

CHOICE AND CHANCE

An Introduction to Inductive Logic

Fourth Edition

BRIAN SKYRMS

UNIVERSITY OF CALIFORNIA, IRVINE



Australia • Canada • Mexico • Singapore • Spain
• United Kingdom • United States

Publisher: Eve Howard
Acquisitions Editor: Peter Adams
Assistant Editor: Kerri Abdinoor
Editorial Assistant: Mindy Newfarmer
Marketing Manager: Dave Garrison
Project Editor: Christal Niederer
Print Buyer: Mary Noel

Permissions Editor: Joohee Lee
Production Service and Composition: Progressive Publishing Alternatives
Cover Designer: Laurie Anderson
Cover and Text Printer/Binder: Webcom Ltd.

COPYRIGHT © 2000 Wadsworth, a division of Thomson Learning. Thomson Learning is a trademark used herein under license.

ALL RIGHTS RESERVED. No part of this work covered by the copyright hereon may be reproduced or used in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, Web distribution, or information storage and retrieval systems—without the written permission of the publisher.

Printed in Canada

2 3 4 5 6 7 03 02 01 00

For permission to use material from this text, contact us by

Web: <http://www.thomsonrights.com>

Fax: 1-800-730-2215

Phone: 1-800-730-2214

For more information, contact
Wadsworth/Thomson Learning
10 Davis Drive
Belmont, CA 94002-3098
USA
<http://www.wadsworth.com>

International Headquarters

Thomson Learning
290 Harbor Drive, 2nd Floor
Stamford, CT 06902-7477
USA

UK/Europe/Middle East/ South Africa

Thomson Learning
Berkshire House
168-173 High Holborn
London WC1V 7AA
United Kingdom

Asia

Thomson Learning
60 Albert Complex #15-01
Albert Complex
Singapore 189969

Canada

Nelson/Thomson Learning
1120 Birchmount Road
Scarborough, Ontario M1K 5G4
Canada

Contents

Preface		vii
CHAPTER I	BASICS OF LOGIC	1
I.1	Introduction	1
I.2	The Structure of Simple Statements	1
I.3	The Structure of Complex Statements	2
I.4	Simple and Complex Properties	8
I.5	Validity	9
CHAPTER II	PROBABILITY AND INDUCTIVE LOGIC	12
II.1	Introduction	12
II.2	Arguments	12
II.3	Logic	14
II.4	Inductive versus Deductive Logic	17
II.5	Epistemic Probability	23
II.6	Probability and the Problems of Inductive Logic	26
CHAPTER III	THE TRADITIONAL PROBLEM OF INDUCTION	30
III.1	Introduction	30
III.2	Hume's Argument	30
III.3	The Inductive Justification of Induction	35
III.4	The Pragmatic Justification of Induction	44
III.5	Summary	49
CHAPTER IV	THE GOODMAN PARADOX AND THE NEW RIDDLE OF INDUCTION	51
IV.1	Introduction	51
IV.2	Regularities and Projection	52
IV.3	The Goodman Paradox	56
IV.4	The Goodman Paradox, Regularity, and the Principle of the Uniformity of Nature	60
IV.5	Summary	67

CHAPTER V	MILL'S METHODS OF EXPERIMENTAL INQUIRY AND THE NATURE OF CAUSALITY	69
V.1	Introduction	69
V.2	Causality and Necessary and Sufficient Conditions	69
V.3	Mill's Methods	73
V.4	The Direct Method of Agreement	74
V.5	The Inverse Method of Agreement	79
V.6	The Method of Difference	85
V.7	The Combined Methods	90
V.8	The Application of Mill's Methods	94
V.9	Sufficient Conditions and Functional Relationships	97
V.10	Lawlike and Accidental Conditions	101
CHAPTER VI	THE PROBABILITY CALCULUS	109
VI.1	Introduction	109
VI.2	Probability, Arguments, Statements, and Properties	109
VI.3	Disjunction and Negation Rules	111
VI.4	Conjunction Rules and Conditional Probability	117
VI.5	Expected Value of a Gamble	128
VI.6	Bayes' Theorem	130
VI.7	Probability and Causality	134
CHAPTER VII	KINDS OF PROBABILITY	137
VII.1	Introduction	137
VII.2	Rational Degree of Belief	137
VII.3	Utility	138
VII.4	Ramsey	141
VII.5	Relative Frequency	144
VII.6	Chance	146
CHAPTER VIII	PROBABILITY AND SCIENTIFIC INDUCTIVE LOGIC	151
VIII.1	Introduction	151
VIII.2	Hypothesis and Deduction	151
VIII.3	Quantity and Variety of Evidence	153
VIII.4	Total Evidence	154
VIII.5	Convergence to the Truth	156
Answers to Selected Exercises		158
Index		171

Preface to the Fourth Edition

My first year out of graduate school I was asked to teach a beginning undergraduate course in Inductive Logic. The only books available were monographs that were clearly unsuitable as texts. I struggled through the course as best I could. The following summer I wrote *Choice and Chance*, motivated largely by an instinct for survival. My goal was to keep the exposition clear, uncluttered, and accessible while, at the same time, getting things right. I also wanted a book that imparted logical skills while remaining focused on basic philosophical problems. Judging by the gratifying reception that the book has been given, many other philosophers have the same basic priorities.

I have tried to keep these priorities in mind in successive revisions of the book, while keeping the material up to date. In preparing the fourth edition of *Choice and Chance* I have reorganized, simplified, and extended the material in the third edition. I have also added *Answers to Selected Exercises* at the end of the book to provide the student with quick feedback on his or her understanding of the material.

There is a new first chapter, *Basics of Logic*, which is a brief introduction to the logic of “and,” “or,” and “not.” This material was previously in the chapter on *Mill’s Methods of Experimental Inquiry*, but I think it is better for the student to see it right away. The old chapter on *Coherence* is gone. It seemed to me to contain material that was too complicated for a book at this level. Some essentials of utility theory and Ramsey’s treatment of degree of belief from that chapter survive in the discussion of degree of belief in the chapter on *Kinds of Probability*. The last chapter, *Probability and Scientific Inductive Logic*, is entirely new. It revisits questions raised earlier in the book, shows how probabilistic reasoning ties them together, and provides the beginnings of some constructive answers. I hope that students will view it as an invitation to further study.

In preparing this edition, I was helped by suggestions of Professors Jeffrey Barrett, University of California, Irvine; Joseph Keim Campbell, Washington State University; Brian Domino, Eastern Michigan University; and George Rainbolt, Georgia State University. I was able to consult an instructor’s guide, which Professors Richard Ratliff and Martin Frické of the University of Otago prepared for the third edition, and which they were kind enough to send me. Many improvements to the third edition, which are carried over to this edition, are due to Professors Frans van der Bogert, Appalachian State University; Ellery Eells, University of Wisconsin, Madison; Sidney Luckenbach, California State University, Northridge; David Seimens, Jr., Los Angeles Pierce College; Frederick Suppe, University of Maryland; and Georg Dorn, University of Salzburg. Thanks also to Mindy Newfarmer for seeing the manuscript into print.

Brian Skyrms
University of California, Irvine

I

Basics of Logic

I.1. INTRODUCTION. Deductive logic is the logic of “and,” “or,” and “not.” It is useful for classifying, sorting, or searching and can be used for searching a library, using an Internet search engine, or searching any sort of database. The help section for one Internet search engine explains that searching for “Mary AND NOT lamb” finds documents containing “Mary” but not containing “lamb.” A library database can also be searched for “Aztec OR Toltec” for a history report. The logic of “and,” “or,” and “not” gives us a taste of deductive logic, with which we can compare inductive logic. Deductive logic will also be useful in the analysis of Mill’s methods of experimental inquiry in Chapter V, and in the treatment of probability in Chapter VI.

I.2. THE STRUCTURE OF SIMPLE STATEMENTS. A *statement* is a sentence that makes a definite claim. A straightforward way of making a claim is to (1) identify what you are talking about, and (2) make a claim about it. Thus, in the simple statement “Socrates is bald,” the proper name “Socrates” identifies who we are talking about and the predicate “is bald” makes our claim about him. In general, expressions that identify what we are talking about are called *referring expressions* and the expressions used to make factual claims about the things we are talking about are called *characterizing expressions*. Thus, the name “Socrates” refers to a certain individual, and the predicate “is bald” characterizes that individual.

Although proper names are an important type of referring expression, there are others. Pronouns such as “I,” “you,” “he,” and “it” are referring expressions often used in ordinary speech, where context is relied upon to make clear what is being talked about. Sometimes whole phrases are used as referring expressions. In the statement “The first President of the United States had wooden false teeth,” the phrase “The first President of the United States” is used to refer to George Washington. He is then characterized as having wooden false teeth (as in fact he did).

Although statements are often constructed out of one referring expression, as in the examples above, sometimes they are constructed out of more than one referring expression, plus an expression that characterizes the relationship between the things referred to. For instance, the statement “Mercury is hotter than Pluto” contains two referring expressions—“Mercury” and “Pluto”—and one characterizing expression—“is hotter than.” Characterizing expressions that characterize an individual thing are called *property expressions* or

one-place predicates. “Is bald,” “is red,” and “conducts electricity” are examples of property expressions. Characterizing expressions that characterize two or more individual things in relation to one another are called *relational expressions* or *many-place predicates*. “Is hotter than,” “is a brother of,” “is to the north of,” and “is between” are examples of relational expressions.

