



WHAT
SCIENCE
IS AND
HOW IT
REALLY
WORKS

JAMES C. ZIMRING

What Science Is and How It Really Works

JAMES C. ZIMRING

University of Virginia



CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE
UNIVERSITY PRESS

University Printing House, Cambridge CB2 8BS, United Kingdom

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre,
New Delhi – 110025, India

79 Anson Road, #06–04/06, Singapore 079906

Cambridge University Press is part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of education, learning, and research at the highest international levels of excellence.

www.cambridge.org

Information on this title: www.cambridge.org/9781108476850

DOI: 10.1017/9781108569149

© James C. Zimring 2019

This publication is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press.

First published 2019

Printed in the United Kingdom by TJ International Ltd, Padstow Cornwall

A catalogue record for this publication is available from the British Library.

Library of Congress Cataloging-in-Publication Data

Names: Zimring, James C., 1970– author.

Title: What science is and how it really works / James C. Zimring (University of Virginia, Charlottesville).

Description: Cambridge ; New York, NY : Cambridge University Press, 2019. | Includes bibliographical references and index.

Identifiers: LCCN 2018058439 | ISBN 9781108476850 (hardback : alk. paper) | ISBN 9781108701648 (pbk. : alk. paper)

Subjects: LCSH: Science–Methodology. | Reasoning.

Classification: LCC Q175.32.R45 Z56 2019 | DDC 501–dc23

LC record available at <https://lccn.loc.gov/2018058439>

ISBN 978-1-108-47685-0 Hardback

ISBN 978-1-108-70164-8 Paperback

Cambridge University Press has no responsibility for the persistence or accuracy of URLs for external or third-party internet websites referred to in this publication and does not guarantee that any content on such websites is, or will remain, accurate or appropriate.

Contents

<u>Acknowledgments</u>	<i>page ix</i>
Introduction	1
<u>PART I</u>	<u>17</u>
1 <u>The Knowledge Problem, or What Can We Really “Know”?</u>	19
2 <u>Adding More Building Blocks of Human Reasoning to the Knowledge Problem</u>	44
3 <u>Holistic Coherence in Thinking, or Describing a System of How Humans Reason and Think</u>	66
<u>PART II</u>	<u>103</u>
4 <u>How Scientific Reasoning Differs from Other Reasoning</u>	105
5 <u>Natural Properties of a Rule-Governed World, or Why Scientists Study Certain Types of Things and Not Others</u>	130
6 <u>How Human Observation of the Natural World Can Differ from What the World Really Is</u>	155
7 <u>Detection of Patterns and Associations, or How Human Perceptions and Reasoning Complicate Understanding of Real-World Information</u>	185

8	<u>The Association of Ideas and Causes, or How Science Figures Out What Causes What</u>	208
	PART III	239
9	<u>Remedies That Science Uses to Compensate for How Humans Tend to Make Errors</u>	241
10	<u>The Analysis of a Phantom Apparition, or Has Science Really Been Studied Yet?</u>	279
11	<u>The Societal Factor, or How Social Dynamics Affect Science</u>	299
12	<u>A Holistic World of Scientific Entities, or Considering the Forest and the Trees Together</u>	335
13	Putting It All Together to Describe “What Science Is and How It Really Works”	353
	About the Author	381
	Index	382

Acknowledgments

I am delighted to acknowledge the long list of people to whom I am indebted for all of their help, support, encouragement, and critical feedback. The list is sufficiently long that I will inevitably omit someone whom I need to thank; if you are that person, please accept my appreciation and apology. It should not be taken that those whom I acknowledge necessarily hold my views; my goal is to express my thanks without making them involuntarily complicit.

I have heartfelt gratitude for those who have been my scientific teachers and mentors, who patiently supported and guided my intellectual development, and without whom I would never have experienced science at all. These include Michael Carmichael, Dennis Liotta, Margaret Offermann, Judith Kapp, Christopher Hillyer, Tristram Parslow, James AuBuchon, and others. Your kindness, patience, and guidance have been a *sine qua non* of my development. A special thanks to David Jollow for patiently guiding my development for 30 years.

I would be remiss not to acknowledge the debt I owe to the students who took my course on this topic, first at Emory University and later at the University of Washington. As is typically the case in teaching, I likely learned as much, if not more, from my students than they learned from me. I am, and shall remain, grateful for all the critical debate and intellectual engagement from a truly wonderful group of young and energetic minds. I hope each of you have found or are finding the scientific careers and pursuits that you seek.

I have a special and heartfelt thanks to give to each of the students, fellows, staff scientists, and lab members with whom I have worked in the labs that I have been a part of or that I have overseen. I have shared scientific journeys with each of you, navigating ever-changing webs of belief, which has been the joy of my professional

life. I give special thanks to Krystalyn Hudson, Seema Patel, Maxime Desmarets, Justine Liepkalns, Chantel Cadwell, Matthew Schniederjan, Arielle Medford, Linda Kapp, Rene Shaw, Brian Harcourt, Jennifer Harcourt, Gloria Kelly, Florence Roan, Kyle McKenna, Travis Spurling, Christopher Gilson, Kate Henry, Geetha Mylvaganam, Alice Long, Nikki Smith, Ashley Bennett, Traci Chadwick, Suzanna Schott, Kathryn Girard-Pierce, Jeanne Hendrickson, Sean Stowell, Hayley Waterman, Karen de Wolski, Heather Howie, Ariel Hay, Jacqueline Poston, Xiaohong Wang, and Kimberly Anderson.

To my professional friends and colleagues, who have patiently tolerated my incessant babble on the topic of this book and who have continued to support and encourage me, I do not have the words to express my profound gratitude. There are many of you and I appreciate you all; in particular, I would like to thank Steven Spitalnik, Patrice Spitalnik, Eldad Hod, David Masopust, Vaiva Vezys, Jaap Jan Zwaginga, Gary Falcon, Theresa Prestwood, Zoe Falcon, Suzanna Schott, Taylor Maxson, Angelo D'Alessandro, Carla Fowler, Dana Devine, Krystalyn Hudson, Melissa Hudson, Richard Francis, Tiffany Thomas, Jeanne Hendrickson Jill Johnsen, Stephanie Eisenbarth, and John Luckey.

I would like to express my extreme gratitude to those who contributed to the editing, development, and production of this manuscript. My thanks to Susan Beckett, Robert Swanson, and especially Al Desetta for his expert help in editing several versions of this book. I give my special thanks to members of Cambridge University Press (Katrina Halliday, Anastasia Toynbee, Aleksandra Serocka, and others who have so graciously accepted and developed this project), and to the anonymous peer-reviewers whose identities I shall never know but whose comments and criticisms were essential and very much appreciated. I especially thank Dr. Steven Spitalnik, Dr. Eldad Hod, Dr. Jacqueline Poston, Dr. Jeffrey Kasser, and Dr. Paul Humphreys for their insights and comments on a number of components of the text.

Most importantly, and above all else, I would like to give my everlasting love and gratitude to my family. To my mother and father for literally making me who I am and always supporting and encouraging my life's journey, to my wife Kim for taking that journey with me with endless love and support, and to my daughter Alex, the light of my life, who reminds me every day of the wonder that our world holds. To my brother Jon for a lifetime of debate and laughter, to Michelle, Matthew, Joshua, Kayla, David, Margaret, John, and Nicole for everything you are.

James C. Zimring

Note added on proof:

While this book was already in press, and the final typesetting had been done, a new book was published by Dr. Lee McIntyre entitled "The Scientific Attitude: Defending Science from Denial, Fraud, and Pseudoscience." In this work, Dr. McIntyre advances novel concepts and defines the attitude by which scientists approach evidence as a necessary component of science, which can help to defend it against deniers and pseudoscientists. Had "The Scientific Attitude" been available to me while this book was still being written, I would have certainly referenced it highly, especially in sections discussing how scientists handle data and the general property of attempting to mitigate error as a characteristic of science. These ideas are fully developed in Dr. McIntyre's book that was published six months prior to the current work.

