under the circumstances) to bring about the phenomenon we are examining. But confidence is, notoriously, not an indicator of truth.

The track record of this sense of understanding is not entirely treacherous, however. This sense sometimes seems to be associated with very secure explanations as well. Darwin claimed that the explanation of the distribution and origin of species in terms of natural section was "so satisfactory" (Darwin 1859, 476). Avogadro drew a similar inference to the best explanation for the existence of molecules. Diverse domains and measurement procedures led to the unification of molecular and atomic phenomena. Perrin commented on the confidence with which the molecular hypothesis could be held, and did so in distinctly psychological terms:

> "Our wonder is aroused at the very remarkable agreement found between values derived from the consideration of such widely different phenomena. Seeing that not only is the same magnitude obtained by each method when the conditions under which it is applied are varied as much as possible, but that the numbers thus established also agree among themselves, without discrepancy, for all methods employed, the real existence of the molecule is given a probability bordering on certainty." (1913, 215–216)

This subjective sense of understanding may be conveyed by a psychological impression that the explanatory mechanisms are transparent and coherent, or that the explanation seems plausible, and so should be confidently accepted. At the moment, there is neither a satisfying formal account of explanation, nor agreement about the important informal criteria for good explanation, producing what one review casts as "an embarrassment for the philosophy of science" (Newton-Smith, 2000, 132).

Current theories of explanation may leave us embarrassed, but we should be at least as embarrassed by our "feels-right" diagnostic standard for the acceptance of an explanation. Sadly, it has been the fate of humans

---

signals of good explanations, such as simplicity (consider Watson's claim about his and Crick's completed DNA model that "[A] structure this pretty just had to exist" (1968, 205)), or rendering the unfamiliar, familiar. These standards, however, are complicated. What counts as simple or familiar is theory-dependent, and not surprisingly, verdicts in particular cases are controversial. For example, the familiarity account may seem regularly violated by such cases as quantum mechanics, in which it appears that the familiar is being explained in terms of the less familiar. But this appearance may be misleading. A finding unfamiliar to the lay public may convey a feeling to the scientist that the pieces of a theoretical puzzle have just fallen into place. Only a comprehensive overview can treat all of these issues. The reader is referred to Salmon (1992, 14), where the reduction of the unfamiliar to the familiar is explicitly discussed.

to have more counterfeit than genuine experiences of "understanding". Upon reflection, we are aware of this unreliability. What, then, beckons us to make such judgments? There is, in fact, a very substantial literature in cognitive psychology establishing an overconfidence effect for scientists and laypeople alike (see Fischhoff, Slovic, and Lichtenstein 1977), an effect that conveys the mere "feeling of knowing" (see Koriat 1993, 1995). This overconfidence is not the mere result of individual differences in personality or of clinical delusions of grandeur; it is the normal consequence of routine cognitive activity.

What conception of understanding is important to theories of explanation? The traditional account casts understanding as the result of a process that is fundamentally reconstructive: understanding requires that the individual be able to piece together bits of information in their cognitive possession. Reconstructive accounts of explanation agree on one point: Understanding is centrally involved in explanation, whether as an intellectual goal or as a means of unifying practice. Peter Achinstein asserts a "fundamental relationship between explanation and understanding" (1983, 16). David Lewis requires of an explanation that "the recipient understands and believes what he is told" (1993, 185). Wesley Salmon proposes that scientific understanding is achieved in two ways: by "fitting phenomena into a comprehensive scientific world-picture" (1998, 77), and by detailing and thereby exposing the "inner mechanisms" of a process (1998, 77). Michael Friedman claims that the relation of phenomena that "gives understanding of the explained phenomenon" is "the central problem of scientific explanation" (1988, 189). Peter Railton associates the scientific understanding of quantum mechanics with "the ability to fill out arbitrarily bits of ideal explanatory text" (Railton 1993, 170–171), which produces a mass of theoretical detail about the object of explanation. Philip Kitcher relates understanding and explanation so closely that elucidation of this connection in a theory of explanation "should show us *how* scientific explanation advances our understanding" (1988, 168). James Woodward claims that a theory of explanation should "identify the structural features of such explanation which function so as to produce understanding in the ordinary user" (1993, 249).

While these declarations associate explanation with understanding, none of these accounts have much to say about the precise *nature* of understanding, and certainly none of them entail that a sense of understanding is *criterial* for good explanation. But most of them do draw upon our ordinary, ungarnished notion of understanding, in terms of "detailing" and "identifying" the inner mechanisms and structural features of interest. And this contingent fact is a significant feature in contemporary accounts of explanation. Terms like "detailing", "identifying" and "fitting" imply that explaining a phenomenon is a metacognitive process, not an implicit or otherwise inarticulate process or skill. Accordingly, explanation is portrayed as detective work, in which the investigator consciously pieces together otherwise disparate facts into a coherent, global picture. But scientists are not little epistemologists and, apart from the lay appetite for scientists' grand Nobellizing on the "big" issues, working scientists' explanatory concerns are local rather than global. Given the philosophical tendency to focus on the explicit detailing of nomic subsumption, establishing fit, and describing mechanisms, Ronald Giere's characterization of philosophical work on explanation nicely captures the current situation: "[M]ost philosophical writing on 'scientific explanation' is not really about explanation *within* science, but about the use of scientific knowledge in the explanation of events in everyday life" (1988, 104–105).[4]

2. The Traditional Epistemic Analysis of Explanatory Understanding. If the research findings on judgment and decision-making examined below are accepted, then the sense of understanding introduces a systematic, but regrettably inaccurate, index of intellectual achievement. In particular, this sense of understanding results from hindsight and overconfidence biases, as well as mistaken attachment to the idea that transparency is routinely achievable. These biases diminish the truth-tracking role of explanation. It is not just philosophical theories of explanation that have accorded to the sense of understanding an essential role in explanation. Psychological theories of explanation, too, appeal to the important role of a sense of understanding in both everyday and scientific explanation. Like some global, unifying accounts of explanation in the philosophy of science, a prominent psychological account focuses on the unified conceptual framework it provides: ". . . [I]n everyday use an explanation is an account that provides a conceptual framework for a phenomenon (e.g., fact, law, theory) that leads to a feeling of understanding in the reader-hearer" (Brewer et al. 1998, 120). And scientific explanations are no different in this respect; they should "provide a feeling of understanding" (1998, 121), in addition to whatever objective virtues the explanation might possess. The operative assumption may be that this sense of understanding is an epistemic virtue; the fact that an explanation conveys this sense of understanding offers a reason for thinking it is also a true, or at least a good, explanation. How-