The basic way to construct a simple statement is to combine referring and characterizing expressions to make the appropriate factual claim. In the next section it will be seen how these simple statements can be combined with logical connectives to form complex statements.

Exercises

Pick out the referring and characterizing expressions in the following statements. State whether each characterizing expression is a property expression or a relational expression.

1. Tony loves Cleo.
2. Dennis is tall.
3. This book is confusing.
4. Arizona is between New Mexico and California.
5. Los Angeles is bigger than Chicago.

I.3. THE STRUCTURE OF COMPLEX STATEMENTS. Consider the two simple statements “Socrates is bald” and “Socrates is wise.” Each of these statements is composed of one referring expression and one characterizing expression. From these statements, together with the words “not,” “and,” and “or,” we can construct a variety of complex statements:

Socrates is *not* bald.

Socrates is bald *and* Socrates is wise.

Socrates is bald *or* Socrates is wise.

Socrates is *not* bald *or* Socrates is wise.

Socrates is bald *and* Socrates is wise *or* Socrates is *not* bald *and*

Socrates is *not* wise.

The words “not,” “and,” and “or” are neither referring nor characterizing expressions. They are called *logical connectives* and are used together with referring and characterizing expressions to make complex factual claims.

We can see how the logical connectives are used in the making of complex factual claims by investigating how the truth or falsity of a complex statement depends on the truth or falsity of its simple constituent statements. A simple statement is true just when its characterizing expression correctly

characterizes the thing or things it refers to. For instance, the statement “Socrates is bald” is true if and only if Socrates is in fact bald; otherwise it is false. Whether a complex statement is true or not depends on the truth or falsity of its simple constituent statements *and* the way that they are put together with the logical connectives. Let us see how this process works for each of the connectives.

Not. We *deny* or *negate* a simple statement by placing the word “not” at the appropriate place within it. For instance, the denial or negation of the simple statement “Socrates is bald” is the complex statement “Socrates is not bald.” Often we abbreviate a statement by using a single letter; for example, we may let the letter “*s*” to stand for “Socrates is bald.” We may deny a statement by placing a sign for negation, “ \sim ,” in front of the letter that abbreviates that statement. Thus, “ $\sim s$ ” stands for “Socrates is not bald.” Now it is obvious that when a statement is true its denial is false, and when a statement is false its denial is true. Using the shorthand introduced above, we can symbolize this information in the following *truth table*, where T stands for true and F for false:

p	$\sim p$
T	F
F	T

What this table tells us is that if the statement “ p ” is true, then its denial, “ $\sim p$,” is false. If the statement “ p ” is false, then its denial, “ $\sim p$,” is true. The truth table is a summary of the way in which the truth or falsity of the complex statement depends on the truth or falsity of its constituent statements.

And. We form the *conjunction* of two statements by putting the word “and” between them. Each of the original statements is then called a *conjunct*. A conjunction is true just when both of the conjuncts are true. Using the symbol “ $\&$ ” to abbreviate the word “and” we can represent this in the following truth table:

p	q	$p\&q$
T	T	T
T	F	F
F	T	F
F	F	F

Here we have four possible combinations of truth and falsity that the constituent statements “ p ” and “ q ” might have, and corresponding to each combination we have an entry telling us whether the complex statement “ $p\&q$ ” is true or false for that combination. Thus, in the case where “ p ” is true and “ q ”

is true, " $p \& q$ " is also true. Where " p " is true and " q " is false, " $p \& q$ " is false. Where " p " is false and " q " is true, " $p \& q$ " is again false. And where both " p " and " q " are false, " $p \& q$ " remains false.

Or. The word "or" has two distinct uses in English. Sometimes " p or q " means "either p or q , but not both," as in "I will go to the movies or I will stay home and study." This is called the *exclusive* sense of "or." Sometimes " p or q " means " p or q or both," as in "Club members or their spouses may attend." This is called the *inclusive* sense of "or." We are especially interested in the inclusive sense of "or," which we shall represent by the symbol " \vee ." " $p \vee q$ " is called a *disjunction* (or alternation), with " p " and " q " being the *disjuncts*. The truth table for disjunction is:

p	q	$p \vee q$
T	T	T
T	F	T
F	T	T
F	F	F

By reference to the truth tables for " \sim ," " $\&$," and " \vee " we can construct a truth table for any complex statement. Consider the complex statement "Socrates is not bald or Socrates is wise." This complex statement contains two simple constituent statements: "Socrates is bald" and "Socrates is wise." We may abbreviate the first statement as " s " and the second as " w ." We can then symbolize the complex statement as " $\sim s \vee w$." We may use the following procedure to construct a truth table for this complex statement:

Step 1: List all the possible combinations of truth and falsity for the simple constituent statements, " s ," " w ."

Step 2: For each of these combinations, find whether " $\sim s$ " is true or false from the truth table for negation.

Step 3: For each of the combinations, find whether " $\sim s \vee w$ " is true or false from step 2 and the truth table for disjunction.

The result is the following truth table for " $\sim s \vee w$ ":

	Step 1		Step 2	Step 3
	s	w	$\sim s$	$\sim s \vee w$
Case 1:	T	T	F	T
Case 2:	T	F	F	F
Case 3:	F	T	T	T
Case 4:	F	F	T	T

This truth table tells us exactly what factual claim the complex statement makes, for it shows us in which cases that statement is true and in which it is false.

Because a truth table tells us what factual claim is made by a complex statement, it can tell us when two statements make the same factual claim. Let us examine the truth table for “ $(s \& w) \vee (\sim s \& w)$ ”:

	s	w	$\sim s$	$s \& w$	$\sim s \& w$	$(s \& w) \vee (\sim s \& w)$
Case 1:	T	T	F	T	F	T
Case 2:	T	F	F	F	F	F
Case 3:	F	T	T	F	T	T
Case 4:	F	F	T	F	F	F

Note that in reading across the truth table we start with the simple constituent statements, proceed to the next largest complex statements, until we finally arrive at the complex statement that is the goal. The truth table shows that the final complex statement is true in cases 1 and 3 and false in cases 2 and 4. But notice that the simple statement “ w ” is also true in cases 1 and 3 and false in cases 2 and 4. This shows that the simple statement “ w ” and the complex statement “ $(s \& w) \vee (\sim s \& w)$ ” make the same factual claim. To claim that Socrates is either bald and wise or not bald and wise is just a complicated way of claiming that Socrates is wise. When two statements make the same factual claim, they are *logically equivalent*.

Truth tables may also be used to show that two complex statements make conflicting factual claims. For example, the claim made by the statement “ $\sim s \& \sim w$ ” obviously conflicts with the claim made by the statement “ $s \& w$.” Socrates cannot both be bald and wise and be not bald and not wise. This conflict is reflected in a truth table for both statements:

	s	w	$\sim s$	$\sim w$	$s \& w$	$\sim s \& \sim w$
Case 1:	T	T	F	F	T	F
Case 2:	T	F	F	T	F	F
Case 3:	F	T	T	F	F	F
Case 4:	F	F	T	T	F	T

The statement “ $s \& w$ ” is true only in case 1, while the statement “ $\sim s \& \sim w$ ” is true only in case 4. There is no case in which both statements are true. Thus, the two statements make conflicting factual claims. When two statements make conflicting factual claims, they are *inconsistent* with each other, or *mutually exclusive*.

There are some peculiar complex statements that make no factual claim whatsoever. If we say “Either Socrates is bald or Socrates is not bald” we have really not said anything at all about Socrates. Let us see how this situation is reflected in the truth table for “ $sv\sim s$ ”:

	s	$\sim s$	$sv\sim s$
Case 1:	T	F	T
Case 2:	F	T	T

The reason why the statement “ $sv\sim s$ ” makes no factual claim is that it is true no matter what the facts are. This is illustrated in the truth table by the statement being true in all cases. When a complex statement is true, no matter what the truth values of its constituent statements are, that statement is called a *tautology*.

At the opposite end of the scale from a tautology is the type of statement that makes an impossible claim. For instance, the statement “Socrates is bald and Socrates is not bald” must be false no matter what the state of Socrates’ head. This is reflected in the truth table by the statement being false in all cases:

	s	$\sim s$	$s\&\sim s$
Case 1:	T	F	F
Case 2:	F	T	F

Such a statement is called a *self-contradiction*. Self-contradictions are false no matter what the facts are, in contrast to tautologies, which are true no matter what the facts are. Statements that are neither tautologies nor self-contradictions are called *contingent statements* because whether they are true or not is contingent on what the facts are. A contingent statement is true in some cases and false in others.

The purpose of this section has been to convey an understanding of the basic ideas behind truth tables and the logical connectives. We shall apply these ideas in our discussion of Mill’s methods and the theory of probability.

The main points of this section are:

1. Complex statements are constructed from simple statements and the logical connectives “ \sim ,” “ $\&$,” and “ v .”
2. The truth tables for “ \sim ,” “ $\&$,” and “ v ” show how the truth or falsity of complex statements depends on the truth or falsity of their simple constituent statements.