There are whole academic fields (separate from the practice of science itself) that focus precisely on how scientists study nature; in particular, the philosophy, history, psychology, anthropology, and sociology of science, as well as integrated studies that combine these different areas to give a broader view. Whereas my fellow scientists and I study nature, these fields study us. Sadly, in my experience, practicing scientists are seldom very much aware of the details of such fields and in many cases may not know that such fields even exist, other than having heard rumors to that effect. This is not to say that scientists aren't very good at practicing science; theoretical knowledge of the inner workings of an internal combustion engine and the ability to drive a car are separate areas of understanding, and one can be an expert in one while being entirely ignorant of the other. Many scientists can recognize good science when they see it and can call out flawed science when it is encountered, or at least they believe they can, as this is the basis for the entire peer-review system in science. However, because they can judge science does not necessarily mean they can clearly articulate the theoretical underpinnings of what defines science, of how science works and/or how science fails.

If even professional scientists are not typically trained in the general underpinnings of scientific knowledge claims, how is there any hope for nonscientists to understand what level of confidence they should (or should not) place in science and the claims it makes. Moreover, how can people be expected to distinguish valid scientific claims from all manner of flim-flam and fluff? The goal of this work is to help lay people, students of science, and professional scientists understand and explain how science works in general, the strengths and weaknesses of scientific thinking, and the extents and limits of scientific knowledge claims.

Despite the weight that the label of science may carry with many people, it is an utter fiction that there is (or ever has been) a uniform consensus among scientists (or anyone else for that matter) as to what precisely defines science. This question has been tackled over the years by many great scholars and yet there is not a clear and

unequivocal answer. Nevertheless, much progress has been made, and this has generated a greater understanding of characteristics of science, its practice, and its strengths and limitations. The goal of this work is to communicate a broad view of that progress. This is an ambitious goal to be sure, but the difficulty of the task does not diminish its importance. What has been learned is surprising, counter-intuitive, and complex. Ultimately, it speaks not only to science but to the human condition itself.

GOALS AND ORGANIZATION OF THIS WORK

When I introduce the reader to concepts generated by the outside fields that study science itself, I am reflecting the insights, innovations, and contributions of others – standing on the shoulders of giants. I name the giants when I can and have taken care to try and point out to the reader what the original source of many concepts are and what resources and further reading one might do to explore the more granular details of different specialized areas of focus. However, the richness of these different fields goes so deep, that much will be neglected – other works devoted to the finer nuances and details of each component field are abundant, and one need only seek them out. Herein I attempt to synthesize key ideas into a unified framework, hopefully making it coherent to the reader. In addition to explaining the progress of those who study science, I contribute the perspective of the thing being studied – a view of great utility. I speak with the voice of the bacterium on the observations, interpretations, and theories of the microbiologist, for I believe that I know what it is to grow in the chaotic ferment of the microbial culture.

The book is organized into three parts. In Part I (Chapters 1–3), the individual working parts of scientific reasoning and logic are described (and then an attempt is made to draw a picture of scientific reasoning as a whole). In Part II (Chapters 4–8), flaws that undermine natural human observation, perception, and reasoning will be described. In Part III (Chapters 9–13), I will explore how scientific processes and methods try to address these flaws, attempting a

distinction between scientific and nonscientific thinking. An overarching theme of the final part of the book is how science mitigates the tendencies of normal human thinking to “get the world wrong” in particular situations.

The first goal of this book is to help guide nonscientists in having reasonable expectations of what science can and can't do. Scientific claims are often regarded with either too much confidence or too much skepticism by different groups of the general public. This book strives to lay the groundwork for a healthy balance in how to weigh scientific knowledge claims. The second goal of this book is to help professional scientists gain a better understanding and codification of the strengths and weaknesses of their craft and the role they play in portraying it. Of high importance to this latter audience is the recognition that it is quite intoxicating, from an ego standpoint, for scientists to be regarded as the arbiters of “true” knowledge. This has been described as “the Legend” of science.¹

The extreme version of the Legend claims that “science aims at discovering the truth, the whole truth, and nothing but the truth about the world” – a less grandiose version of the Legend states that science is “directed at discovering truth about those aspects of nature that impinge most directly upon us, those that we can observe (and, perhaps, hope to control).” While it is argued that the Legend has been abandoned by those who study science, the Legend (or a slightly weakened version of it) seems very much alive among some in the lay public. In my experience, scientists themselves hold onto a version of the Legend, and while it is less extreme than a philosophical truth, it nevertheless has some component of being “truer” than that which is not science. In my view, scientists should neither seek nor accept the extreme versions of the Legend, which are pleasant in the short term but harmful in the long term, and ultimately destructive as they lead to unsupportable claims. Failures to live up to hyperbolic

¹ Kitcher P. 1995. *The Advancement of Science: Science Without Legend, Objectivity Without Illusions*. New York: Oxford University Press.

attributes only leads to anti-Legend, those who claim with vitriolic hostility that science is, at best, nothing at all, and at worst, a grand conspiracy to dupe the world. Rather, we must seek a balanced and honest view based on realistic assessments. Science's greatest apparent weaknesses are, in actuality, its greatest strengths; as professional scientists we should embrace this and not seek to minimize or ignore it. In the greatest traditions of scientific scholarship, let the existing data of what science is inform us as to the properties of science itself – let us look it in the eye, unflinching, and without spin or propagandist inclinations.

A RIDICULOUSLY BRIEF HISTORY OF SCIENCE VS. NONSCIENCE

Early on (dating back to antiquity) and arguably from a position of great overconfidence, scholars of science often stated that science (or natural philosophy as it was called before the 1830s) dealt with facts, whereas other schools of thought dealt with opinions. However, as many established scientific facts were later rejected by subsequent generations, they came to be understood as fallible and thus not so different from opinions.² Yet the realization that scientific facts are imperfect doesn't mean that they don't have a different character than nonscientific knowledge claims. But if they do (which is not a given), why is that so, and what is the justification for such a view? Later thinkers gravitated to the notion that if it is not fact that distinguishes science from other ways of knowing, then it must be the manner by which scientific claims are generated and/or evaluated that distinguishes science from nonscience. In other words, the method that science uses to create knowledge has a special character that is different from nonscientific approaches.

² Laudan L. 1983. "The Demise of the Demarcation Problem." In Cohen RS, Laudan L (Eds.), *Physics, Philosophy and Psychoanalysis: Essays in Honor of Adolf Grunbaum*. pp. 111-27. Dordrecht: D. Reidel.

For the reasons just stated, most modern attempts at defining science have focused on *methods* or *modes of thinking* that distinguish scientific activities from nonscientific activities rather than the specific content of scientific knowledge claims. However, while one often encounters discussions of “the scientific method” and its application to investigation, there is a lack of agreement about what precisely this method entails, and there are those who argue that the very notion of a scientific method is itself an utter myth.³ It has further been argued that different areas of scientific study favor different types of method(s), and thus one cannot precisely define “science” or “the scientific method” per se.

Moreover, it has been argued that even if some broader characteristics can help identify a method as scientific, precisely demarcating how science differs from other ways of thinking is neither possible nor useful.⁴ It has even been claimed that science flourishes only with a distinct lack of required methodology, and that attempts to codify a scientific process will only serve to destroy it – in other words, the only rule of science is that “anything goes.”⁵ However, this latter view is somewhat radical and is certainly not embraced by most professional scientists, as evidenced by certain generally agreed-upon standards used in the practice of peer review of reports of scientific discoveries, grant applications, and research.

Scholars of science have often rejected any definition that would render the great historical scientists as “nonscientific.” While this seems logical, it presupposes that those who have made the most progress and achieved the most recognition were those acting most scientifically – a question we shall explore in detail. Perhaps more importantly, it assumes that the scientific method, however we define it, has been stable over time, a claim that seems hard to justify.