4. Beyond these observations, I will not argue for the claim that the contemporary focus on global features of explanation invites reliance on a subjective sense of "fit" or "overall coherence" as a cue to acceptable explanation—a cue that explanation can stop. Because this feature is not explicitly defended in any particular theory of explanation, this paper does not treat any specific account as the target of my critique. Instead, I will assume that the above observations establish that the internalist notion of understanding is typically at play in many of the influential contemporary theories of explanation. Positions asserting the contrary must match the evidence presented here.

ever, in order to be an epistemic virtue, it should be at least positively correlated with accurate causal descriptions, or the goodness of an explanation. But it isn't; or so I will argue. Instead, this sense is a consequence of largely nonepistemic forces, such as the demonstrated psychological tendency of overconfidence, or the desire to reduce, as William James put it, "that peculiar feeling of inward unrest known as indecision" ([1890] 1981, Chapter 26 (Will), 1136).[5] Formulating a unified, consistent story is one way to eradicate that peculiar feeling of inward unrest. An explanation pleasantly discharges that feeling of intellectual unease.[6] But unity and consistency with background knowledge is a poor substitute for accuracy and truth, as the litany of false but unifying and consistent theories in the history of science should warn. A good story, and so a good explanation, is coherent. Coherence appears to consist of three features: completeness, plausibility and consistency. One might suppose that a plausibility requirement addresses the worry about false theories. However, judgments of plausibility are themselves theory-dependent, and so if your theory is badly false you cannot expect your explanations to fare much better.[7]

Most of the widely discussed accounts of explanation have been objectivist: What makes an explanation good concerns a property that it has independent of the psychology of the explainers; it concerns features of external objects, independent of particular minds. These properties might be formal or they might be causal. As unassailable as these objectivist credentials might be, even the most trenchantly objectivist philosophers of science are tempted by the allure of internal access. Hempel treated explanations as formal arguments, weakening the formal criteria in response to counterexamples as seemed fit. Hempel's Deductive-Nomological (D-N) model of explanation attempted to capture this subjective component in the psychological experience of expectability: "[G]iven particular circumstances and the laws in question, the occurrence of the phenomenon *was to be expected;* and it is in this sense that the explanation enables us to *understand why* the phenomenon occurred" (1965, 337). Of course, the crucial phrase "understand why" must be glossed in a nomic way here.[8]

5. Some of the most ill-fated enterprises were mounted and sustained by thinkers with supreme confidence in their explanations. Consider the alchemist Paracelsus, who claimed to have found the "Universal Medicine" in his *Archidoxis: Comprised in Ten Books,* on the grounds that "By means thereof I have cured the leprosy, venereal disease, dropsy, the falling sickness, colic, scab, and similar afflictions; also lupus, cancer, noli-me-tangere, fistulas, and the whole race of internal diseases, more surely than one could believe." (Paracelsus 1665 Chapter VII)

6. Gopnik (1998, 121) likens to orgasm the satisfaction conveyed by an explanation.

7. Thagard (1989) provides a deep and detailed account of coherence.

8. Against the covering-law model, Scriven (1962) emphasized that we often embrace particular statements as explanations even when no statement about the phenomenon

Wesley Salmon, too, expressed the centrality of understanding to explanation, even in the context of his treatment of explanation as "causal-mechanical" and "ontic": "To understand the world and what goes on in it, we must expose its inner workings. To the extent that causal mechanisms operate, they explain how the world works" (Salmon 1984, 133). Presumably 'understand' here is cast in terms of internal access; 'exposing' the world's inner workings normally requires the ability to represent implementations of (in this case) mechanisms. That is, what makes the knowledge appropriately expository is that significant parts of its complexity can be represented, and what affords understanding is that *we* can represent it. This latter feature is made possible by subjective access.

In more recent work, Salmon became more explicitly permissive about the requirements of explanatory understanding, and located the intellectual value of scientific explanations in their power to achieve a number of different goals, all of which

> "enhance our understanding of the world. Our understanding is enhanced (1) when we obtain knowledge of the hidden mechanisms, causal or other, that produce the phenomena we seek to explain, (2) when our knowledge of the world is so organized that we can comprehend what we know under a smaller number of assumptions than previously, and (3) when we supply missing bits of descriptive knowledge that answer why-questions and remove us from the particular sorts of intellectual predicaments. Which of these is *the* function of scientific explanation? None *uniquely* qualifies. . . ." (1989, 134–135)

In his wonderfully accessible "The Importance of Scientific Understanding", Wesley Salmon (1998) distinguished between what he called "scientific" and "psychological" senses of understanding. The scientific sense of understanding

> "involves the development of a world-picture, including knowledge of the basic mechanisms according to which it operates, that is based on objective evidence—one that we have good reason to suppose actually represents, more or less accurately, the way the world is." (1998, 90)

The objectivity of explanation is undermined, however, if the "good reason to suppose" condition is interpreted in internalist terms.

Fitting these mechanisms into a scientific "world-picture" is an admirable aim, and it animated the philosophes of the Enlightenment. But nascent specialization later dashed those hopes, making the theoretical

---

to be explained was derivable from it. Toulmin (1961) is another classic alternative to the formal treatment of the received view of explanation that dominated at the time.

arcaneness of scientific knowledge at once a powerful source of theoretical progress and an equally potent source of suspicion among those wanting scientific matters to be decidable by any person who was intelligent and generally educated. The admirable aim was not achieved, and the science that Enlightenment figures so loved left them behind.