3. With the aid of the truth tables for “ \sim ,” “ $\&$,” and “ \vee ,” a truth table may be constructed for any complex statement.
4. The truth table for a complex statement will have a case for each possible combination of truth or falsity of its simple constituent statements. It will show in each case whether the complex statement is true or false.
5. The factual claim made by a complex statement can be discovered by examining the cases in which it is true and those in which it is false.
6. If two statements are true in exactly the same cases, they make the same factual claim and are said to be logically equivalent.
7. If two statements are such that there is no case in which they are both true, they make conflicting factual claims and are said to be inconsistent with each other, or mutually exclusive.
8. If a statement is true in all cases, it is a tautology; if it is false in all cases, it is a self-contradiction; otherwise it is a contingent statement.

Exercises

1. Using truth tables, find which of the following pairs of statements are logically equivalent, which are mutually exclusive, and which are neither:
 - a. $p, \sim\sim p$.
 - b. $\sim p \vee \sim q, p \& q$.
 - c. $p \& \sim q, \sim(p \& q)$.
 - d. $\sim p \vee q, p \& \sim q$.
 - e. $(p \vee p) \& q, p \& (q \vee q)$.
 - f. $\sim(\sim p \vee q), p \& \sim q$.
2. Using truth tables, find which of the following statements are tautologies, which are self-contradictions, and which are contingent statements:
 - a. $\sim\sim p \vee \sim p$.
 - b. $p \vee q \vee r$.
 - c. $(p \vee p) \& \sim(p \vee p)$.
 - d. $(p \vee \sim q) \vee \sim(p \vee \sim q)$.
 - e. $p \& q \& r$.
 - f. $\sim\sim(p \vee \sim p)$.
 - g. $\sim p \vee p \vee q$.

I.4. SIMPLE AND COMPLEX PROPERTIES. Just as complex statements can be constructed out of simple ones using the logical connectives, so complex properties (or property expressions) can be constructed out of simple ones using “and,” “or,” and “not.” These complex properties are the categories used in “Boolean search” of databases. For example, from “Persian Gulf country,” “Iraq,” and “Iran” you can form the complex property “Persian Gulf country AND NOT (Iraq OR Iran).” We will use capital letters to abbreviate property expressions.

We can use a method to examine complex properties which is quite similar to the method of truth tables used to examine complex statements. Whether a complex property is present or absent in a given thing or event depends on whether its constituent simple properties are present or absent, just as the truth or falsity of a complex statement depends on the truth or falsity of its simple constituent statements. When the logical connectives are used to construct complex properties, we can refer to the following presence tables, where “*F*” and “*G*” stand for simple properties and where “*P*” stands for “present” and “*A*” for “absent”:

Table I		Table II			Table III		
<i>F</i>	$\sim F$	<i>F</i>	<i>G</i>	<i>F</i> & <i>G</i>	<i>F</i>	<i>G</i>	<i>F</i> v <i>G</i>
P	A	P	P	P	P	P	P
A	P	P	A	A	P	A	P
		A	P	A	A	P	P
		A	A	A	A	A	A

Note that these tables are exactly the same as the truth tables for the logical connectives except that “present” is substituted for “true” and “absent” is substituted for “false.” With the aid of these presence tables for the logical connectives, we can construct a presence table for any complex property in exactly the same way as we constructed truth tables for complex statements. The presence table for a complex property will have a case for each possible combination of presence or absence of its simple constituent properties. For each case, it will tell whether the complex property is present or absent. As an illustration, we may construct a presence table for “ $\sim F$ v*G*”:

	<i>F</i>	<i>G</i>	$\sim F$	$\sim F$ v <i>G</i>
Case 1:	P	P	A	P
Case 2:	P	A	A	A
Case 3:	A	P	P	P
Case 4:	A	A	P	P

There are other parallels between the treatment of complex statements and the treatment of complex properties. Two complex properties are *logically equivalent* if they are present in exactly the same cases; two properties are *mutually exclusive* if there is no case in which they are both present. When a property is present in all cases (such as “ $Fv\sim F$ ”) it is called a *universal property*. A universal property is analogous to a tautology. When a property is absent in all cases, it is called a *null property*. A null property is analogous to a self-contradiction. The properties in which we are most interested in inductive logic are those which are neither universal nor null. These are called *contingent* properties.

Exercises

1. Using presence tables, find which of the following pairs of properties are logically equivalent, which are mutually exclusive, and which are neither:
 - a. $\sim FvG, \sim\sim Gv\sim F$.
 - b. $\sim Fv\sim G, \sim(F\&G)$.
 - c. $\sim FvG, F\&\sim G$.
 - d. $Fv\sim(F\&G), \sim(F\&G)\&F$.
 - e. $\sim F\&\sim G, \sim(FvG)$.
 - f. $\sim(FvGvH), FvGvH$.
 - g. $F\&\sim G, \sim(F\&G)$.
2. Using presence tables, find out which of the following properties are universal, which are null, and which are contingent:
 - a. $\sim FvGvF$.
 - b. $(FvF)\&\sim(FvF)$.
 - c. $\sim(Fv\sim F)$.
 - d. $(Fv\sim G)\&(Gv\sim F)$.
 - e. $FvGvH$.
 - f. $\sim(F\&\sim G)v\sim(Gv\sim F)$.

I.5. VALIDITY. We can use the truth tables of section I.3 to investigate whether one statement (the conclusion) follows logically from some others (the premises). If it does, we have a valid argument; otherwise we don't. An argument is *valid* if the conclusion is true in every case in which the premises are all true. The argument:

p
therefore, $p\&q$

is not valid because there is a case in which the premise, “ p ,” is true and the conclusion, “ $p \& q$,” is false. It is case 2 in the following truth table:

	p	q	$p \& q$
Case 1:	T	T	T
Case 2:	T	F	F
Case 3:	F	T	F
Case 4:	F	F	F

But the argument:

p
therefore, $p \vee q$

is valid because every case in which the premise, “ p ,” is true (cases 1 and 2 in the following truth table) is a case in which the conclusion, “ $p \vee q$,” is true.

	p	q	$p \vee q$
Case 1:	T	T	T
Case 2:	T	F	T
Case 3:	F	T	T
Case 4:	F	F	F

Here is an example of a valid argument with two premises:

$\sim p$
 $p \vee q$
therefore, q

We can establish its validity by looking at the following truth table:

	p	q	$\sim p$	$p \vee q$
Case 1:	T	T	F	T
Case 2:	T	F	F	T
Case 3:	F	T	T	T
Case 4:	F	F	T	F

First we find the cases where both premises, “ $\sim p$ ” and “ $p \vee q$,” are true. Only case 3 qualifies. Then we can check that the conclusion, “ q ,” is true in case 3.

Here you have a little taste of deductive logic. In the next chapter we consider a larger picture that includes both inductive and deductive logic.

Exercises

Check the following arguments for validity using truth tables:

1. p
therefore, $p \vee p$
2. p
therefore, $p \& p$
3. $\sim(p \vee q)$
therefore, $\sim q$
4. $\sim(p \& q)$
therefore, $\sim p$
5. $\sim p$
 $\sim q$
therefore, $\sim(p \vee q)$
6. $\sim p$
 $\sim q$
 $p \vee q \vee r$
therefore, r

II

Probability and Inductive Logic

II.1. INTRODUCTION. What is logic as a whole and how do inductive and deductive logic fit into the big picture? How does inductive logic use the concept of probability? What does logic have to do with arguments? In this chapter we give a preliminary discussion of these large questions to provide a perspective from which to approach the rest of the book.

II.2. ARGUMENTS. The word “argument” is used to mean several different things in the English language. We speak of two people having an argument, of one person advancing an argument, and of the value of a mathematical function as depending on the values of its arguments. One of these various senses of “argument” is selected and refined by the logician for the purposes at hand.

When we speak of a person advancing an argument, we have in mind his giving certain reasons designed to persuade us that a certain claim he is making is correct. Let us call that claim which is being argued for the *conclusion* of the argument, and the reasons advanced in support of it the *premises*. If we now abstract from the concrete situation in which one person is trying to convince others and consider the bare bones of this conception of an argument, we arrive at the following definition: An *argument* is a list of *sentences*, one of which is designated as the conclusion, and the rest of which are designated as premises.

But if we consider the matter closely, we see that this definition will not do. Questions, commands, and curses can all be expressed by sentences, but they do not make factual claims nor can they stand as reasons supporting such claims. Suppose someone said, “The Dirty Sox star pitcher has just broken both his arms and legs, their catcher has glaucoma, their entire outfield has come down with bubonic plague, and their shortstop has been deported. Therefore, they cannot possibly win the pennant.” He would clearly be advancing an argument, to the effect that the Dirty Sox cannot win the pennant. But if someone said, “How’s your sister? Stand up on the table. May you perish in unspeakable slime!” he would, whatever else he was doing, *not* be advancing an argument. That is, he would not be advancing evidence in support of a *factual claim*.

Let us call a sentence that makes a definite factual claim a *statement*. “Hannibal crossed the Alps,” “Socrates was a corruptor of youth,” “Every body attracts every other body with a force directly proportional to the sum of their

masses and inversely proportional to the square of their distance,” and “The moon is made of avocado paste” are all statements, some true, some false. We may now formulate the logician’s definition of an argument:

Definition 1: An argument is a list of *statements*, one of which is designated as the conclusion and the rest of which are designated as premises.