³ Bauer HH. 1992. *Scientific Literacy and the Myth of the Scientific Method*. Urbana and Chicago: University of Illinois Press.

⁴ Laudan L. 1983.

⁵ Feyerabend P. 1975. *Against Method: Outline of an Anarchist Theory of Knowledge*. New York: New Left Books.

In order to understand science, it is necessary to jettison the unrealistic hyperbole that has been mistakenly assigned to it (from a number of sources), and that has been perpetuated by practicing scientists and science enthusiasts. The flaws of science need to be called out, and greater attention must be directed to its problems, weaknesses, and the limits of what it can show us. We must turn the scientific microscope back on itself and dissect the specimen with a critical and analytic eye, without succumbing to the tendency to give descriptions that are unjustifiably favorable. By describing science as it really is, warts and all, we can simultaneously view science realistically and more accurately differentiate it from other ways of thinking. That science is imperfect and flawed doesn't mean that it isn't distinct from other knowledge systems or that it can't be described, if not defined. Likewise, its flaws and imperfections don't prevent it from being the most effective means (thus far) of exploring and understanding nature. Just as democracy may be "the worst form of government except for all the others,"⁹ the same could be said for science's role in understanding the natural world. The goal of this book is to view science more realistically, not to make vain and misguided attempts to defend a grandiose view that is out of step with the actual entity.

The depiction of science as a logical and orderly process, governed by a specific method and leading to firm facts about nature, although regrettably a distortion of how science is really carried out, is the byproduct of how scientific findings are communicated among scientists. The reasons for this will be expanded upon later, but for now it's important to note that while such distortions may be necessary to communicate scientific findings efficiently, it is profoundly damaging to present this illusion of science rather than the reality of how it is practiced. Practicing scientists typically understand that the distortion does not reflect the reality of the situation. However, those

⁹ This quote is attributed to Winston Churchill, but apparently he was quoting an earlier source that remains unidentified.

outside a field (or outside science altogether) may miss this distinction, believing science to be other than it is. Indeed, this distortion has likely contributed to the genesis of the Legend in the first place. In trying to understand science we can mistake the mirage for the desert. We must be willing to accept that the tempting oasis is merely an image, and we must focus instead on understanding the desert itself, which is the reality of the situation.

MISSING THE FOREST FOR THE TREES

Attempts to describe and understand science have often focused on its component parts, which is a necessary process of any deep analysis of an entity. Efforts have been made to distinguish science from nonscience based on logical constructs, the sociology of science, the psychology of science, and the history of science. However, while each of these areas plays a central role in what it means to practice modern science, none of them tell us the whole story. Trying to understand science exclusively through analysis of its parts is like the ancient story of the three blind humans each studying a different part of an elephant.¹⁰ The person feeling the legs may think it's a tree, the person feeling the tail may think it's a rope, and the person feeling the trunk may assume it's a snake. Each is correct in their observations, but to understand what an elephant really is requires a broader view that merges the component parts into a greater system. This has been recognized in recent decades, and academic disciplines that attempt to generate an overall synthesis of what science is and how it works have emerged.

Modern science is a combination of multiple working parts, including: advanced instrumentation and approaches to observation, human perception and cognition, computational analysis, the application of logic and reasoning, and the effect of social bodies on how

¹⁰ This story can be found at least as far back as the Buddhist text *Udana* 6.5 from the middle of the first millennium BCE (and likely earlier) (https://en.wikipedia.org/wiki/Blind_men_and_an_elephant).

investigation is conducted and interpreted. To grasp modern science, each of these factors must be accounted for; we need to understand both the forest and the trees, as neither has its fullest meaning without the other. In the last century alone, technologies and methodologies have greatly increased the scope of what we can observe (and also misinterpret) in ways never before possible. Cognitive psychology has taught us much about common errors in human reasoning, perception, and observation. Computational capacity allows us to generate and analyze previously overwhelming volumes of data and to make comparisons and analyses far beyond the capacity of the human mind and, in doing so, to lead to new errors an individual human mind would have a hard time making. Statistics has made great strides in its ability to analyze levels of error and to quantify uncertainty, to determine the nature of underlying mechanisms by the distribution of data, and to evaluate associations. Philosophers of science and logic have provided us with a much clearer understanding of the strengths and shortcomings of reasoning than ever before, as well as novel insights into the nature of evidence and the extent to which one can actually verify or reject an idea. Sociologists and anthropologists have learned a great deal about the effect of group dynamics and scientific societies on thinking. Linguists and philosophers of language have identified sources of ambiguity and miscommunication. To understand the current limits of science, each of these must be examined.

The understanding that science is a complex machine, with multiple working parts that need to be understood both individually and in aggregate, is essential. One cannot understand how an internal combustion engine works as a whole without understanding the function of a spark plug. Yet the full implications of what a spark plug does are unintelligible without a preexisting understanding of the entire engine. To break into a circle of codependent knowledge such as this, one may have to visit, and then revisit, the parts and the whole. In this book, the different individual components of science are described first. Later chapters then illustrate the interactions of individual parts

as a system. For this reason, the reader is encouraged to loop back to earlier parts of the book, if needed, during the development of the narrative. In this way, the full implications of the properties of the individual parts of science may become clearer in later sections, when the whole system is described.

SCIENCE IS AN EXTENSION OF HUMAN THINKING
 THAT ALSO VIOLATES ASPECTS OF HUMAN
 THINKING

Science is often portrayed as something very different from normal human behavior or thinking. It seems fair to suspect that science is indeed distinct from normal human thought in at least some ways; after all, most humans are not scientists. However, just because science may differ from typical thinking in some fundamental ways doesn't mean that it is entirely foreign to human cognition. Rather, only a very small number of differences between science and normal human thinking distinguish the two. This may explain part of the difficulty in attempting to demarcate science from nonscience. Because they are so closely related, it's easy to point to apparent exceptions that violate any potential distinction. In other words, defining characteristics attributed to science have been rejected by some precisely because it is easy to find the same characteristics in thinking that is agreed to be nonscientific. Likewise, thinking that is ubiquitous in nonscientific pursuits has been easy to identify as an important component of methods of practicing scientists. Large categorical differences between science and nonscience are not to be found; rather, the small differences between them hold tremendous weight but can also be difficult to pin down.

The differences between scientific and normal thinking, while small, are nevertheless both fundamental to science and also deeply baked into normal human cognition. This in part explains why scientists must undergo so much formal training; they must first become aware of certain normal human tendencies and then learn how to manage and/or overcome them. One must learn how to ignore certain

parts of what it is to think like a human – a difficult task for a thinking human to accomplish. While humans have been on this Earth for a long time, modern science has only been an activity of ours for about 400 years. The generation of science has been a development of human understanding, not of biology – prehistoric humans had essentially the same brains we do today but didn't have advanced science and technology. Scientists are engaging in a learned practice no different than any technique or skill that is developed and refined over time. But the reasoning and thinking that has been learned have subtle but essential parts that are different than our natural (or traditional) reasoning, otherwise we would have had science all along. Exploration of these small differences is key, as well as developing an understanding that because they may violate our natural thinking they feel "wrong" or "counterintuitive," while actually being quite correct.

DEFLATING SCIENCE TO A REALISTIC AND THEREFORE DEFENSIBLE ENTITY

It is an aim of this work to deflate common scientific hyperbole to a realistic and therefore more defensible description. Despite its accomplishments, the ability of science to predict nature is always limited in multiple ways. Scientific predictions and conclusions can never be certain, never be perfect, are certainly never infallible – nothing is ever "proven" in a formal sense of the word. Even things that science labels "Laws of Nature" are themselves reversible if later understanding arises that requires their modification or even wholesale rejection. Ironically, when combined with its other properties (and this is key), it is precisely the recognition of its own limits and imperfections and the practices to which such recognitions then give rise, that constitutes science's greatest strength.