Salmon's discussion tempts many questions, mostly urgent requests for further elaboration of the character of this global understanding. Because there is significant and honest disagreement among scientists about what the correct scientific world-picture is, it is not clear what constitutes fitting phenomena into the scientific world-picture. What Salmon might intend here is the attempted integration of the best-tested and influential views of the recent history of science: quantum mechanics, atomic theory, the germ theory of disease, Darwinian evolution, and so on. Taken together, these theories form a truly impressive monument to the scope, detail, and accuracy of modern science; perhaps they constitute, as Salmon put it, a scientific world picture. But if explanatory scientific understanding requires seeing "how we can fit them [phenomena] into the general scheme of things, that is, into the *scientific world-picture*" (1998, 87), then most people are incapable of explanatory scientific understanding, including most scientists. Indeed, when scientists piece together phenomena, they do so by focussing on the detailed findings of their (usually) narrow specialization. In contemporary science, global unification arises spontaneously from coordinated piecemeal efforts, not from a meta-level at which the philosopher or reflective scientist assembles remote domains (Miller, 1987). In fact, in light of the arcaneness of contemporary theoretical knowledge, no single individual is so situated.[9] Does Salmon's talk of global explanation *entail* this internalist conception of understanding? Of course not. Natural definitions seldom have such entailments. But these accounts of explanation are the work of philosophers, not psychologists. We should expect philosophical accounts of understanding to begin, and perhaps end, with the primitive internalist account of subjective sense.

But this convenient account of justification should not be mistaken for an accurate description of scientific practice. When a scientist (or anyone else, for that matter) ultimately accepts an explanation, it is more likely that the scientist, without any such detailed internal representation of a

9. We might argue for a kind of "community internalism", according to which explanatory justification terminates in the transparent understanding of an appropriate person—as testimonial knowledge might terminate in the first-hand observer—deferred scientific explanatory knowledge terminates in the transparent representation of the expert in the arcane field. I will leave unanswered the plausibility of this "community internalism". The nature of deferred knowledge is a deep and interesting one, however, and has been addressed in a variety of ways by Hardwig (1985, 1991), Humphreys (2000), Kitcher (1990, 1993), and Trout (1992, 1998).

global (and so diverse) subject matter, simply defers to outside experts or appeals to textbook knowledge.[10] But when philosophers piece together a global explanation, they look to the subjective sense of fit as a cue for an acceptable explanation.

Michael Friedman's account of explanation, too, is tied to the goal of understanding, and it stipulates special criteria for increased understanding:

> "I claim that this is the crucial property of scientific theories we are looking for; this is the essence of scientific explanation—science increases our understanding of the world by reducing the total number of independent phenomena that we have to accept as ultimate or given. A world with fewer independent phenomena is, other things equal, more comprehensible than one with more." (Friedman 1988, 195)

The connection that Friedman draws between understanding and number of independent phenomena concerns cognitive efficiency and tractability. If phenomenon P is reduced to Q, then Q is more basic, and so more likely itself to be irreducible. If what makes a phenomenon irreducible is that it can't be explained in terms of anything else more basic, then it is inexplicable. A theory (or explanation) that has fewer inexplicable commitments would seem to be rationally preferable to one that has more. As a psychological hypothesis, it may be easier to comprehend a theory that has just one, rather than two or more, irreducible mysteries.

Friedman proposes that the kind of understanding provided by scientific knowledge is global rather than local. Accordingly, accounts of explanation that identify the objects of understanding as individual phenomena can be expected to fail: "Scientific explanations do not confer intelligibility on individual phenomena by showing them to be somehow natural, necessary, familiar, or inevitable" (Friedman 1988, 197). It is for this reason that scientific understanding reduces the number of ultimate (and thus inexplicable) commitments.[11]

Philip Kitcher advances a novel and detailed version of the unification approach, adding crucial categories of learning to this account. Kitcher's approach is crafted to honor the work-a-day details of scientific practice. On Kitcher's account, comprehensibility does not emerge spontaneously from the mere reduction of independent theoretical commitments; comprehensibility depends upon the possession of background knowledge and

10. There are those, such as Wittgensteinians and some pragmatists, who are suspicious of any explanatory account of behavior that adverts (ineliminably) to internal representations. We might note that contemporary psychology treats this as a settled issue, decidedly in favor of the causal role of internal representations. In empirical matters of the mind, I will cast my lot with psychology.

11. Peter Lipton points out that the understanding associated with explanation is simply additional knowledge, not some kind of "superknowledge" (1993, 207).

the training required for the assimilation of superficially different phenomena to common categories:

> “Understanding the phenomena is not simply a matter of reducing the “fundamental incomprehensibilities” but of seeing connections, common patterns, in what initially appeared to be different situations. Here the switch in conception from premise-conclusion pairs to derivations proves vital. *Science advances understanding of nature by showing us how to derive descriptions of many phenomena, using the same patterns of derivation again and again, and, in demonstrating this, it teaches us how to reduce the number of types of facts we have to accept as ultimate (or brute).* [fn deleted]” (1989, 432)

Kitcher captures this objective, nonfoundationalist and pedagogical feature of explanation by asserting that knowledge of a theory

> “involves the internalization of the argument patterns associated with it, and that, in consequence, an adequate philosophical reconstruction of a scientific theory requires us to identify a set of argument patterns as one component of the theory. This is especially obvious when the theory under reconstruction is not associated with any “grand equations” and when reconstructions of it along traditional lines produce a trivialization that is remote from the practice of science.” (1989, 438)

The identification and internalization of argument patterns is an effect of learning, of scientific pedagogy. Of course, familiarity with the esoteric argument patterns that are characteristic of a theory can be acquired only as an arcane skill. I will have more to say about the implicit learning of esoteric detail in section 3.

Unlike other descriptions of unification, for Kitcher the role of explanation issues from dirtying our hands with local details of scientific practice, rather than from philosophical rumination or the supposed scientific aesthetic preference for grand coherence. Explanation plays an epistemic role in science not chiefly because global scientific worldviews set an explicit, top-down research agenda, but because working scientists must piece together local findings within the very small handful of models and theories they use and evaluate.

Once reconstituted as an openly causal account, this treatment of explanation has sufficient detail and descriptive accuracy to accommodate the history of science and current practice. Any such argument patterns—if they are to represent the esoteric practice of science—must be pedagogically valuable and highly enthymematic. Also, there is no reason that the internalization of argument patterns cannot be understood in a more straightforward metaphysical way, in terms of the causal features of internal representations of the explainer and the relevant parts of the world being represented.