In ordinary speech we seldom come right out and say, “A, B, C are my premises and D is my conclusion.” However, there are several “indicator words” that are commonly used in English to point out which statement is the conclusion and which are the premises. The word “therefore” signals that the premises have been run through, and that the conclusion is about to be presented (as in the Dirty Sox example). The words “thus,” “consequently,” “hence,” “so,” and the phrase “it follows that” function in exactly the same way.

In ordinary discourse the conclusion is sometimes stated first, followed by the premises advanced in support of it. In these cases, different indicator words are used. Consider the following argument: “Socrates is basically selfish, since after all Socrates is a man, and all men are basically selfish.” Here the conclusion is stated first and the word “since” signals that reasons in support of that conclusion follow. The words “because” and “for” and the phrase “by virtue of the fact that” are often used in the same way. There is a variation on this mode of presenting an argument, where the word “since” or “because” is followed by a list of premises and then the conclusion; for example, “Since all men are basically selfish and Socrates is a man, Socrates is basically selfish.”

These are the most common ways of stating arguments in English, but there are other ways, too numerous to catalog. However, you should have no trouble identifying the premises and conclusion of a given argument if you remember that:

The conclusion states the point being argued for and the premises state the reasons being advanced in support of the conclusion.

Since in logic we are interested in clarity rather than in literary style, one simple, clear method of stating an argument (and indicating which statements are the premises and which the conclusion) is preferred to the rich variety of forms available in English. We will put an argument into standard logical form simply by listing the premises, drawing a line under them, and writing the conclusion under the line. For example, the argument “Diodorus was not an Eagle Scout, since Diodorus did not know how to tie a square knot and all Eagle Scouts know how to tie square knots” would be put into standard logical form as follows:

Diodorus did not know how to tie a square knot.
 All Eagle Scouts know how to tie square knots.

Diodorus was not an Eagle Scout.

Exercises

1. Which of the following sentences are statements?
 - a. Friends, Romans, countrymen, lend me your ears.
 - b. The sum of the squares of the sides of a right triangle equals the square of the hypotenuse.
 - c. Hast thou considered my servant Job, a perfect and an upright man?
 - d. My name is Faust; in all things thy equal.
 - e. $E = mc^2$.
 - f. May he be boiled in his own pudding and buried with a stick of holly through his heart.
 - g. Ptolemy maintained that the sun revolved around the Earth.
 - h. Ouch!
 - i. Did Sophocles write *Electra*?
 - j. The sun never sets on the British Empire.
2. Which of the following selections advance arguments? Put all arguments in standard logical form.
 - a. All professors are absent-minded, and since Dr. Wise is a professor he must be absent-minded.
 - b. Since three o'clock this afternoon I have felt ill, and now I feel worse.
 - c. Candidate X is certain to win the election because his backers have more money than do Candidate Ys, and furthermore Candidate X is more popular in the urban areas.
 - d. Iron will not float when put in water because the specific gravity of water is less than that of iron.
 - e. In the past, every instance of smoke has been accompanied by fire, so the next instance of smoke will also be accompanied by fire.

II.3. LOGIC. When we *evaluate* an argument, we are interested in two things:

- i. Are the premises true?
- ii. Supposing that the premises are true, what sort of support do they give to the conclusion?

The first consideration is obviously of great importance. The argument “All college students are highly intelligent, since all college students are insane, and all people who are insane are highly intelligent” is not very compelling, simply because it is a matter of common knowledge that the premises are false. But important as consideration (i) may be, it is not the business of a logician to judge whether the premises of an argument are true or false.¹ After all, any statements whatsoever can stand as premises to some argument, and the logician has no special access to all human knowledge. If the premises of an argument make a claim about the internal structure of the nucleus of the carbon atom, one is likely to get more reliable judgments as to their truth from a physicist than from a logician. If the premises claim that a certain mechanism is responsible for the chameleon’s color changes, one would ask a biologist, not a logician, whether they are true.

Consideration (ii), however, is the logician’s stock in trade. Supposing that the premises are true, does it follow that the conclusion must be true? Do the premises provide strong but not conclusive evidence for the conclusion? Do they provide any evidence at all for it? These are questions which it is the business of logic to answer.

Definition 2: *Logic* is the study of the strength of the evidential link between the premises and conclusions of arguments.

In some arguments the link between the premises and the conclusion is the strongest possible in that the truth of the premises *guarantees* the truth of the conclusion. Consider the following argument: “No Athenian ever drank to excess, and Alcibiades was an Athenian. Therefore, Alcibiades never drank to excess.” Now if we suppose that the premises “No Athenian ever drank to excess” and “Alcibiades was an Athenian” are true, then we must also suppose that the conclusion “Alcibiades never drank to excess” is also true, for there is no way in which the conclusion could be false while the premises were true. Thus, for this argument we say that the truth of the premises guarantees the truth of the conclusion, and the evidential link between premises and conclusion is as strong as possible. This is in no way altered by the fact that the first premise and the conclusion are false. What is important for evaluating the strength of the evidential link is that, *if* the premises were true, the conclusion would also have to be true.

In other arguments the link between the premises and the conclusion is not so strong, but the premises nevertheless provide some evidential support for the conclusion. Sometimes the premises provide strong evidence for the conclusion, sometimes weak evidence. In the following argument the truth of

¹ Except in certain very special cases which need not concern us here.

the premises does not guarantee the truth of the conclusion, but the evidential link between the premises and the conclusion is still quite strong:

Smith has confessed to killing Jones. Dr. Zed has signed a statement to the effect that he saw Smith shoot Jones. A large number of witnesses heard Jones gasp with his dying breath, "Smith did it." Therefore Smith killed Jones.

Although the premises are good evidence for the conclusion, we know that the truth of the premises does not *guarantee* the truth of the conclusion, for *we can imagine circumstances in which the premises would be true and the conclusion false.*

Suppose, for instance, that Smith was insane and that he confessed to every murder he ever heard of, but that this fact was generally unknown because he had just moved into the neighborhood. This peculiarity was, however, known to Dr. Zed, who was Jones's psychiatrist. For his own malevolent reasons, Dr. Zed decided to eliminate Jones and frame Smith. He convinced Jones under hypnosis that Smith was a homicidal maniac bent on killing Jones. Then one day Dr. Zed shot Jones from behind a potted plant and fled.

Let it be granted that these circumstances are highly improbable. If they were not, the premises could not provide strong evidential support for the conclusion. Nevertheless, the circumstances are not impossible and thus the truth of the premises does not guarantee the truth of the conclusion.

The following is an argument in which the premises provide some evidence for the conclusion, but in which the evidential link between the premises and the conclusion is much weaker than in the foregoing example:

Student 777 arrived at the health center to obtain a medical excuse from his final examination. He complained of nausea and a headache. The nurse reported a temperature of 100 degrees. Therefore, student 777 was really ill.

Given that the premises of this argument are true, it is not as improbable that the conclusion is false as it was in the preceding argument. Hence, the argument is a weaker one, though not entirely without merit.

Thus we see that arguments may have various *degrees of strength*. When the premises present absolutely conclusive evidence for the conclusion—that is, when the truth of the premises guarantees the truth of the conclusion—then we have the strongest possible type of argument. There are cases ranging from this maximum possible strength down to arguments where the premises contradict the conclusion.

Exercises:

Arrange the following arguments in the order of the strength of the link between premises and conclusion.

1. No one who is not a member of the club will be admitted to the meeting.
I am not a member of the club.

I will not be admitted to the meeting.
2. The last three cars I have owned have all been sports cars. They have all performed beautifully and given me little trouble. Therefore, I am sure that the next sports car I own will also perform beautifully and give me little trouble.
3. My nose itches; therefore I am going to have a fight.
4. Brutus said that Caesar was ambitious, and Brutus was an honorable man. Therefore Caesar must have been ambitious.
5. The weatherman has said that a low-pressure front is moving into the area. The sky is gray and very overcast. On the street I can see several people carrying umbrellas. The weatherman is usually accurate. Therefore, it will rain.

II.4. INDUCTIVE VERSUS DEDUCTIVE LOGIC. When an argument is such that the truth of the premises guarantees the truth of the conclusion, we shall say that it is deductively valid. When an argument is not deductively valid but nevertheless the premises provide good evidence for the conclusion, the argument is said to be inductively strong. How strong it is depends on how much evidential support the premises give to the conclusion. In line with the discussion in the last section, we can define these two concepts more precisely as follows:

Definition 3: An argument is *deductively valid* if and only if it is *impossible* that its conclusion is false while its premises are true.

Definition 4: An argument is *inductively strong* if and only if it is *improbable* that its conclusion is false while its premises are true, and it is not deductively valid. The *degree* of inductive strength depends on how improbable it is that the conclusion is false while the premises are true.²

The sense of “impossible” intended in Definition 3 requires clarification. In a sense, it is impossible for me to fly around the room by flapping my arms;

²Although the “while” in Definition 3 may be read as “and” with the definition remaining correct, the “while” in Definition 4 should be read as “given that” and not “and.” The reasons for this can be made precise only after some probability theory has been studied. However, the sense of Definition 4 will be explained later in this section.

this sense of impossibility is called *physical impossibility*. But it is not physical impossibility that we have in mind in Definition 3. Consider the following argument:

George is a man.
George is 100 years old.
George has arthritis.

George will not run a four-minute mile tomorrow.