There are many systems of belief that provide much better explanations of experience than does science. Indeed, some systems can explain why anything and everything occurs; science makes no such claim, and those who would suggest that science currently has these ambitions are misguided. Other systems often claim to know

PART I

I The Knowledge Problem, or What Can We Really “Know”?

And since no one bothered to explain otherwise, he regarded the process of seeking knowledge as the greatest waste of time of all.

– Norton Juster, *The Phantom Tollbooth*

GENERAL DESCRIPTION OF THE KNOWLEDGE PROBLEM

Francis Bacon, one of the luminaries of modern science, is thought to have said that “knowledge is power.” Since Bacon made that statement, it has become abundantly clear that humans have a very distinct and difficult “knowledge problem.” There is a fundamental defect in how we come to know anything, and while this is recognized as a problem, the depths of the problem are seldom appreciated and even less frequently discussed. At first glance such a statement may seem ridiculous. What is the problem in saying someone knows something? I know where I am and what I’m doing. I know the names and faces of my friends, family, and acquaintances. I know how to drive a car, how to cook (at least somewhat), and how to pay bills. In fact, just to navigate the tasks of daily life one has to “know” a great number of things.

The knowledge problem in its classic form is not a challenge to one’s knowledge of the things that one has observed or the techniques one has acquired. No one questions that you know that you have a car, that you know you’re married, or that you know you own a collection of Elvis Presley commemorative plates from the Franklin Mint that you inherited from your grandmother.¹ However, the word

¹ Much time and energy has been spent in classic philosophy debating whether we can actually know anything of the external world. However, in everyday life, it is generally accepted that our experience is the result of some external reality that is actually out there.

“knowledge” takes on a very different character as soon as one goes beyond that which is directly observed or experienced. Substantial problems emerge the moment that knowledge claims are extended to things that have not yet been observed, either now or in the past or future. An additional and separate problem emerges once one claims knowledge regarding the relationships and associations between things. Over the past two millennia, as historians, sociologists, anthropologists, and philosophers have analyzed how claims to knowledge arise, develop, and collapse, there has been an expanding appreciation of just how limited our ability is to know.

Since ancient times humans have been on a quest for a higher form of knowledge that can make universal claims. Knowledge in its most ambitious form consists of fundamental truths about which we can be certain, about which we cannot be mistaken or wrong, of which we are sure. Facts and understanding of this kind can be forever considered true; we no longer need to worry about their validity, as these are things that must be so. We can put them in the “true folder” on our computer and forget about having to continue questioning them, they are certain. This is the meaning of knowledge in its extreme form, and it is with this form of knowledge that the knowledge problem is most pronounced.

For some people, it is both unacceptable and distasteful to admit that there is no certain knowledge; they hold particular ideas and convictions with absolute certainty, and there are many systems of belief constructed on such premises. For others, certain knowledge does not and need not exist, as it is not required to navigate the world and enjoy one’s life.² From a pragmatic point of view, if an understanding works, then it is useful, even if in substance it only reflects some misunderstanding. In reading this book, I would ask those who believe in certain knowledge to have an open mind when we analyze

² Even for skeptics of knowledge, many would hold that forms of mathematics and logic constitute certain knowledge; the potential problems and limitations of this view are discussed later.

the basis of its claims to certainty. I would likewise ask the pragmatist to consider that the problem of certain knowledge does not confine itself to ivory tower epistemology, but extends its tentacles into pragmatic knowledge, as we shall see. There are serious implications and ramifications associated with the view that if a theory works, then it is a useful theory regardless of its “truth.”

PREDICTING THE UNOBSERVED

The knowledge problem is most evident when we discuss our ability to predict that which has not yet been observed. Most people would say that they “know” the Sun will rise tomorrow. However, can we call this a certainty? It seems very likely, but it has also been predicted that the time will come (hopefully far in the future) when the Sun runs out of fuel, swells massively, and consumes the Earth. We don’t know if this prediction is true, but it is consistent with our best understanding and we cannot rule it out. It is also possible that the Earth will explode due to some internal process with its molten core, which we had not anticipated. A massive comet that our telescopes have not observed may crash into the Earth and destroy the planet. These examples seem a bit extreme, but consider the 230,000 people who died in the tsunami in Sumatra in 2004, which resulted from an undersea megathrust earthquake that had not been anticipated. The most reasonable prediction on that day was that it would be an average day, like so many days before it, not that a massive wave was going to destroy many thousands of lives; tragically, such was the case.³

Another problem with the concept of knowledge is caused by the association of things. Humans are highly skilled at observing patterns of associations. Whenever there are dark clouds, sounds of thunder, and flashes of lightning, we consider it more likely to rain than if such things are not observed. The more people smoke

³ Wikipedia. n.d. “2004 Indian Ocean earthquake and tsunami.” https://en.wikipedia.org/wiki/2004_Indian_Ocean_earthquake_and_tsunami

that can give us profound advantages over those who are less experienced or completely inexperienced. If you had to subject yourself to a surgical procedure, would you prefer a surgeon who had successfully performed the same operation hundreds of times or a doctor who had never done the operation even once? The second or third time you travel through a foreign country, travel through an airport, or even go to a restaurant, you have abilities that you didn't previously have. Basically, you "know the drill" – where the bathrooms are, what the different lines are for, what documents you need, and what the culture is like. Do you remember your first day of high school? For many of us it was a terrifying thing for a number of reasons, not the least of which was not knowing how to navigate an unfamiliar environment (forget for a moment that the madness of adolescence was clouding our feeble minds). However, as days and weeks went by, we became familiar with the place and the process, and were able to navigate a system and structure that we previously found confusing and intimidating.⁵

Induction is a natural form of human thinking that is practiced routinely and often unknowingly, and is required for everyday navigation of the world. It is basically the use of experience to predict events that one has not yet encountered. I distinctly remember an argument I had with my daughter in our kitchen one Saturday morning. She was 7 years old at the time and quite displeased with whatever it was I was telling her. She folded her arms across her chest, scrunched her face in frustration, and blurted out, "You can't predict the future!" My response was, "Of course I can. I predict that if I push this salt shaker off the counter, it will fall." Which I proceeded to do. She responded, "That's not what I meant. You can't *really* predict the future." She summarily dismissed my argument and stomped off in frustration. This incident illustrates a point that gives thinking animals with

⁵ As was very much the case in my own experience, gaining the ability to navigate does not imply any manner of success or social acceptance; however, at the very least, I had a better idea of what humiliations to expect.

memories a profound advantage over other kinds of creatures. In fact, I had predicted the future and the prediction had held. It wasn't a stunning or unexpected prediction, and it was in a very limited context, but the fact remains that I had predicted the outcome of an event that had not yet occurred, and my prediction was spot on correct. I foresaw that the salt shaker would fall, as every previous salt shaker I had ever dropped had fallen; I had induced the prediction.

In more general terms, induction can be described as predicting the quality or behavior of the unobserved based on the observed. When you are only concerned with what has already been observed, that's not induction, it's description. In other words, if one were to restrict statements of knowledge to that which has already been experienced, the observations speak for themselves. I might simply state that every salt shaker I have dropped from my hand has fallen. It would actually be safer to state that I *perceived* every salt shaker I remember dropping to have fallen. If one restricts statements to the already observed, then one can make very clear statements about the perceived properties, but no predictions about the future are being made. Again, this is not induction but observation and only leads to encyclopedic information about things and situations already encountered. In this case, knowledge is no longer power, or at the very least a far less useful power, to the extent that power is the ability to predict and control – the ability to promote or prevent something.

Induction's immense power comes precisely from its ability to predict the future – that which has not yet occurred or been observed. However, this power comes with a tremendous vulnerability. The successful prediction of the falling salt shaker depended, as does all induction over time, on patterns in the future resembling patterns in the past. I have dropped a great many things in my life and almost all of them have fallen; indeed, every salt shaker I have ever dropped has fallen. So, it is easy to induce that when you drop things that are not otherwise supported, they fall. (The exception would be things that are less dense than air, e.g., helium balloons.) Yet just because things

have behaved one way in the past does not necessarily mean they will continue to do so in the future – this assumption is the Achilles' heel of induction.