3. Understanding, Learning, and Transparent Capture of Information. 'Understanding' has many senses. Consider a sense of understanding associated with an explanation of why jets fly. According to one sense of understanding, I understand why planes fly if I know Bernoulli's principle. In such a case, I must have some knowledge of relevant background conditions, of course—how much the plane weighs, the area of the wings, etc. But in a stronger sense of understanding, I can know Bernoulli's principle and still not understand why planes fly, because my knowledge that Bernoulli's principle applies in this case is not associated with a specific kind of subjective state or feeling; it does not have the phenomenology of understanding.

Understanding is usually the consequence of a learning process, generally construed. But there are many modes of learning. We can learn by feedback, with explicit instruction or not. We can learn by analogy, representing to oneself or "internalizing" a model of a process, and then generating various scenarios about the performance of the system under a range of circumstances. Models, of the sort described above, provide a kind of metaphor by which one can transfer a well-understood structure to one less well-understood. Indeed, learning is typically characterized as the "transfer of structure", where this means either the transfer of informational structure from the environment to an individual (as when one learns first-hand that bees sting), from one individual to others (as when a teacher explains the Pythagorean Theorem to students), or from one psychological subsystem to another (as when one uses their visual information about facial expression to correctly interpret an auditory sequence of spoken language).

We often, perhaps routinely, learn without awareness. In fact, we acquire enormous amounts of information quite incidentally about unattended dimensions of objects. This point can be pressed even further: Surprisingly little learning occurs via the learner transparently representing to itself the as-yet unintellected object. In the last 40 years or so, cognitive psychology has produced a spate of research establishing the importance of implicit learning in a variety of cognitive domains. Consider the research on implicit learning of particular voices. You begin by exposing individuals over several days to a series of words uttered by a number of people. You then ask them to distinguish between words, half of which were, and half of which were not, among the original test set. In this case, reaction times are faster, and error rates lower, for words that were uttered by the same people as those in the original test set. Thus, the same words, uttered by voices different from those in the original test set, required more perceptual and cognitive effort to identify. The redundancy gains achieved by memory for concrete vocal detail were the product of implicit learning. Throughout the training sessions, the participants were not aware that they were learning vocal detail of the particular word-tokens; they thought they were simply learning word-types, as the instructions asked them to do (see Nygaard et. al. 1994, Goldinger 1996, and Church and Schacter 1994). In effect, subjects developed a tacit model of each talker's vocal characteristics, and their cognitive mechanisms engaged that model in matching the input to a prior lexical entry. Because their model was accurate, they enjoyed increased speed and accuracy. Implicit learning has been demonstrated on a variety of perceptual and cognitive tasks, such as lexical decision, picture naming, object decision, word association, category instance generation, and answering general knowledge questions (Baddeley 1996). The lesson here is the same as that for the implicit learning of particular voices: Awareness of the dimensions attended to is not required for learning. This sensitivity to relevant dimensions—a recognitional capacity—is certainly one form of understanding.

If we are trying to sort a collection of instances into natural classes or "kinds" of objects—a taxonomic task required by any science—we must rely on feedback about negative instances in order to improve our discriminative sensitivity. In other words, we must see some non-gazelles if we are to sort objects into classes of gazelles and ibis. There is one major exception to this generalization, and this exception provides the basis of the pedagogical account of explanation I favor. We can learn without exposure to negative instances only when, to quote a famous study on medical diagnostic classification, we have a theoretical model that gives us "access to the major causal influences, possesses accurate measuring instruments to assess them, and uses a well-corroborated theory to make the transition from the theory to fact (that is, when the expert has access to a specific model)" (Dawes, Faust, and Meehl 1989, 1670). In short, for a combination of accuracy and efficiency, there is no explanatory substitute for an accurate model, or a good theory. Let us now turn to the cognitive limits that make the contingent adoption of an accurate model so important.

4. The Sense of Understanding as the Product of Hindsight and Overconfidence Biases. If evidence from the history of science shows that sense of understanding is not produced by a reliable relation between belief and truth, where does this subjective sense of understanding derive from? I will argue that in a significant number of cases it comes from two well-documented psychological biases—hindsight and overconfidence—biases that are difficult to correct, and which survive different experimental methods, test items, and classes of people.[12] The history of science is a rich

12. The present paper focuses exclusively on the influence of psychological biases on

source of examples. In a classic moment of hindsight and overconfidence, Ptolemy claimed that "it is idle to seek for causes for the motion of objects towards the centre, once it has been so clearly established from the actual phenomena that the earth occupies the middle place in the universe, and that all heavy objects are carried towards the earth" (1984, 43/H22). Ptolemy's conservative appeal to coherence with (in this case, false) background beliefs, together with his dismissive treatment of alternatives, displays his influence by hindsight. But what is the evidence that there is a hindsight *bias?*

4.1. Hindsight Bias. Explanation accounts for events after the fact. People are notably unaware of the influence that outcome information has on them. This is precisely the retrodictive epistemic position of the explainer. The traditional manner of establishing the hindsight bias begins by asking subjects to estimate the likelihood of various outcomes of an upcoming event, and then retesting them after the event, asking them to recall how likely they had found each of the possible outcomes the first time around. Fischhoff and Beyth (1975) did just that in an early study of the hindsight bias. Prior to President Nixon's trip to China and the Soviet Union in 1972, subjects were asked how likely they found a variety of possible outcomes (for example, whether Nixon would meet Mao, that the Soviet Union and US would establish a joint space program, etc.). Two weeks to six months after the trip the subjects were asked to fill out the same questionnaire. They were asked to recall the probabilities they assigned initially to the same events and, if they couldn't recall, to assign the probability they would have assigned immediately before Nixon's trip. They were also asked if each of the listed outcomes had, in fact, occurred.

The results were a striking demonstration of the distorting influence of hindsight. For those outcomes that subjects thought had occurred, they remembered their estimates as more accurate than they in fact were. For those outcomes thought not to have occurred, subjects recalled their estimates as having been lower than they in fact were. The effect seems to strengthen with the passage of time. After three to six months, 84 percent of the subjects displayed hindsight biases. Therefore, after learning the results of Nixon's trip, subjects believed the outcomes were more predictable than they actually were.