Although it is physically impossible for the conclusion of the argument to be false (that is, that he will indeed run a four-minute mile) while the premises are true, the argument, although a pretty good one, is *not* deductively valid.

To uncover the sense of impossibility in the definition of deductive validity, let us look at an example of a deductively valid argument:

No gourmets enjoy banana –tuna fish soufflés with chocolate sauce.
Antoine enjoys banana –tuna fish soufflés with chocolate sauce.

Antoine is not a gourmet.

In this example it is impossible in a stronger sense—we shall say *logically impossible*—for the conclusion to be false while the premises are true. What sort of impossibility is this? For the conclusion to be false Antoine would have to be a gourmet. For the second premise to be true he would also have to enjoy banana –tuna fish soufflés with chocolate sauce. But for the first premise to be true there must be no such person. Thus, to suppose the conclusion is false is to contradict the factual claim made by the premises. To put the matter a different way, the factual claim made by the conclusion is already implicit in the premises. This is a feature of all deductively valid arguments.

If an argument is deductively valid, its conclusion makes no factual claim that is not, at least implicitly, made by its premises.

Thus, it is logically impossible for the conclusion of a deductively valid argument to be false while its premises are true, because to suppose that the conclusion is false is to contradict some of the factual claims made by the premises.

We can now see why the following argument is not deductively valid:

George is a man.
George is 100 years old.
George has arthritis.

George will not run a four-minute mile tomorrow.

The factual claim made by the conclusion is *not* implicit in the premises, for there is no premise stating that no 100-year-old man with arthritis can run a four-minute mile. Of course, we all believe this to be a fact, but there is nothing in the premises that claims this to be a fact; if we *added* a premise to this effect, *then* we would have a deductively valid argument.

The conclusion of an *inductively strong argument*, on the other hand, ventures beyond the factual claims made by the premises. The conclusion asserts more than the premises, since we can describe situations in which the premises would be true and the conclusion false.

If an argument is inductively strong, its conclusion makes factual claims that go beyond the factual information given in the premises.

Thus, an inductively strong argument risks more than a deductively valid one; it risks the possibility of leading from true premises to a false conclusion. But this risk is the price that must be paid for the advantage which inductively strong arguments have over deductively valid ones: the possibility of discovery and prediction of new facts on the basis of old ones.

Definition 4 stated that an argument is inductively strong if and only if it meets two conditions:

- i. It is improbable that its conclusion is false, while its premises are true.
- ii. It is not deductively valid.

Condition (ii) is required because all deductively valid arguments meet condition (i). It is *impossible* for the conclusion of a deductively valid argument to be false while its premises are true, so the probability that the conclusion is false while the premises are true is zero.

Condition (i), however, requires clarification. The “while” in this condition should be read as “given that,” not as “and,” so that the condition can be rephrased as:

- i. It is improbable that its conclusion is false, *given that* its premises are true.

But just what do we mean by “given that?” And why is “It is improbable that its conclusion is false *and* its premises true” an incorrect formulation of condition (i)? What is the difference, in this context, between “and” and “given that”? At this stage these questions are best answered by examining several examples of arguments. The following is an inductively strong argument:

There is intelligent life on Mercury.
There is intelligent life on Venus.

There is intelligent life on Earth.
 There is intelligent life on Jupiter.
 There is intelligent life on Saturn.
 There is intelligent life on Uranus.
 There is intelligent life on Neptune.
 There is intelligent life on Pluto.

There is intelligent life on Mars.

Note that the conclusion is not by itself probable. It is, in fact, probable that the conclusion is false. But it is improbable that the conclusion is false *given that* the premises are true. That is, if the premises were true, then on the basis of that information it would be probable that the conclusion would be true (and thus improbable that it would be false). This is not affected in the least by the fact that some of the premises themselves are quite improbable. Thus, although the conclusion taken by itself is improbable, and some of the premises taken by themselves are also improbable, the conclusion is probable *given the premises*. This example illustrates an important principle:

The type of probability that grades the inductive strength of arguments—we shall call it *inductive probability*—does not depend on the premises alone or on the conclusion alone, but on the *evidential relation* between the premises and the conclusion.

Hopefully we have now gained a certain intuitive understanding of the phrase “given that.” Let us now see why it is incorrect to replace it with “and” and thus incorrect to say that an argument is inductively strong if and only if it is improbable that its conclusion is false *and* its premises are true (and it is not deductively valid). Consider the following argument, which is not inductively strong:

There is a 2000-year-old man in Cleveland.

There is a 2000-year-old man in Cleveland who has three heads.

Now it is quite probable that the conclusion is false *given that* the premise is true. Given that there is a 2000-year-old man in Cleveland, it is quite likely that he has only one head. Thus, the argument is *not* inductively strong. But it is improbable that the conclusion is false *and* the premise is true. For the conclusion to be false and the premise true, there would have to be a non-three-headed 2000-year-old man in Cleveland, and it is quite improbable that there is *any* 2000-year-old man in Cleveland. Thus, it is improbable that the conclusion is false *and* the premise is true, simply because it is improbable that the premise is true.

We now see that the inductive strength of arguments cannot depend on the premises alone. Thus, although it is improbable that the conclusion is false *and* the premises true, it is probable that the conclusion is false *given that* the premises are true and the argument is *not* inductively strong.

An argument might be such that it is improbable that the premises are true *and* the conclusion false, simply because it is improbable that the conclusion is false; that is, it is probable that the conclusion is true. It is important to note that such conditions do not guarantee that the argument is inductively strong. Consider the following example of an argument that has a probable conclusion and yet is *not* inductively strong:

There is a man in Cleveland who is 1999 years and 11-months-old and in good health.

No man will live to be 2000 years old.

Now the conclusion itself is highly probable. Thus, it is improbable that the conclusion is false and consequently improbable that the conclusion is false *and* the premise true. But if the premise were true it would be likely that the conclusion would be false. By itself the conclusion is probable, but given the premise it is not.

The main points of this discussion of inductive strength can be summed up as follows:

1. The inductive probability of an argument is the probability that its conclusion is true given that its premises are true.
2. The inductive probability of an argument is determined by the evidential relation between its premises and its conclusion, not by the likelihood of the truth of its premises alone or the likelihood of the truth of its conclusion alone.
3. An argument is inductively strong if and only if:
 - a. Its inductive probability is high.
 - b. It is not deductively valid.

We defined logic as the study of the strength of the evidential link between the premises and conclusions of arguments. We have seen that there are two different standards against which to evaluate the strength of this link: deductive validity and inductive strength. Corresponding to these two standards are two branches of logic: deductive logic and inductive logic. *Deductive logic* is concerned with tests for deductive validity—that is, rules for deciding whether or not a given argument is deductively valid—and rules for constructing deductively valid arguments. *Inductive logic* is concerned with tests


for measuring the inductive probability, and hence the inductive strength, of arguments and with rules for constructing inductively strong arguments.

Some books appear to suggest that there are two different types of arguments, deductive and inductive, and that deductive logic is concerned with deductive arguments and inductive logic with inductive arguments. That is, they suggest the following classification, together with the assumption that every argument falls in one and only one category:

	Deductive arguments	Inductive arguments
Good	Valid	Strong
Bad	Invalid	Weak

Nothing, however, is further from the truth, for, as we have seen, all inductively strong arguments are deductively invalid.

It is more correct to picture arguments as being arranged on a scale of descending strength, as follows:



Deductive and inductive logic are not distinguished by the different types of arguments with which they deal, but by the different standards against which they evaluate arguments.

Exercises:

Decide whether each of the following arguments is deductively valid, inductively strong, or neither:

1. The meeting place is either the gym or the cafeteria.
 The meeting place is not the gym.

 The meeting place is the cafeteria.

2. A good meal has always made me feel better.

A good meal today will make me feel better.
3. Many great leaders have been crazy.

Everyone who isn't a leader is sane.
4. On all the birthdays I have ever had I have been less than 30 years old.

On my next birthday I will be less than 30 years old.
5. No pigs can fly.
Some horses can fly.

Some horses aren't pigs.

II.5. EPISTEMIC PROBABILITY. We have seen that the concept of inductive probability applies to arguments. The inductive probability of an argument is the probability that its conclusion is true given that its premises are true. Thus, the inductive probability of an argument is a measure of the strength of the evidence that the premises provide for the conclusion. It is correct to speak of the inductive probability of an argument, but incorrect to speak of the inductive probability of statements. Since the premises and conclusion of any argument are statements, it is incorrect to speak of the inductive probability of a premise or of a conclusion.

There is, however, some sense of probability in which it is intuitively acceptable to speak of the probability of a premise or conclusion. When we said that it is improbable that there is a 2000-year-old man in Cleveland, we were relying on some such intuitive sense of probability. There must then be a type of probability, other than inductive probability, that applies to statements rather than arguments.

Let us call this type of probability *epistemic probability* because the Greek stem *episteme* means knowledge, and the epistemic probability of a statement depends on just what our stock of relevant knowledge is. Thus, *the epistemic probability of a statement can vary from person to person and from time to time*, since different people have different stocks of knowledge at the same time and the same person has different stocks of knowledge at different times. *For me*, the epistemic probability that there is a 2000-year-old man now living in Cleveland is quite low, since I have certain background knowledge about the current normal life span of human beings. I feel safe in using this statement as an example of a statement whose epistemic probability is low because I feel safe in assuming that your stock of background knowledge is similar in the relevant respects and that *for you* its epistemic probability is also low.