At first, this problem with induction seems a very common-sense sort of thing that doesn't set off any alarms. Everyone knows that things change, that things don't generally stay the same forever, and that there are times when past experience no longer applies. However, the gravity of this problem may be highly underestimated. A classic example of this problem with induction was put forth by Bertrand Russell, who described a chicken raised by a farmer. Every day of the chicken's life the farmer came out and fed the chicken. We'll assume that the chicken couldn't talk to the other chickens or to anyone else for that matter, and therefore the chicken's specific life constituted the entirety of its information. Hence, from the chicken's point of view, on every day that had ever existed, the farmer approached the chicken and gave it food. It would be a very reasonable induction for the chicken to predict that on the following day the farmer would once again give it food. Regrettably, when the next morning arrives, the farmer wrings the chicken's neck, plucks its feathers, and cooks it for supper – a tragic failure of induction to be sure – at least for the chicken.⁶ The example of the chicken is highly applicable to human behavior. I have not yet died in a car accident; thus, I do not predict that I will die in a car accident today, and I feel comfortable driving – a mistake made by the more than 3,000 humans who die in car accidents each day worldwide. The assumption that the future will resemble the past is a highly useful assumption, but it is by no means certain. In some cases, it is almost inevitably false.

A practical example of falsely assuming that the future will resemble the past can be found with the advent of antibiotics. When

⁶ The very same event may have been a failure of induction for the chicken and a successful induction for the farmer, who might frequently eat chickens he is raising. Of course, his perspective of the event that is being repeated is different – the raising of a chicken over time vs. day-to-day feedings.

penicillin was first used therapeutically in humans, it was observed that the administration of penicillin in patients infected with gonorrhoea was uniformly efficacious in killing off the bacteria. One might be tempted to conclude a general principle – that penicillin kills gonorrhoea. In fact, this became an accepted practical truth, and penicillin was listed by the medical community as the definitive treatment for gonorrhoea. However, given the selective pressure of widespread penicillin use, some strains of gonorrhoea acquired resistance to penicillin through evolutionary processes. Thus, whereas essentially 100% of gonorrhoea was observed to be sensitive to penicillin in the past, such is not the case at the current time, a clear example of the fallibility of induction in being able to predict the future.

The problem of predicting future events by induction can be expanded to include the assumption that relevant modifiers of future situations will also resemble the past. In other words, the assumption that all things are equal – that one is always comparing apples to apples. I have a vivid recollection of the first time I gave my daughter a helium balloon (she was 9 months old at the time). When I handed it to her, she was extremely upset that the balloon “fell up” instead of falling down. It was an unpredicted event, because up to that point in her life 100% of everything she dropped had fallen down. Thus, it would have been very reasonable for her to predict that the balloon, like every other object, would in fact fall down. Induction failed in this case because the generalized rule did not happen to extend to this particular situation (i.e., helium balloons are not the same as other dropped objects). This problem with induction was not that the future didn’t resemble the past, but rather that situational changes in the future didn’t line up with the past. If I were to hold a salt shaker in my hand and then let it go while I was a passenger in the international space station, I would likely observe a very different result than in my kitchen on Earth. However, in the case of both the helium balloon and the space station, the future exactly resembled the past – as far as we know, helium balloons have thus far always floated in the atmosphere of Earth and salt shakers have always floated in outer

space; the failure of my prediction was that I didn't understand how other circumstances and modifiers had changed.

The problem of background circumstances is ubiquitous and takes place in everyday interactions. We all know how frustrating it is to receive unsolicited advice from strangers that doesn't seem to apply to our situations. Most of us have seen a child having a meltdown in a public place and the parents (or other responsible adults) struggling to calm the child. For those of us who have been that struggling parent, it seems that onlookers have a variety of responses: sympathy, relief that it is not their problem, annoyance at being disturbed, and disapproval of the child, the parents, or both. In many cases the onlookers are critical of how the parents are handling the situation and, in some cases, can't resist giving "helpful" advice.

The problem with such advice is that every child is different, every parent is different, every child-parent dynamic is different, and there are all manner of specific modifiers that may affect a given situation (i.e., the child's pet might have died, the child might be on the autism spectrum, or the family may have different cultural norms or be facing some unusual stress, etc.). In most cases the person giving the advice has limited experience with a small number of children and yet feels comfortable generalizing his or her advice to this child, and maybe to all children. Of course, there are some generalities to human behavior, and certain advice may very well apply, but for obvious reasons, it may not. This is most acutely felt when one's parents offer advice (often unsolicited and typically obnoxious) about how their grandchildren are being raised, because their advice is no longer applicable (and in some cases no longer legal), coming from a generation when corporal punishment was not only allowed but encouraged, when car seats had not been invented, and when there was no problem with chain smoking in the nursery. Likewise, conversely, it is easy for children to criticize their parents' past behavior when held up to current norms, which were not in existence when the behavior in question was taking place. In all these cases generalizations about what "should be done" are being drawn that may not be valid, because

Humans cannot help but induce in all aspects of life, as it is one of the fundamental ways by which we navigate the world. A person who has complete amnesia, or who cannot form new memories and is thus deprived of the ability to induce due to lack of conscious memory and experience, is at a tremendous disadvantage in the world. Induction has thus far been superior to random guessing or untargeted trial and error; however, as stated previously, it is not a path to certain knowledge about unobserved things, and it can be (and will be) tragically wrong at times.

REJECTING POPULATIONS BASED ON SINGLE INSTANCES

The experience of life is distinct and particular to each of us. Forget for a moment that even when faced with the same experience, we may each perceive it differently; clearly, we each encounter a particular set of conditions and life events, and we each have a different interface with the world. While we may also incorporate the information of others through communication, we still have direct access to only a very small slice of the pie that is our world. Most of what is in existence (the universe) is simply not accessible to us, and we know little of even that to which we do have access. What percentage of people do you actually know in your hometown, on your street, or in your workplace? For the nearly 50% of Americans who live in large cities, it is likely that you know very few of the people in your general proximity and very little about them. Certainly, none of us has met a significant percentage of the approximately 7 billion people on Earth, seen a significant amount of the 197 million square miles of the Earth, encountered a significant number of the animals on Earth, etc. Yet, in order to use the power of induction to help us navigate the world, it is necessary that we make generalizations of some kind. Basing such generalizations on the small amount of data we have seems like a better guess than basing it on no data at all.

While we may be stuck doing our best to navigate the world with what we have, it is nevertheless a big problem to reject factual

claims made about populations by using minuscule sample sizes. Nevertheless, this seems to be an enduring human trait. One reads that, on average, smoking increases the risk of getting lung cancer. This is a population-based argument. A group of smokers will have a rate of lung cancer that is 23 times more likely (for males) and 13 more likely (for females) than for similar populations who do not smoke.¹¹ However, this situation is often offered as the answer to the question: Does smoking cause lung cancer? When faced with such a statement, it is common to hear, “Well, you may say that smoking causes lung cancer, but my grandfather smoked four packs of unfiltered cigarettes for 35 years, and he never got lung cancer.” This may have been the case for the grandfather, and that’s a great thing for him, but it is irrelevant to the claim that smoking, on average, increases the risk of lung cancer. The claim was not that smoking causes lung cancer (i.e., if a person smokes, then he or she will get lung cancer) in the same way that removing someone’s head causes death.¹² By definition, if smoking increases rates of lung cancer to anything less than 100%, then the population-based argument can be true even though some people will smoke their whole lives and never get lung cancer.¹³

Positive assertions of generalizations are no less based on minuscule data sets than are the rejections of assertions. One might go to two different restaurants and have a wonderful dining experience at one of them and a horrible dining experience at the other. Based on this experience, one rates the first restaurant as good and the second restaurant as horrible. However, the first restaurant may have gotten the wrong shipment of food that day, receiving excellent ingredients

¹¹ This statistic is in reference to small cell and non-small cell lung cancer (80%–90% of lung cancers). U.S. Department of Health and Human Services. “The Health Consequences of Smoking – 50 Years of Progress: A Report of the Surgeon General, 2014.” www.surgeongeneral.gov/library/reports/50-years-of-progress/index.html

¹² At least so far, with the current limits of technology.