Learning the outcomes of scientific theory testing—the preamble to any explanation—places scientists in a position similar to the hindsight subjects. Explanation is retrospective. When constructing an explanation, we draw one line through many events, a line that could have been drawn in

explanation specifically. Solomon (1992) discusses empirical findings concerning rationality in theory evaluation.

different ways, even if we know the point of termination. As we draw this line, we are not very accurate judges of how much we are affected by information about an outcome, be it a mass extinction, an explosion, or infection. But once the line is drawn, we conceptualize the event as inevitable, and thus people tend to say that the event was fairly predictable all along. Thus, the hindsight bias is also known as the "I-knew-it-all-along effect." In particular, people tend to overestimate how probable they thought the event was before it occurred.

The hindsight bias is all-too familiar. Many of us have had the experience of telling a colleague over lunch about an experimental finding about human behavior. The reported finding is immediately met with a knowing chuckle: "You didn't need to run an experiment to know that." Your interlocutor, insensitive to the effect that the reported outcome has just had on him or her, claims to have "known it all along." Of course, your interlocutor had no such knowledge. Most psychological processes and social behavior are complex, and you don't gain knowledge of them by mere reflection or casual observation; that is why experiments are needed.[13]

It is the hindsight bias that lies behind our tendency to confuse predictability with statistical contingency. Prediction is often thought to lift the epistemic burden from explanation. And while predictive accuracy can be an important index of scientific integrity, prediction is both epistemically over-rated and often difficult to secure. For instance, knowing the factors contributing to an effect in a complex system does not necessarily allow you to predict the outcome. Plane crashes are especially good examples of the limited fare supplied by prediction. Upon analysis of cockpit transcripts, a handful of important factors are identified that contributed to the crash. But as Robyn Dawes puts the point: "[T]hese factors would not allow us to predict future crashes very well at all. Airplane crews are often fatigued; bad weather occurs frequently; miscommunication is not that unusual, nor are temporary breakdowns of radio communication or panic at the last minute" (1999, 37). These effects are too high-frequency to satisfy a low-probability condition for prediction. But more impor-

13. One conversational strategy would clearly illustrate the falsity of the contention that they knew it all along, while at the same time demonstrate the pervasiveness of the hindsight bias. It goes as follows. Tell your interlocutor about another experimental finding, but lie about the punchline. For example, tell them that Darley and Latane (1968) found that the more people are present when someone needs help, the more likely an individual is to provide help. This is false, and when your interlocutor provides the "obvious" explanation for this counterfeit effect, you can reveal the real result, thereby correcting what is no doubt the inadvertent arrogance of your interlocutor. As luck would have it, this conversational strategy will probably never be widely implemented, so grotesquely does it violate norms of social politeness.

tantly, the hindsight bias causes us to acquiesce in the belief that we have an understanding of an effect. After that, we regard our search as complete. This unjustified conceit is what makes the hindsight bias so damaging to the search for accurate explanations. And the hindsight bias is not alone in the aid and comfort of indolent arrogance.[14] The overconfidence bias feeds the lazy beast as well.

4.2 Overconfidence Bias. If error is the constant companion of inquiry, so is overconfidence. Lay adults are systematically prone to believing that they are right when they are not. The literature demonstrating overconfidence is very large, and counterexamples to the effect have been difficult to produce. To cite one representative example of lay overconfidence, Fischhoff, Slovic and Lichtenstein (1977) asked subjects to indicate the most frequent cause of death in the U.S., and to estimate their confidence that their choice was correct (in terms of "odds"). It turns out that, when subjects set the odds of their answer's correctness at 100:1, they were correct only 73% of the time. Remarkably, even when they were so certain as to set the odds between 10,000:1 and 1,000,000:1, they were correct only between 85% and 90% of the time. It is important to note that the overconfidence effect is systematic and directional. It is highly replicable, and survives changes in task and setting. And, the effect is in the direction of over rather than underconfidence. Given this, it is not surprising that expert training is not the key to deliverance from the overconfidence bias. Physicists, economists, and demographers have all been observed to suffer from this bias, even when reasoning about the content of their special discipline (Henrion and Fischoff 1986). It is of little consolation, then, that good explanation has been associated with the "feeling" or "sense" of understanding. The same sense of understanding, of intellectual conviction, accompanied the above subjects' incorrect answers. And there is no reason to think that answers to *explanatory* why-questions are different in a way that allows them to evade that fact.

In fact, the "settled" subjective feeling of understanding that is associated with overconfidence (and so confidence) may be the subjective state that prompts a stopping rule, or a more informal decision that we can stop explaining or considering alternative explanations of an event, on the grounds that we now understand the relevant causes. Overconfidence, then, may be a truly disastrous component in explanatory reasoning. Indeed, our judgments of accuracy are systematically correct only when we have a good theory or model of the process we are making judgments about. This tendency toward overconfidence would not be so damaging if

14. For a more sanguine interpretation of the hindsight bias that concedes its occurrence but attempts to limit its significance, see Gigerenzer et al. (1999).

our judgments were "calibrated", a term used to designate the extent to which confidence matches accuracy. But in science we seldom have such accuracy information available to us.[15]

Often, we hold a primitive view, one that is independent of or prior to any existing, empirically supported theory, especially when we are innocent of the arcane causes of our behavior. In this situation, the most important determinant of the plausibility of the explanation is the accuracy of the causal theory used to generate it. I take a similar approach to explanation. Accurate explanation appears to be a truth-conducive factor in science; we seem to have scientific knowledge that is explanatory, and this knowledge helps us to generate further knowledge. Unless we rest content with an account of scientific explanation resulting from pure philosophical analysis, we need a scientifically respectable account of scientific explanation itself.

People who believe p, and others who believe $\sim$p, can both believe overconfidently. True, the priority assigned to this sense of understanding is responsible for much error in the history of science. But there is a systematic and directional effect of this error: Overconfidence extends the life of a false belief, and propagates a true belief. After all, Copernicus claims, of two alternatives to heliocentrism, that "the mind shudders at either of these suppositions" (1952, 514). This is merely an interesting personal report, even though heliocentrism is an accurate theory. The important question is whether heliocentrism is true, not whether envisioning an alternative is too intellectually painful to bear.[16]

Why do people find it so difficult to learn from the lessons of overconfidence? The answer has two parts. First, it is difficult to learn from individual cases when we typically don't get (or don't retain) systematic feedback about the quality of our judgments. To mention just one restriction on feedback, we can't compare the long-term outcomes of our actual decisions against alternative decisions we didn't implement. Second, the overconfidence is general and persistent, and this attachment to subjective

15. There are studies, such as Arkes, Christensen, Lai, and Blumer (1987) that demonstrate that overconfidence can be eliminated in a decision-maker when persistently exposed to a rigorous schedule of specially prepared feedback. However, there is little resemblance between this setting and those of routine explanation, and scientists (or anyone else for that matter) are not subjected to any such exercises in calibration. So this finding offers no hope to the trenchant defender of the centrality of the sense of understanding.