It is easy to imagine a situation in which the background knowledge of two people would differ in such a way as to generate a difference in the epistemic probability of a given statement. For example, the epistemic probability

that Pegasus will show in the third race may be different for a fan in the grandstand than for Pegasus' jockey, owing to the difference in their knowledge of the relevant factors involved.

It is also easy to see how the epistemic probability of a given statement can change over time for a particular person. The fund of knowledge that each of us possesses is constantly in a state of flux. We are all constantly learning new things directly through experience and indirectly through information which is communicated to us. We are also, unfortunately, continually forgetting things that we once knew. This holds true for societies and cultures as well as for individuals, and human knowledge is continually in a dynamic process of simultaneous growth and decay.

It is important to see how upon the addition of new knowledge to a previous body of knowledge the epistemic probability of a given statement could either increase or decrease. Suppose we are interested in the epistemic probability of the statement that Mr. X is an Armenian and the only relevant information we have is that Mr. X is an Oriental rug dealer in Allentown, Pa., that 90 percent of the Oriental rug dealers in the United States are Armenian, and that Allentown, Pa., is in the United States. On the basis of this stock of relevant knowledge, the epistemic probability of the statement is equal to the inductive probability of the following argument:

Mr. X is an Oriental rug dealer in Allentown, Pa.

Allentown, Pa., is in the United States.

Ninety percent of the Oriental rug dealers in the United States are Armenian.

Mr. X is an Armenian.

The inductive probability of this argument is quite high. If we are now given the new information that although 90 percent of the Oriental rug dealers in the United States are Armenian, only 2 percent of the Oriental rug dealers in Allentown, Pa., are Armenian, while 98 percent are Syrian, the epistemic probability that Mr. X is Armenian decreases drastically, for it is now equal to the inductive probability of the following argument:

Mr. X is an Oriental rug dealer in Allentown, Pa.

Allentown, Pa., is in the United States.

Ninety percent of the Oriental rug dealers in the United States are Armenian.

Ninety-eight percent of the Oriental rug dealers in Allentown, Pa., are Syrian.

Two percent of the Oriental rug dealers in Allentown, Pa., are Armenian.

Mr. X is an Armenian.

The inductive probability of this argument is quite low. Note that the decrease in the epistemic probability of the statement “Mr. X is an Armenian” results not from a change in the inductive probability of a given argument but from the fact that, upon the addition of new information, a *different* argument with more premises becomes relevant in assessing its epistemic probability.

Suppose now we are given still more information, to the effect that Mr. X is a member of the Armenian Club of Allentown and that 99 percent of the members of the Armenian Club are actually Armenians. Upon addition of this information the epistemic probability that Mr. X is an Armenian again becomes quite high, for it is now equal to the inductive probability of the following argument:

Mr. X is an Oriental rug dealer in Allentown, Pa.

Allentown, Pa., is in the United States.

Ninety percent of the Oriental rug dealers in the United States are Armenian.

Ninety-eight percent of the Oriental rug dealers in Allentown, Pa., are Syrian.

Two percent of the Oriental rug dealers in Allentown, Pa., are Armenian.

Mr. X is a member of the Armenian Club of Allentown, Pa.

Ninety-nine percent of the members of the Armenian Club are Armenian.

Mr. X is an Armenian.

Notice once more that the epistemic probability of the statement changes because, with the addition of new knowledge, it became equal to the inductive probability of a new argument with additional premises.

Epistemic probabilities are *important* to us. They are the probabilities upon which we base our decisions. From a stock of knowledge we will arrive at the associated epistemic probability of a statement by the application of inductive logic. Exactly how inductive logic gets us epistemic probabilities from a stock of knowledge depends on how we characterize a stock of knowledge. Just what knowledge is; how we get it; what it is like once we have it; these are deep questions. At this stage, we will work within a simplified model of knowing—the Certainty Model.

The Certainty Model: Suppose that our knowledge originates in observation; that observation makes particular sentences (observation reports) certain and that the probability of other sentences is attributable to the certainty of these. In such a situation we can identify our stock of knowledge with a *list*

of sentences, those observation reports that have been rendered certain by observational experience. It is then natural to evaluate the probability of a statement by looking at an argument with all our stock of knowledge as premises and the statement in question as the conclusion. The inductive strength of that argument will determine the probability of the statement in question. In the certainty model, the relation between epistemic probability and inductive probability is quite simple:

Definition 5: The *epistemic* probability of a statement is the *inductive* probability of that argument which has the statement in question as its conclusion and whose premises consist of all of the observation reports which comprise our stock of knowledge.

The Certainty Model lives up to its name by assigning epistemic probability of one to each observation report in our stock of knowledge.

The certainty of observation reports may be something of an idealization. But it is a useful idealization, and we will adopt it for the present. Later in the course we will discuss some other models of observation.

Exercise

1. Construct several new examples in which the epistemic probability of a statement is increased or decreased by the addition of new information to a previous stock of knowledge.

II.6. PROBABILITY AND THE PROBLEMS OF INDUCTIVE LOGIC. Deductive logic, at least in its basic branches, is well-developed. The definitions of its basic concepts are precise, its rules are rigorously formulated, and the interrelations between the two are well understood. Such is not the case, however, with inductive logic. There are no universally accepted rules for constructing inductively strong arguments; no general agreement on a way of measuring the inductive strength of arguments; no precise, uncontroversial definition of inductive probability. Thus, inductive logic cannot be learned in the sense in which one learns algebra or the basic branches of deductive logic. This is not to say that inductive logicians are wallowing in a sea of total ignorance; many things are known about inductive logic, but many problems still remain to be solved. We shall try to get an idea of just what the problems are, as well as what progress has been made toward their solution.

Some of the main problems of inductive logic can be framed in terms of the concept of inductive probability. I said that there is no precise, uncontroversial definition of inductive probability. I did give a definition of inductive probability. Was it controversial? I think not, but, if you will remember, it

was imprecise. I said that the inductive probability of an argument is the probability that its conclusion is true, given that its premises are true. But at that point I could not give an exact definition of "the probability that an argument's conclusion is true, given that its premises are true." I was, instead, reduced to giving examples so that you could get an intuitive feeling for the meaning of this phrase. The logician, however, is not satisfied with an intuitive feeling for the meaning of key words and phrases. He wishes to analyze the concepts involved and arrive at precise, unambiguous definitions. Thus, one of the problems of inductive logic which remains outstanding is, what exactly is inductive probability?

This problem is intimately connected with two other problems: How is the inductive probability of an argument measured? And, what are the rules for constructing inductively strong arguments? Obviously we cannot develop an exact measure of inductive probability if we do not know precisely what it is. And before we can devise rules for constructing inductively strong arguments, we must have ways of telling which arguments measure up to the required degree of inductive strength. Thus, the solution to the problem of providing a precise definition of inductive probability determines what solutions are available for the problems of determining the inductive probabilities of arguments and constructing systematic rules for generating inductively strong arguments.

Let us call a precise definition of inductive probability, together with the associated method of determining the inductive probability of arguments and rules for constructing inductively strong arguments, an *inductive logic*. Thus, different definitions of inductive probability give rise to different inductive logics. Now we are not interested in finding just any system of inductive logic. We want a system that accords well with common sense and scientific practice. We want a system that gives the result that most of the cases that we would intuitively classify as inductively strong arguments do indeed have a high inductive probability. We want a system that accords with scientific practice and common sense, but that is more precise, more clearly formulated, and more rigorous than they are; a system that codifies, explains, and refines our intuitive judgments. We shall call such a system of inductive logic a *scientific inductive logic*. The problem that we have been discussing can now be reformulated as *the problem of constructing a scientific inductive logic*.

The second major problem of inductive logic, and the one that has been more widely discussed in the history of philosophy, is *the problem of rationally justifying the use of a system of scientific inductive logic* rather than some other system of inductive logic. After all, there are many different possible inductive logics. Some might give the result that arguments that we think are inductively strong are, in fact, inductively weak, and arguments that we think

inductively weak are, in fact, inductively strong. That is, there are possible inductive logics which are diametrically opposed to scientific inductive logic, which are in total disagreement with scientific practice and common sense. Why should we not employ one of these systems rather than scientific induction?

Any adequate answer to this question must take into account the uses to which we put inductive logic (or, at present, the vague intuitions we use in place of a precise system of inductive logic). One of the most important uses of inductive logic is to frame our expectations of the future on the basis of our knowledge of the past and present. We must use our knowledge of the past and present as a guide to our expectations of the future; it is the only guide we have. But it is impossible to have a *deductively valid* argument whose premises contain factual information solely about the past and present and whose conclusion makes factual claims about the future. For the conclusion of a deductively valid argument makes no factual claim that is not already made by the premises. Thus, the gap separating the past and present from the future cannot be bridged in this way by deductively valid arguments, and if the arguments we use to bridge that gap are to have any strength whatsoever they must be inductively strong.

Let us look a little more closely, then, at the way in which inductive logic would be used to frame our expectations of the future. Suppose our plans depend critically on whether it will rain tomorrow. Then the reasonable thing to do, before we decide what course of action to take, is to ascertain the epistemic probability of the statement. "It will rain tomorrow." This we do by putting all the relevant information we now have into the premises of an argument whose conclusion is "It will rain tomorrow" and ascertaining the inductive probability of that argument. If the probability is high, we will have a strong expectation of rain and will make our plans on that basis. If the probability is near zero, we will be reasonably sure that it will not rain and act accordingly.