¹³ There is also a percentage of people who never smoke and still get lung cancer. This compounds the problem of reconciling population claims with individual data due to a logical fallacy, because the claim that smoking increases rates of lung cancer in no way suggests that smoking is the only cause of lung cancer; it is not claimed that smoking is necessary for lung cancer to occur.

instead of the bargain basement, outdated food they normally buy to save money. In contrast, the second restaurant may have had both cooks and half their servers call in sick that day. The statements we tend to make are not that one particular meal was good at one place and poor at the other. Rather, we conclude one restaurant is good and the other bad, a generalized statement.

During the 2016 American presidential election, much regrettable rhetoric was passed around regarding whether or not individuals of the Muslim faith, or even of Middle Eastern background (regardless of faith), should be allowed into the country, or whether they should even be eligible to run for president if already citizens based on the assertion that people of the Muslim faith tend to be terrorists. As tragic as terrorist events have been in the Western world (and I use the Western world merely as a basis of comparison, not meaning to imply terrorism is any less tragic anywhere else), the perpetrators of these acts represent a very small number of individuals out of a world population of 1.6 billion Muslims (22% of all living humans on Earth). Surely, one cannot draw a meaningful generalization about 1.6 billion people based on the actions of a handful of individuals. If one were to look at Muslim-related terrorism in the United States, fewer than 20 individuals in recent years have engaged in terrorist acts out of 1.8 million Muslims in the country. This in no way rejects the observation that terrorist acts can be carried out by people of this group or that some extreme variation of ideology may drive the actions of these few individuals. However, this is a very small quantity of evidence to justify broader generalizations about Muslims. If anything, we can conclude that 99.9% of Muslims in the United States are not terrorists, the very opposite of what the rhetoric was suggesting. Moreover, this situation is a prime example of the availability heuristic (heuristics will be discussed in Chapter 4) combined with the base rate fallacy. When someone perpetrates a terrorist act, the media tells us the characteristics of that person. However, the media seldom (if ever) tells us the number of people with the same characteristics who don't carry out such acts.

This tendency to draw generalized knowledge from scant data may be the best we can do as individuals, as performing population-based studies is not a typical activity of humans; even if we were so inclined to engage in systematic study, most of us have neither the resources nor the ability to do so. However, the fact that individuals are doing the best they can doesn't mean their best is necessarily doing it well. Moreover, even when we do have access to the population data (e.g. with Muslims and terrorism), we are prone to ignore it. As will be discussed in more detail later in the book, it can be argued that the study of science has focused on (and analyzed) only a very few scientists and drawn general conclusions based on them. Moreover, by focusing on the scientists who have made the most progress (or at least are the most famous), those who study science bias themselves to the extreme of the population, potentially hobbling any ability to capture what scientists do in general (or as a group).

WHY PROBABILITY-BASED THINKING DOESN'T HELP WITH INDUCTION

A common approach to the problem of induction, which is often invoked in response to the previously stated concerns, is to state induced knowledge claims in probabilistic terms. This applies both to making statements about unobserved entities in the present and also across time. For example, if one had observed 99 ravens and all were black, then one might induce the statement that "all ravens are black." However, if the 100th raven observed was not black, we wouldn't throw up our hands in frustration at having no knowledge of ravens. Rather, one would simply modify the knowledge claim by saying that "99% of observed ravens are black." This maneuver is simply restating the data to modify a principle about all ravens. This can then be used to predict unobserved events from a probabilistic point of view; you can't tell what color the next raven you encounter will be, but you can say that 99% of the time it will be black and 1% of the time it will be nonblack – not with absolute certainty regarding the next raven, but with predictive power regarding a whole population and the relative likelihood of what color the next raven will be.

A probabilistic point of view can't predict an individual event, but there is no reason it can't make predictions about populations with great accuracy.

Although probability statements may bring comfort to some people, they fail to help much with the knowledge problem itself and with the issue of induction. The reason probability determinations do not help with the knowledge problem is that even if the probability statement is true with a capital “T”, it cannot provide the ability to predict next events with certainty. While a probability statement can tell you the odds that the next raven will be black, the next raven can only be either black or nonblack.¹⁴ Being able to state the likelihood that the next raven will be black is a type of prediction. Nevertheless, even if one has absolute knowledge of a population, it does not speak to specific cases, and thus one still cannot predict particular events. When most people talk to their doctor they don't want to know what their probability of getting cancer is; they want to know whether or not *they* themselves will get cancer.

Another problem with probability statements is that, much like simple induction itself, one can never rule out things changing in the future. After observing another 100 ravens, the 99% probability determination may change again, and in fact will change, unless 99 of the next ravens are black and one is nonblack. Thus, while the 99% probability determination may be better than random guessing, it is not knowledge about which we cannot be wrong. Let us retreat even further from our desire for absolute knowledge and stipulate that the more ravens we observe, the better and better our probability determination will become.¹⁵ This seems a justifiable statement (often called the law of large numbers). This is just another way of saying that the closer we get to having observed every raven, the closer we

¹⁴ This example uses categorical classifications and assumes that there are distinct colors as opposed to simply being a continuum of colors. Although it can be debated whether clean and distinct categories truly exist in nature, humans nevertheless tend to think in categorical terms, and there certainly does seem to be some basis (if not an absolute basis) for categories.

¹⁵ This resembles a more Bayesian approach.

additional bears (nonpolar bears) also live at the North Pole. Hence, bears at the North Pole may consist of some polar bears and some brown bears. This possibility does not necessarily make the statement false, as it doesn't guarantee that brown bears will be at the North Pole; however, it doesn't rule it out and thus allows for the possibility that the conclusion is incorrect. In other words, the conclusion is not necessarily true and is thus a fallacy that doesn't lead to certain knowledge.

Like induction, deduction is a common tool of human reasoning, without which we wouldn't navigate the world as well as we do. While Aristotle may have first named and characterized deduction, it is not something that Aristotle invented. Rather, he described a process that, like induction, is a normal part of everyday human thinking. Deductive thought in humans can be found in children as young as preschoolers.¹⁷ This is not to say that humans are perfect deducers; indeed, a whole body of studies has shown that we tend to deduce incorrectly, especially in certain circumstances.¹⁸

The correct application of formal logic is a highly complex and well-developed field, much of which is difficult to learn and certainly not intuitive. Nevertheless, like induction, deduction is a normal part of human thinking that we deploy as part of our navigation of the world. However, errors in deduction are also a normal human trait. Moreover, when we make such errors, we often feel as though we have reasoned our way to a correct conclusion, even though we have actually failed to do so. It is for this reason that logicians have invented specific ways to express logical statements, have defined different types of logic and the rules by which they work, and have made tremendous progress in such thinking. Indeed, much of mathematics can be described as a deductive language.

¹⁷ Hawkins RD, Pea J, Glick J, Scribner S. 1984. "Merds That Laugh Don't Like Mushrooms: Evidence for Deductive Reasoning by Preschoolers." *Developmental Psychology* 20: 584–94.

¹⁸ Evans, J St BT. 2017. "Belief Bias in Deductive Reasoning." In Rüdiger PF (Ed.). *Cognitive Illusions*. pp. 165–81. New York: Routledge.