16. The reader can decide whether Kepler was lacking in confidence: "If you forgive me, I shall rejoice; if you are enraged with me, I shall bear it. See, I cast the die, and I write the book. Whether it is to be read by the people of the present or of the future makes no difference: let it await its reader for a hundred years, if God Himself has stood ready for six thousand years for one to study him." (Kepler [1619] 1997, 391)

evaluation of complex incoming evidence has specific consequences for learning. Perhaps this is best seen in the "interview effect", a very prominent example of this overconfidence in our powers of subjective evaluation. When "experts", such as hiring and admissions officers, are able to review applicants in unstructured interviews, they are still outperformed by statistical prediction rules that take no account of the interviews. In fact, unstructured interviews actually degrade the reliability of human prediction (Bloom and Brundage 1947; DeVaul et al. 1957; Oskamp 1965; Milstein et al. 1981). That is, people degrade the reliability of their predictions by availing themselves of unstructured interviews; the interview information used is irrelevant to (and thus dilutive of) accurate prediction about future performance. Although the interview effect is one of the most robust findings in psychology, highly educated people ignore its obvious practical implications. This occurs in part due to our overconfidence in our subjective ability to "read" people, but is not limited to it. We suppose that our insight into human nature or into the subtleties of nature is so powerful that we can plumb the depths of not only an individual, but a domain, or control the processes used to evaluate it.

Our little conceits are abetted by a classic and systematic frailty in interpreting probabilistic information. This conceit accounts for the general finding that, in the face of a half century of experiments showing the inferiority of human judgment to well-tested statistical prediction rules in such contexts, many disciplinary experts and others still base judgments on subjective impressions and unmonitored evaluation of the evidence. Resistance to these findings runs very deep, and typically comes in the form of a self-serving bias we might call Peirce's Problem.[17] Peirce ([1878] 1982, 281–2) raised what is now the classic worry about frequentist interpretations of probability: How can a probability claim (say, the claim that 99 out of 100 cards are red) be relevant to a judgment about a particular case (whether the next card will be red)? After all, the next card will be red or not, and the other 99 cards can't change that fact. Those who resist the findings of inferior performance are typically quite willing to admit that *in the long run,* a simple statistical prediction rule will be right more often than human experts. In short, their (over)confidence in subjective powers of reflection leads them to deny that we should believe we are unreliable *in some particular case.* Whatever the long-run performance of humans, so the sentiment goes, in this case I am the right person to make the subjective judgment. And the effect of this lack of discipline is made worse because it is difficult to secure all of the information necessary for clear feedback: We can't compare the long-term outcomes of our actual decisions against the decisions we would have made if we hadn't interviewed the candidates but simply used a rule.

17. See Bishop and Trout (2002).

Our reluctance to envision our fallibility, and then admit error, goes a long way toward accounting for our unwillingness to acknowledge the embarrassments of overconfidence and hindsight biases. On the basis of this evidence, humans appear naturally disposed to exaggerate the powers of our subjective faculties. No matter what this disposition's origin, it won't do to reply to my claims of unreliability that this result seems implausible, for that is exactly the point. For example, the theory predicts that people will find the sad lesson of the interview effect implausible. Despite a moral that should be humbling, we continue to look for, and when it is present we use, that special, subjective sense of illumination as a cue of correctness. By now it should be clear that this cue, by itself, is unreliable, and often isn't present when we have explanatory knowledge. So clearly if explanation imparts understanding, understanding must consist in something other than a subjective sense of metacognitive control.

5. Conclusion. Despite these intellectual warts on the phenomenology of explanatory understanding, many philosophers of science, such as scientific realists, hold that explanation plays a robust, epistemic role in theory construction. But if the traditional, phenomenological sense of understanding is not a valid cue of a good explanation, what accounts for the epistemic function of explanation in theory construction? I have not attempted to answer this question here. While the conception of good explanation I would defend is objectivist, it is virtually unique in its utter disregard for the role of the subjective sense of understanding as a determinant of explanatory goodness.

On a genuinely realist account, understanding is a cluster of epistemic virtues. How could there be intellectual virtue at work without transparent appreciation of the content of one's theoretical beliefs, content that characterizes the causes of the phenomenon to be explained? The parallel with perceptual knowing is instructive here. A nonskeptic would regard the faculty of vision as a justification-conferring faculty even though the successful perceiver may know nothing about visual transduction, or about the vast array of other physiological causes of visual knowledge. It would seem that knowledge does not require conscious access to information that mediates input and behavior. Why, then, should explanation be essentially tied to a subjective sense of understanding? Nonformal approaches to explanation may employ causal, or model-theoretic notions. But they also use sense of understanding as a cue of those specific notions' plausibility. This is a problem faced by all philosophers of science who suppose that the feeling of understanding attending the intellectual assembling of mechanisms constitutes a reason for thinking our account is correct.

The fact is, our history is littered with inaccurate explanations we confidently thought were obviously true: the explanation for mental illness in terms of demonic possession, the humoral theory of illness, and so on. The sense of understanding would be epistemically idle phenomenology were it not so poisonous a combination of seduction and unreliability. It actually does harm, sometimes making us squeamish about accepting true claims that we don't personally understand, and more often operating in the opposite direction, causing us to overconfidently accept false claims because they have a kind of anecdotal or theoretical charm.