Now although it is doubtful that anyone carries out the formal process outlined above when he plans for the future, it is hard to deny that, if we were to make our reasoning explicit, it would fall into this pattern. Thus, the making of rational decisions is dependent, via the concept of epistemic probability, on our inductive logic. The second main problem of inductive logic, then, leads us to the following question: How can we rationally justify the use of scientific inductive logic, rather than some other inductive logic, as an instrument for shaping our expectations of the future?

The two main problems of inductive logic are:

1. The construction of a system of scientific inductive logic.
2. The rational justification of the use of that system rather than some other system of inductive logic.

It would seem that the first problem must be solved before the second, since we can hardly justify the use of a system of inductive logic before we know what it is. Nevertheless, I shall discuss the second problem first. It makes sense to do this because we can see why the second problem is *such* a problem without having to know all the details of scientific inductive logic. Furthermore, philosophers historically came to appreciate the difficulty of the second problem much earlier than they realized the full force of the first problem. This second problem, the traditional problem of induction, is discussed in the next chapter.

III

The Traditional Problem of Induction

III.1. INTRODUCTION. In Chapter II we saw that inductive logic is used to shape our expectations of that which is as yet unknown on the basis of those facts that are already known; for instance, to shape our expectations of the future on the basis of our knowledge of the past and present. Our problem is the rational justification of the use of a system of scientific inductive logic, rather than some other system of inductive logic, for this task.

The Scottish philosopher David Hume first raised this problem, which we shall call the *traditional problem of induction*, in full force. Hume gave the problem a cutting edge, for he advanced arguments designed to show that no such rational justification of inductive logic is possible, no matter what the details of a system of scientific inductive logic turn out to be. The history of philosophical discussion of inductive logic since Hume has been in large measure occupied with attempts to circumvent the difficulties he raised. This chapter examines these difficulties and the various attempts to overcome them.

II.2. HUME'S ARGUMENT. Before we can meaningfully discuss arguments which purport to show that it is impossible to rationally justify scientific induction, we must be clear on what would be required to rationally justify a system of inductive logic. Presumably we could rationally justify such a system if we could show that it is well-suited for the uses to which it is put. One of the most important uses of inductive logic is in setting up our predictions of the future. Inductive logic figures in these predictions by way of *epistemic probabilities*. If a claim about the future has high epistemic probability, we predict that it will prove true. And, more generally, we expect something more or less strongly as its epistemic probability is higher or lower. The epistemic probability of a statement is just the inductive probability of the argument which embodies all available information in its premises. Thus, the epistemic probability of a statement depends on two things: (i) the stock of knowledge, and (ii) the inductive logic used to grade the strength of the argument from that stock of knowledge to the conclusion.

Now obviously what we want is for our predictions to be correct. If we could get by with deductively valid arguments we could be assured of true predictions all the time. Deductively valid arguments lead from true premises always to

true conclusions and the statements comprising our stock of knowledge are known to be true. But deductively valid arguments are too conservative to leap from the past and present to the future. For this sort of daring behavior we will have to rely on inductively strong arguments—and we will have to give up the comfortable assurance that we will be right all the time.

How about most of the time? Let us call the sort of argument used to set up an epistemic probability an *e*-argument. That is, an *e*-argument is an argument which has, as its premises, some stock of knowledge. We might hope, then, that inductively strong *e*-arguments will give us *true conclusions most of the time*. Remember that there are *degrees* of inductive strength and that, on the basis of our present knowledge, we do not always simply predict or not-predict that an event will occur, but anticipate it with various *degrees of confidence*. We might hope further that inductively *stronger* *e*-arguments have true conclusions *more often* than inductively *weaker* ones. Finally, since we think that it is useful to gather evidence to enlarge our stock of knowledge, we might hope that inductively strong *e*-arguments give us true conclusions more often when the stock of knowledge embodied in the premises is great than when it is small.

The last consideration really has to do with justifying epistemic probabilities as tools for prediction. The epistemic probability is the inductive probability of an argument embodying *all* our stock of knowledge in its premises. The requirement that it embody *all* our knowledge, and not just some part of it, is known as the Total Evidence Condition. If we could show that basing our predictions on more knowledge gives us better success ratios, we would have justified the total evidence condition.

The other considerations have to do with justifying the other determinant of epistemic probability—the inductive logic which assigns inductive probabilities to arguments.

We are now ready to suggest what is required to rationally justify a system of inductive logic:

Rational Justification

Suggestion I: A system of inductive logic is rationally justified if and only if it is shown that the arguments to which it assigns high inductive probability yield true conclusions from true premises most of the time, and the *e*-arguments to which it assigns higher inductive probability yield true conclusions from true premises more often than the arguments to which it assigns lower inductive probability.

It is this sense of rational justification, or something quite close to it, that Hume has in mind when he advances his arguments to prove that a rational justification of scientific induction is impossible.

If scientific induction is to be rationally justified in the sense of Suggestion I, we must establish that the arguments to which it assigns high inductive probability yield true conclusions from true premises most of the time. By what sort of reasoning, asks Hume, could we establish such a conclusion? If the argument that we must use is to have any force whatsoever, it must be either deductively valid or inductively strong. Hume proceeds to show that neither sort of argument could do the job.

Suppose we try to rationally justify scientific inductive logic by means of a deductively valid argument. The only premises we are entitled to use in this argument are those that state things we know. Since we do not know what the future will be like (if we did, we would have no need of an inductive logic on which to base our predictions), the premises can contain knowledge of only the past and present. But if the argument is deductively valid, then the conclusion can make no factual claims that are not already made by the premises. Thus, the conclusion of the argument can only refer to the past and present, not to the future, for the premises made no factual claims about the future. Such a conclusion cannot, however, be adequate to rationally justify scientific induction.

To rationally justify scientific induction we must show that e-arguments to which it assigns high inductive probability yield true conclusions from true premises most of the time. And "most of the time" does not mean most of the time in only the past and present; it means most of the time, *past, present, and future*. It is conceivable that a certain type of argument might have given us true conclusions from true premises in the past and might cease to do so in the future. Since our conclusion cannot tell us how successful arguments will be in the future, it cannot establish that the e-arguments to which scientific induction assigns high probability will give us true conclusions from true premises *most of the time*. Thus, we cannot use a deductively valid argument to rationally justify induction.

Suppose we try to rationally justify scientific induction by means of an inductively strong argument. We construct our argument, whatever it may be, and present it as an inductively strong argument. "Why do you think that this is an inductively strong argument?" Hume might ask. "Because it has a high inductive probability," we would reply. "And what system of inductive logic assigns it a high probability?" "Scientific induction, of course." What Hume has pointed out is that if we attempt to rationally justify scientific induction by use of an inductively strong argument, we are in the position of having to *assume* that scientific induction is reliable in order to prove that scientific induction is reliable; we are reduced to begging the question. Thus, we cannot use an inductively strong argument to rationally justify scientific induction.

A common argument is that scientific induction is justified because it has been quite successful in the past. On reflection, however, we see that this argument is really an attempt to justify induction by means of an inductively strong argument, and thus begs the question. More explicitly, the argument reads something like this:

Arguments that are judged by scientific inductive logic to have high inductive probability have given us true conclusions from true premises most of the time in the past.

Such arguments will give us true conclusions from true premises most of the time, past, present, and future.

It should be obvious that this argument is not deductively valid. At best it is assigned high inductive probability by a system of scientific inductive logic. But the point at issue is whether we should put our faith in such a system.

We can view the traditional problem of induction from a different perspective by discussing it in terms of the *principle of the uniformity of nature*. Although we do not have the details of a system of scientific induction in hand, we do know that it must accord well with common sense and scientific practice, and we are reasonably familiar with both. A few examples will illustrate a general principle which appears to underlie both scientific and common-sense judgments of inductive strength.

If you were to order steak in a restaurant, and a friend were to object that steak would corrode your stomach and lead to quick and violent death, it would seem quite sufficient to respond that you had often eaten steak without any of the dire consequences he predicted. That is, you would intuitively judge the following argument to be inductively strong:

I have eaten steak many times and it has never corroded my stomach.

Steak will not now corrode my stomach.

Suppose a scientist is asked whether a rocket would work in reaches of space beyond the range of our telescopes. She replies that it would, and to back up her answer appeals to certain principles of theoretical physics. When asked what evidence she has for these principles, she can refer to a great mass of observed phenomena that corroborate them. The scientist is then judging the following argument to be inductively strong:

Principles A, B, and C correctly describe the behavior of material bodies in all of the many situations we have observed.

Principles A, B, and C correctly describe the behavior of material bodies in those reaches of space that we have not as yet observed.

There appears to be a common assumption underlying the judgments that these arguments are inductively strong. As a steak eater you assume that the future will be like the past, that types of food that proved healthful in the past will continue to prove so in the future. The scientist assumes that the distant reaches of space are like the nearer ones, that material bodies obey the same general laws in all areas of space. Thus, it seems that underlying our judgments of inductive strength in both common sense and science is the presupposition that nature is uniform or, as it is sometimes put, that like causes produce like effects throughout all regions of space and time. Thus, we can say that a system of scientific induction will base its judgments of inductive strength on the presupposition that *nature is uniform* (and in particular that the future will resemble the past).