While very powerful, deduction does not solve the knowledge problem. The first thing to note, which is a fundamental limit to deductive reasoning, is that it doesn't generate information about the unobserved; rather, it only reveals complexities that are already contained within the premises, but which may not be intuitively obvious until the deductive reasoning is fully carried out. In other words, no new information has been generated that wasn't already contained within the premises; nevertheless, without the syllogism, the fullest meaning of the stated facts could not be demonstrated and may not be appreciated. This seems to be a real limitation to deduction, as without the ability to make any predictions about the unobserved, our ability to predict or control is limited. However, this limitation can be overcome if the premises are universal, thus allowing the deduction of universal conclusions. In other words, consider premises that include the type of language of "every A is a B" or "no A is a B." Based upon such universal premises, one can deduce knowledge statements that apply to every instance of A, even instances that have not been experienced. Thus, one *is* deducing knowledge of the unobserved. This is one reason why deductivist thinkers tend to prefer premises of a universal type (e.g., all As are Bs), for without such universal premises the conclusions are not universal. If the conclusions are not universal, then one cannot make statements (with certainty) about unobserved things. If one has not achieved certainty of unobserved things, then one has not gained true knowledge (at least as we have defined it), and the knowledge problem remains unsolved.

If deduction can generate true knowledge so long as it uses premises of a universal nature, then where is the problem? The problem is in being able to determine a justifiable premise of a universal nature. For centuries, a number of notable philosophers have believed that humans have some inherent ability to recognize natural truths. However, in recent times, neurology's and cognitive psychology's understanding of human perception and thinking has advanced to the point that we now appreciate that humans can be pretty terrible at perceiving the world right in front of them, let alone coming up

with universal statements of truth (this is explored in detail in later sections). If there is a single error in a premise upon which a deduced system of knowledge is built, then the whole system may come crumbling to the ground. If the premises are not certain, then the knowledge is not certain, no matter how good the reasoning. If there is no reliable source for certain premises, then deductive thinking cannot solve the knowledge problem.

Some of our greatest institutions have solved the premise problem by simply stating that a given premise is true. For example, the U.S. Declaration of Independence states: "We hold these truths to be self-evident, that all men are created equal, that they are endowed by their Creator with certain unalienable Rights, that among these are Life, Liberty, and the pursuit of Happiness." In other words, these truths are self-evident because we say so (so there!), and we will now build a system of beliefs based in part on this premise.¹⁹ If the truths really were self-evident, then this might be okay, but what justification is there for such an assertion other than the authors stating that they hold them to be self-evident? – in other words, their opinion, and what is the justification that their opinions are correct? In the same fashion, many religions are based on an irrefutable premise that a given god or gods exist. Likewise, many systems of belief, even without a formal deity, state the presence of a force, or energy, or structure in the universe. Such premises of gods or forces are certainly not without evidence; indeed, evidence of the divine can be obtained through an experience of god, through observable consequences that would come about in the event that there was a god, or even through revelation. One can feel universal sources through a spiritual experience or perceive the effects of such forces in the world.

One might then argue that there is no knowledge problem for philosophies that invoke self-evident premises or for religions that consider the experience of a god or revelation of the divine as sources

¹⁹ No claim is being made here that the American system of government was deduced, just that it makes claims of self-evident and universal premises.

of unequivocal truth. However, as will be explored in detail later, it should be noted that such systems do not typically deduce an entire system of belief, at least not in any formal sense of deduction, from the premises that are stated, and thus it is a kind of an apple and orange situation. Moreover, while feeling or perceiving something can be a compelling force in persuading an individual that the thing exists, perceptions and feelings are highly prone to error and misinterpretation, and thus do not provide a justification of knowledge that stands up to analytic thinking. This is not to say that such justification is not sufficient for theology or spiritual systems of belief, but it is clear that such justification is fallible. How many religions have existed in human history, many of which believed with certainty that they were the one true way? For this to be true, all of them, except one, must be wrong, and it is not clear that any of them need be right. Thus, theological revelation appears to be able to get things wrong. Hence, while religions typically state things in certain terms, and may lead to certain belief, they do not lead to certain knowledge. The making of claims with certainty and explanations of everything will be explored later as one of the criteria by which one can demarcate some categories of non-science from science.

If we stipulate that humans have no access to fundamental premises, or first conditions, through either revelation or innate knowledge of such premises, then how is one to use deduction? If the premises are not certain to be true, then no matter how valid the deductive reasoning, the results are not certain to be true, which undermines the whole deductive program for generating knowledge. One might point to Euclid, who stated certain premises and was then able to deduce a complex geometry that was very useful in describing the natural world. Likewise, Sir Isaac Newton stated certain premises (laws of motion) from which he deduced a system of mechanics that could describe and predict motions of the planets with great accuracy and how forces worked on bodies in general. Isn't the amazing predictive capacity of these systems a validation of the correctness of their premises? Regrettably, as we shall explore later, such is not the

case. Of note, given modern theories of special relativity and the curved nature of time-space, both Newton's and Euclid's systems are considered to be profound intellectual achievements of great theoretical and practical value, but ultimately, these systems are not entirely correct due to the premises being not entirely correct.

At the end of the day, there is no clear way around the major problem of deductive knowledge. In order to have any ability to predict unobserved nature, deduction must make statements that are universal. Due to the problems of induction, universal statements based on experience cannot be justified, and no other source of universal premises seems supportable.

Although both induction and deduction have the problems described, in real life, one nevertheless uses induction and deduction (or at least reasoning that resembles deduction) together to navigate the world. Induction provides the justification for premises based on experience (albeit an imperfect justification). Deductive reasoning helps reason forward from the induced premises to generate all manner of new understanding of association within the induced premises. Hence the combination of induction and deduction certainly leads to new ideas that would have come from neither alone, but fails to solve the problems of either. In aggregate, the knowledge problem is solved by neither induction, deduction, nor a combination of the two.

THE UTILITY OF UNCERTAIN CONCLUSIONS: IS THE KNOWLEDGE PROBLEM REALLY A PROBLEM?

It seems that a solution to the knowledge problem is not likely to be forthcoming. However, how much of an impediment is this? It brings us to the question: What makes useful knowledge, and does useful knowledge have to be universally certain to be meaningful? Many thinkers have adopted a pragmatist school of thought that has placed value on scientific theories if the theories work in the real world. If a theory predicts the natural world, then it is a useful theory, regardless of whether or not it is ultimately true. Knowledge may be flawed, to be sure, and it may not result in any kind of absolute, objective truth.

What, then, is retroduction? It consists of the process by which one generates ideas regarding the causes of things already observed; in Peirce's words, retroduction is "the process of forming explanatory hypotheses." In other words, retroduction enters our thinking with regards to the association between observed effects and the things we speculate caused those effects. Retroductions are to be found everywhere and in every walk of life, as a ubiquitous part of normal human thinking. One awakens one morning to see snow on the ground that was not present the night before, and retroduces that it must have snowed during the night. One arrives home to see a spouse's car in the driveway and retroduces that one's spouse has already arrived at home. One receives an email from the address of a friend and retroduces that the friend sent the email. Based on evidence, we are guessing at a cause of the effects we observe.

One could argue that this is just experience-based thinking – the simple carrying out of enumerative induction – for example, all of the previous times when there was snow on the ground it had snowed, previously when one's spouse's car was in the driveway, the spouse was in the house, and when previous emails were sent from a certain address, it was a particular friend who had sent them. However, retroduction is quite distinct from induction. It is not observing something and predicting a general principle from what is being observed; it is not a "more of the same" conclusion. Rather, it is suggesting a *previous* (hence "retro-") entity that led to the observed outcome. With retroduction, one is typically positing something that occurred in the past to explain the cause of an experience, not predicting observations that have not yet been experienced. Moreover, retroduced causes can be entirely novel things that have never been observed (e.g., positing that an invisible evil demon is the cause of an outbreak of an illness). Hence, retroduction is distinct from induction. Whereas induction may lead to generalized conclusions and predictions of as of yet unexperienced things, retroduction makes a retrograde guess at the causes of an already observed phenomenon. Importantly, retroduction need not be limited to temporal and causal

entities but can also apply to laws and principles that explain without invoking cause; however, for the current discussion we will focus on retrodution of causal entities as our main example.