Moreover, the understanding conveyed by a good explanation may be a community achievement. Except for the simplest of events, explanatory understanding is not essentially an achievement of an individual. And any alternative account of explanation that requires the transmission of a sense of understanding must address this criticism. My positive account of scientific explanation asserts that, as a contingent matter of fact, the only feature of an explanation that can render explanation epistemic is its systematic tendency to produce increasingly accurate theories. In effect, only explanations capable of sustaining theoretical progress are good explanations. Such explanations are usually approximately true, rather than merely better than available alternatives, as the latter may simply be the best of a rotten lot. In order to accord explanation the epistemic role it seems to play in successful theory selection in contemporary science, we must abandon our sentimental attachment to the comforting sense of understanding, or, at least, abandon the idea that this sense is a valid cue of truth. This will not be easy. Explanation is a backward-looking affair and, as we have seen, 30 years of research on judgment documents that people are not good at tracking how they are affected by knowledge of outcomes. As Robyn Dawes puts it, "The problem is that there is a many-many relationship between antecedents and consequences in the course of human life. As we retrospect, in contrast, we can create many-one relationships." (1999, 37) Explanation creates many-one relationships.

The preceding assault on the role of a subjective sense of understanding in explanation leaves untouched many admirable features of more orthodox accounts of explanation. Nomic subsumption, unification, and logical derivation all may advance scientific aims. In the same sense that casting problems in terms of an independently well supported theory (e.g., materialism, germ theory of disease, etc.) can help to solve them correctly, an explanation that is accurate—no matter how nontransparent its components—is better than an inaccurate one that merely makes us feel that we understand. Only the former is related to the world in a way that conduces to the success of science. If our subjective sense of understanding seems to be reliably related to truth only when we have a good theory, one must wonder whether that subjective sense of understanding is doing any epistemic work at all. It appears that, when it comes to explanation, there is no substitute for simply being (approximately) right. A consistent scientific realist is driven to this historically contingent, radically objectivist account of explanation. What would this scientific realist account of explanation look like? Given our proclivity to confidently concoct a story when we don't have a well-confirmed account ready-to-hand,[18] saying less is truly more. So, for the time being, I will take the cue.

REFERENCES

Achinstein, Peter (1983), *The Nature of Explanation.* New York: Oxford University Press.

Arkes, Hal, C. Christensen, C. Lai, and C. Blumer (1987), "Two Methods of Reducing Overconfidence", *Organizational Behavior and Human Decision Processes* 39: 133–144.

Baddeley, Alan (1996), *Human Memory: Theory and Practice,* 2nd edition. Boston: Allyn and Bacon.

Bishop, Michael, and J. D. Trout (2002), "50 Years of Successful Predictive Modeling Should be Enough: Lessons for the Philosophy of Science", *Philosophy of Science* 68 (Proceedings), forthcoming.

Bloom, R. F. and E. G. Brundage (1947), "Predictions of Success in Elementary School for Enlisted Personnel", in D. B. Stuit (ed.), *Personnel Research and Test Development in the Naval Bureau of Personnel.* Princeton University Press; Princeton, 233–261.

Brewer, William, C. Chinn, and A. Samarapungavan (1998), "Explanation in Scientists and Children", *Minds and Machines* 8: 119–136.

Church, Barbara, and D. Schacter (1994), "Perceptual Specificity of Auditory Priming: Implicit Memory for Voice Intonation and Fundamental Frequency", *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20: 521–533.

Copernicus, Nicholas. (1952), *Revolutions of Heavenly Spheres,* Book I, Section 4, in M. Adler (ed.), *Great Books of the Western World,* volume 16 (translated by Charles Glenn Wallis). Chicago: William Benton Publishers, 513–514.

Darley, John and B. Latane (1968), "Bystander Intervention in Emergencies: Diffusion of Responsibility", *Journal of Personality and Social Psychology* 8: 377–383.

Darwin, Charles (1859), *The Origin of Species,* 6th ed., reprinted in 1962. New York: Collier.

Dawes, Robyn (1999), "A Message from Psychologists to Economists: Mere Predictability Doesn't Matter Like It Should (Without a Good Story Appended to It)", *Journal of Economic Behavior & Organization* 39, 29–40.

——— (2001), *Everyday Irrationality.* Boulder CO: Westview Press.

Dawes, Robyn, D. Faust, and P. Meehl (1989), "Clinical Versus Actuarial Judgment", *Science* 243: 1668–1674.

DeVaul, R. A., F. Jervey, J. A. Chappell, P. Carver, B. Short, and S. O'Keefe (1957), "Medical School Performance of Initially Rejected Students", *Journal of the American Medical Association* 257: 47–51.

Fischhoff, Baruch, P. Slovic, and S. Lichtenstein (1977), "Knowing with Certainty: The Appropriateness of Extreme Confidence", *Journal of Experimental Psychology: Human Perception and Performance* 3: 552–564.

Fischhoff, Baruch and R. Beyth (1975), " 'I Knew It Would Happen': Remembered Probabilities of Once-Future Things", *Organizational Behavior and Human Performance* 13: 1–16.

Friedman, Michael (1988), "Explanation and Scientific Understanding", in J. C. Pitt (ed.) *Theories of Explanation,* New York: Oxford University Press, 188–198. Originally published in *Journal of Philosophy* 71: 5–19.

18. The desire to discharge the sense of unease occasioned by the unexplained is so great that we often fabricate causal stories just to fill the vacuum. This is so even when the unexplained feature concerns our own internal cognitive workings. See Nisbett and Wilson (1977); and Dawes (2001), especially Chapter 7, "Good Stories".

Giere, Ronald (1988), *Explaining Science: A Cognitive Approach.* Chicago: University of Chicago Press.

Gigerenzer, Gerd, P. Todd, and the ABC Group (1999), *Simple Heuristics that Make Us Smart.* New York: Oxford University Press.

Goldinger, Stephen (1996), "Words and Voices: Episodic Traces in Spoken Word Identification and Recognition Memory", *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22 (5): 1166–1183.

Gopnik, Allison (1998), "Explanation as Orgasm", *Minds and Machines* 8: 101–118.

Hardwig, John (1985), "Epistemic Dependence", *Journal of Philosophy* 82: 335–349.

——— (1991), "The Role of Trust in Knowledge", *Journal of Philosophy* 88: 693–708.

Hempel, Carl (1965), *Aspects of Scientific Explanation.* New York: The Free Press.

Henrion, M., and B. Fischoff (1986), "Assessing Uncertainty in Physical Constants", *American Journal of Physics,* 54: 791–798.

Humphreys, Paul (1989). *Chances of Explanation.* Princeton, NJ: Princeton University Press.

——— (2000), "Explanation vs. Understanding". In J. Fetzer, (ed.), *Science, Explanation, and Rationality: The Philosophy of Carl G. Hempel.* New York: Oxford University Press, 267–286.

James, William ([1890] 1981). *Principles of Psychology,* vol 2. Cambridge: Harvard University Press, Chapter 26 (Will).

Kepler, Johannes ([1619] 1997), *The Harmony of the World,* trans. by E. J. Aiton, A.M. Duncan, and J. V. Field. Philadelphia: American Philosophical Society.

Kitcher, Philip (1988), "Explanatory Unification", in J. C. Pitt (ed.), *Theories of Explanation,* New York: Oxford University Press, 167–187. Originally published in *Philosophy of Science* 48: 507–531.

——— (1989), "Explanatory Unification and the Causal Structure of the World", In Philip Kitcher and Wesley Salmon (eds.), *Scientific Explanation: Minnesota Studies in the Philosophy of Science,* vol. 13. Minneapolis: University of Minnesota Press, 410–505.

——— (1990), "The Division of Cognitive Labor", *Journal of Philosophy* 87: 5–21.

——— (1993), *The Advancement of Science.* New York: Oxford University Press.

Koriat, Asher (1993), "How Do We Know that We Know? The Accessibility Model of the Feeling of Knowing", *Psychological Review* 100(4): 609–639.

——— (1995), "Dissociating Knowing and the Feeling of Knowing: Further Evidence for the Accessibility Model", *Journal of Experimental Psychology: General* 124(3): 311–333.

Lewis, David (1993), "Causal Explanation", in David–Hillel Ruben (ed.) *Explanation,* New York: Oxford University Press, 182–206. Originally published in *In Philosophical Papers, II.* New York: Oxford University Press, pp. 214–240.

Lipton, Peter (1993), "Contrastive Explanation", in David-Hillel Ruben (ed.), *Explanation,* Oxford: Oxford University Press, 207–227. Originally published in D. Knowles (ed.) *Explanation and Its Limits.* Cambridge University Press, 1990.

Miller, Richard W. (1987), *Fact and Method.* Princeton: Princeton University Press.

Milstein, R. M., L. Wildkinson, G. N. Burrow, and W. Kessen (1981), "Admission Decisions and Performance During Medical School", *Journal of Medical Education* 56: 77–82.

Newton-Smith, William H. (2000), "Explanation", in W. H. Newton-Smith (ed.), *A Companion to the Philosophy of Science.* Malden, MA: Blackwell, 127–133.

Nisbett, Richard and T. Wilson (1977), "Telling More Than We Can Know: Verbal Reports on Mental Processes", *Psychological Review* 84(3): 231–259.

Nygaard, Lynne C., M. S. Sommers, and D. B. Pisoni (1994), "Speech Perception as a Talker-Contingent Process", *Psychological Science* 5(1): 42–46.

Oskamp, S. (1965), "Overconfidence in Case Study Judgments", *Journal of Consulting Psychology* 63, 81–97.

Paracelsus, Theophrastus (1665), *Of the Tincture of the Philosophers,* trans. by J. H. Oxon. London: J. H. Oxon. Chapter VII: Concerning the Renovation of Men.

Peirce, Charles S. ([1878] 1982), "Doctrine of Chances", in Charles S. Peirce, *Writings of Charles Sanders Peirce: A Chronological Edition,* vol. 3. Bloomington: Indiana University Press, 276–289.

——— (1908), "A Neglected Argument for the Reality of God", *The Hibbert Journal* 7: 90–112.

Perrin, Jean (1913), *Les Atomes.* Paris: Alcan.
Ptolemy, Claudius. (1984), *Almagest,* trans. by G. J. Toomer. New York: Springer-Verlag.
Railton, Peter (1993), "Probability, Explanation, and Information", in David–Hillel Ruben (ed.) *Explanation,* New York: Oxford University Press, 160–181. Originally published in *Synthese* 48: 233–256.
Rothenberg, Albert (1993), "Creative Homospatial and Janusian Processes in Kekule's Discovery of the Structure of the Benzene Molecule", in John H. Wotiz (ed.) *The Kekule Riddle,* Clearwater, FL: Cache River Press, 286–309. (Translation of Kekule quotations by Francis R. Japp).
Salmon, Wesley (1992), "Scientific Explanation", in Merrilee Salmon et. al, *Introduction to the Philosophy of Science.* Englewood Cliffs, NJ: Prentice-Hall, 7–41.
——— (1984), *Scientific Explanation and the Causal Structure of the World.* Princeton: Princeton University Press.
——— (1989), "Four Decades of Explanation", in Philip Kitcher and Wesley Salmon, (eds.), *Scientific Explanation: Minnesota Studies in the Philosophy of Science,* vol. 13. Minneapolis: University of Minnesota Press, 3–219.
——— (1998), "The Importance of Scientific Understanding", in W. Salmon, *Causality and Explanation,* New York: Oxford University Press, 79–91.
Scriven, M. (1962), "Explanation, Predictions, and Laws", in Herbert Feigl and Grover Maxwell (eds.) *Minnesota Studies in the Philosophy of Science,* Minneapolis, MN: University of Minnesota Press, 170–230.
Solomon, Miriam (1992), "Scientific Rationality and Human Reasoning", *Philosophy of Science* 59: 439–455.
Thagard, Paul (1989), "Explanatory Coherence", *Behavioral and Brain Sciences* 12: 435–467.
Toulmin, Stephen (1961), *Foresight and Understanding.* London: Hutchinson.
Trout, J. D. (1992), "Theory Conjunction and Mercenary Reliance", *Philosophy of Science,* 59: 231–245.
——— (1998), *Measuring the Intentional World: Realism, Naturalism, and Quantitative Methods in the Behavioral Sciences.* New York: Oxford University Press.
Watson, James D. (1968), *The Double Helix.* New York: Atheneum.
Woodward, James (1993), "A Theory of Singular Causal Explanation", in David–Hillel Ruben (ed.), *Explanation,* New York: Oxford University Press, 246–274. Originally published in *Erkenntnis* 21: 231–262.