Although retroductions may posit as of yet unknown causes of observed effects, good retroductions don't just randomly guess. Rather, the retroduced cause or causes should only be entities from which the observed effects would follow, or at the very least be consistent. An example might be taken from cancer epidemiology. Let us say there are an unusually high number of childhood leukemia cases in a small Arkansas town; we will assume this observation to be accurate and correct. One induces the generalization that children living in this particular town are at an increased risk of getting cancer. The process of retrodution would begin with the generation of a hypothesis capable of explaining the induced generalization (i.e., increased likelihood of leukemia).

For example, suppose one retroduced the presence of a carcinogen in the drinking water of the Arkansas town and that this particular carcinogen caused higher rates of leukemia in children. Let us also stipulate that the children in this town frequently consumed the drinking water. All other things being equal, then one can deduce what was observed – that there would be higher rates of leukemia in the children in this town. The term “all other things being equal” is a big logical lift, and it is difficult (if not impossible) to ever justify such a condition in the real world, as we shall explore later. Nevertheless, scientists often think in this way, and even if a formal deduction is not possible, one can at least infer the observed outcome based on this retrodution. It is not an acceptable retrodution to guess at a cause that wouldn't lead to the observed effect. In other words, if we accept that consuming candy has no link to cancer, but retroduce that the higher rates of cancer in the town are due to eating more candy, this retrodution is not consistent with accepted evidence. Even if children eat more candy, it would not lead to that which we are trying to explain. This is not to say that people don't make retroductions that wouldn't lead to what is being explained – but at the very least, these

are bad retrroductions. For the retrduced hypothesis to be scientificly useful, it must at least have the potential to lead to the known effect. To be scientificly useful, it must also lead to other (as of yet unobserved) effects, as we shall explore later in the text.

Retroduction, then, is a third and separate mode of reasoning, and by adding it to induction and deduction an integrated model of scientific reasoning can begin to emerge. However, before such a model is synthesized, the problems of retroduction need to be explored in more detail. Just as retroduction is distinct from induction and deduction, its advantages and problems are also distinct. A specific problem with retroduction is called the “fallacy of affirming the consequent” and necessitates some exploration.

THE FALLACY OF AFFIRMING THE CONSEQUENT

Consider this statement: if A, then B. In other words, if A is the cause, then B must be the result. As an example, one can say that if an individual falls off the roof of a 30-story building and lands on concrete (all other things being equal³), the person will be injured. If Bill unequivocally falls off a 30-story building, you can conclude with certainty that Bill will be injured.⁴ If A, then B; A occurs, therefore B must occur.

However, one cannot logically go in the other direction; in other words, given the statement: If A, then B, one cannot conclude that if B occurs then A must also have occurred. Why is this so? If one finds Bill lying on the sidewalk injured, why can’t you conclude with certainty that Bill fell off the 30-story building alongside the sidewalk? The reason is that all other kinds of causes may have resulted in Bill’s injury (e.g., being hit by a car, having fallen out of a plane, being clumsy and falling down). In other words, as illustrated in Figure 2.1,

³ In this case, “all other things being equal” means without introducing any other modifiers (e.g., the person doesn’t have a parachute, the building is on Earth, the rules of gravity apply)

⁴ It is not being claimed that there has never been a report of someone falling off a 30-story building and not being injured, but as an example, we can accept that all will at least be injured somehow – if only a scratch or bruise.




FIGURE 2.1 Cause, effect, and affirming the consequent

consider that A, C, or D each cause B. If one knows that C or A or D occurred, then one can conclude with certainty that B occurred. However, knowing that B occurred, one cannot conclude with certainty that A occurred, as B might have been caused by C or D (or even by some as of yet unknown cause not in the figure).

Why is this fundamental issue so important? Because a great deal of normal human thinking and interacting with the world is precisely the process of observing B and positing A, which is basically the process of retrodution.

In our personal lives, our professional lives, and in local and world events, we are constantly observing effects and retroduting causes. We observe that the Earth is getting warmer, and people propose different causes. We observe that one country invades another, and we speculate as to the motivation. Each day the stock market goes up or down, and all manner of financial punditry puts forth multiple theories about what caused the change.

At a more basic level, someone wins the lottery and everyone begins to seek a reason, such as where the lucky person purchased the ticket or what clothes the person was wearing when he or she did so. People get sick and we have no explanation, so we begin to suggest causes (e.g., toxins in the water, or radio towers, or plastics, or vaccines). The classic murder mystery is still another example – a dead body is found, and the detectives guess at suspects and try to figure out who murdered the victim.

In each of these cases, many different people simultaneously posit various theories. Although we don't normally use the word, in

all of these cases people are retroducting hypotheses to explain observed effects, which is a common process. Moreover, in doing so, people are evoking the fallacy of affirming the consequent, and this is why the speculation goes on so long. Affirming the consequent is an inexorable logical defect in the process of retroduction itself, and thinking alone cannot remedy this defect. This doesn't mean retroduction may not get the right answer; it just means we can never be certain that it has. It is for this reason that it has been questioned whether retroduction is even a form of logic at all; however, it is certainly a form of thinking, and its utility seems clear (if not at all certain).

So, in the case of the town in Arkansas, how many different hypotheses are equally consistent with the data about high rates of cancer? The answer is that there are an infinite number, limited only by the imagination of the thinker.

1. There is a toxin present in the town that causes cancer.
2. The townspeople are infected with a virus that causes cancer.
3. There is a cancer-causing genetic mutation present at a higher rate than in the general population in the families that live in the town.
4. There is a hole in the ozone layer over this city, resulting in more cancer-causing ultraviolet sunlight affecting the citizens.
5. There are high voltage electrical lines in this city that cause cancer.
6. Large deposits of magnetic rock near the city cause cancer.
7. The combination of toxins and a magnetic field are causing cancer.
8. There is a toxin in the fish living in a nearby lake, and eating this fish causes cancer.
9. Secret government radiation experiments are being carried out on the town's children.
10. Space aliens are abducting the kids and implanting them with cancer-causing probes.
11. Past immoral behavior by the kids' parents has caused bad karma and cancer.
12. God is punishing the town.

Here we see the fundamental problem with retroduction and with all hypothesis-based thinking. For any given observation, there

the cancer outcome will occur – as such, this is not deduction in the traditional sense. Nevertheless, the increased rates are predictable, and for the rest of the text we will use HD to include these types of examples.

HD then takes an additional step, which has emerged as an attempt to address the fallacy of affirming the consequent; in other words, to help narrow down the number of possible causes of an observed effect. This step is to use deduction (or statistical inference in the case of probabilistic hypotheses) to make additional predictions, which have not yet been tested (e.g., no observation or attempt at observation has yet been made; Figure 2.2D). The importance of this additional point cannot be emphasized enough, as it provides a response to the problem of affirming the antecedent (albeit an imperfect response). If new effects that can be predicted from a retroduced cause do not occur, then that cause is no longer a valid retroduction, as its presence would not lead to observed outcomes. While all the retroduced hypotheses may predict the initial observation or observations, different hypotheses ultimately lead to at least some different predictions, by which the hypotheses can be whittled down.⁶ Thus, by identifying additional predictions and testing them, the validity of retroduced ideas can be assessed. Although philosophers and scientists have used HD to characterize scientific thinking, there is nothing uniquely scientific about HD thinking. As is the case with induction, deduction, and retroduction, HD thinking is found in everyday human thinking and problem solving.

Let's look at an example of HD thinking in everyday life. One cold morning you get into your car, turn the key, and the engine doesn't turn over. You have observed that your car doesn't start. You retroduce that the cause is a dead battery. You have just used the first part of the HD method – you have retroduced a hypothesis to explain why your car wouldn't start.

⁶ Arguably, if two hypotheses do not lead to any differences in deducible outcomes, then they are functionally indistinguishable as different hypotheses.