We ought to realize at this point that we have only a vague, intuitive understanding of the principle of the uniformity of nature, gleaned from examples rather than specified by precise definitions. This rough understanding is sufficient for the purposes at hand. But we should bear in mind that the task of giving an *exact* definition of the principle, a definition of the sort that would be presupposed by a system of scientific inductive logic, is as difficult as the construction of such a system itself. One of the problems is that nature is simply not uniform in all respects, the future does not resemble the past in all respects. Bertrand Russell once speculated that the chicken on slaughter-day might reason that whenever the humans came it had been fed, so when the humans would come today it would also be fed. The chicken thought that the future would resemble the past, but it was dead wrong.

The future may resemble the past, but it does not do so in all respects. And we do not know beforehand what those respects are nor to what degree the future resembles the past. Our ignorance of what these respects are is a deep reason behind the total evidence condition. Looking at more and more evidence helps us reject spurious patterns which we might otherwise project into the future. Trying to say exactly *what* about nature we believe is uniform thus turns out to be a surprisingly delicate task.

But suppose that a subtle and sophisticated version of the principle of the uniformity of nature can be formulated which adequately explains the judgments of inductive strength rendered by scientific inductive logic. Then if nature is indeed uniform in the required sense (past, present, and future), e-arguments judged strong by scientific induction will indeed give us true conclusions most of the time. Therefore, the problem of rationally justifying scientific induction could be reduced to the problem of establishing that nature is uniform.

But by what reasoning could we establish such a conclusion? If an argument is to have any force whatsoever it must be either deductively valid or

inductively strong. A deductively valid argument could not be adequate, for if the information in the premises consists solely of our knowledge of the past and present, then the conclusion cannot tell us that nature will be uniform in the future. The conclusion of a deductively valid argument can make no factual claims that are not already made by the premises, and factual claims about the future are not factual claims about the past and present. But if we claim to have established the principle of the uniformity of nature by an argument that is rated inductively strong by scientific inductive logic, we are open to a challenge as to why we should place our faith in such arguments. But we cannot reply "Because nature is uniform," for that is precisely what we are trying to establish.

Let us summarize the traditional problem of induction. It appears that to rationally justify a system of scientific inductive logic we would have to establish that the e-arguments it judges to be inductively strong give us true conclusions most of the time. If we try to prove that this is the case by means of a deductively valid argument whose premises state things we already know, then the conclusion must fall short of the desired goal. But to try to rationally justify scientific induction by means of an argument that scientific induction judges to be inductively strong is to beg the question. The same difficulties arise if we attempt to justify scientific inductive logic by establishing that nature is uniform.

III.3. THE INDUCTIVE JUSTIFICATION OF INDUCTION.

Hume has presented us with a dilemma. If we try to justify scientific inductive logic by means of a deductively valid argument with premises known to be true, our conclusion will be too weak. If we try to use an inductively strong argument, we are reduced to begging the question. The proponent of the inductive justification of induction tackles the second horn of the dilemma. He maintains that we can justify scientific induction by an inductively strong argument without begging the question. Although his attempt is not altogether successful, there is a great deal to be learned from it.

The answer to the question "Why should we believe that scientific induction is a reliable guide for our expectations?" that immediately occurs to everyone is "Because it has worked well so far." Hume's objection to this answer was that it begs the question, that it assumes scientific induction is reliable in order to prove that scientific induction is reliable. The proponents of the inductive justification of induction, however, claim that the answer only *appears* to beg the question, because of a mistaken conception of scientific induction. They claim that if we properly distinguish *levels* of scientific induction, rather than lumping all arguments that scientific induction judges to be strong in one

category, we will see that the inductive justification of induction does not beg the question.

Just what then are these levels of scientific induction? And what is their relevance to the inductive justification of induction? We can distinguish different *levels of argument*, in terms of the things they talk about. Arguments on level 1 will talk about individual things or events; for instance:

Many jub-jub birds have been observed, and they have all been purple.

The next jub-jub bird to be observed will be purple.

Level 1 of scientific inductive logic would consist of rules for assigning inductive probabilities to arguments of level 1. Presumably the rules of level 1 of scientific induction would assign high inductive probability to the preceding argument. Arguments on level 2 will talk about arguments on level 1; for instance:

Some deductively valid arguments on level 1 have true premises.

All deductively valid arguments on level 1 which have true premises have true conclusions.

Some deductively valid arguments on level 1 have true conclusions.

This is a deductively valid argument on level 2 which talks about deductively valid arguments on level 1. The following is also an argument on level 2 which talks about arguments on level 1:

Some arguments on level 1 which the rules of level 1 of scientific inductive logic say are inductively strong have true premises.

The denial of a true statement is a false statement.

Some arguments on level 1 which the rules of level 1 of scientific inductive logic say are inductively strong have premises whose denial is false.

This is a deductively valid argument on level 2 that talks about arguments on level 1, which the rules of level 1 of scientific inductive logic classify as inductively strong.

There are, of course, arguments on level 2 that are not deductively valid, and there is a corresponding second level of scientific inductive logic which consists of rules that assign degrees of inductive strength to *these* arguments. There are arguments on level 3 that talk about arguments on level 2, arguments on level 4 that talk about arguments on level 3, and so on. For each level of argument, scientific inductive logic has a corresponding level of rules.

This characterization of the levels of argument, and the corresponding levels of scientific induction, is summarized in Table III.1.

Table III.1

Levels of argument	Levels of scientific inductive logic
k : Arguments about arguments on level $k-1$.	k : Rules for assigning inductive probabilities to arguments on level k .
⋮	⋮
2: Arguments about arguments on level 1.	2: Rules for assigning inductive probabilities to arguments on level 2.
1: Arguments about individuals.	1: Rules for assigning inductive probabilities to arguments on level 1.

As the table shows, scientific inductive logic is seen not as a simple, homogeneous system but rather as a complex structure composed of an infinite number of strata of distinct sets of rules. The sets of rules on different levels are not, however, totally unrelated. The rules on each level presuppose, in some sense, that nature is uniform and that the future will resemble the past. If this were not the case, we would have no reason for calling the whole system of levels a system of *scientific* inductive logic.

We are now in a position to see how the system of levels of scientific induction is to be employed in the inductive justification of induction. In answer to the question, "Why should we place our faith in the rules of level 1 of scientific inductive logic?" the proponent of the inductive justification of induction will advance an argument on level 2:

Among arguments used to make predictions in the past, e-arguments on level 1 (which according to level 1 of scientific inductive logic are inductively strong) have given true conclusions most of the time.

With regard to the next prediction, an e-argument judged inductively strong by the rules of scientific inductive logic will yield a true conclusion.

The proponent will maintain that the premise of this argument is true, and if we ask why he thinks that this is an inductively strong argument, he will reply that *the rules of level 2 of scientific inductive logic* assign it a high inductive probability. If we now ask why we should put our faith in *these* rules, he will advance a similar argument on level 3, justify that argument by appeal to the rules of scientific inductive logic on level 3, justify those rules by an argument on level 4, and so on.

The inductive justification of induction is summarized in Table III.2. The arrows in the table show the order of justification. Thus, the rules of level 1 are justified by an argument on level 2, which is justified by the rules on level 2, which are justified by an argument on level 3, and so on.

Let us now see how it is that the proponent of the inductive justification of induction can plead not guilty to Hume's charge of begging the question; that is, of presupposing exactly what one is trying to prove. In justifying the rules of level 1, the proponent of the inductive justification of induction does not presuppose that *these* rules will work the next time; in fact, he advances an argument (on level 2) to show that they will work next time. Now it is true that the use of this argument presupposes that the rules of level 2 will work next time. But there is another argument waiting on level 3 to show that the rules of level 2 will work. The use of that argument does not presuppose what it is trying to establish; it presupposes that the rules on level 3 will work. Thus, none of the arguments used in the inductive justification of induction presuppose what they are trying to prove, and the inductive justification of induction does not technically beg the question.

Table III.2

Levels of argument	Levels of scientific inductive logic
⋮	⋮
3: Rules of level 2 of scientific inductive logic have worked well in the past. <hr/> They will work well next time.	3: Rules for assigning inductive probabilities to arguments on level 3.
2: Rules of level 1 of scientific inductive logic have worked well in the past.* <hr/> They will work well next time.	2: Rules for assigning inductive probabilities to arguments on level 2.
1:	1: Rules for assigning inductive probabilities to arguments on level 1.

* The statement "rules of level 1 of scientific inductive logic have worked well in the past," is to be taken as shorthand for "arguments on level 1, which according to the rules of level 1 of scientific inductive logic are inductively strong and which have been used to make predictions in the past, have given us true conclusions, when the premises were true, most of the time." Thus, the argument on level 2 used to justify the rules of level 1 is exactly the same one as put forth in the second paragraph on page 33.

Perhaps how these levels work can be made clearer by looking at a simple example. Suppose our only observations of the world have been of 100 jub-jub birds and they have all been purple. After observing 99 jub-jub birds, we advanced argument jj-99:

We have seen 99 jub-jub birds and they were all purple.

The next jub-jub bird we see will be purple.

This argument was given high inductive probability by rules of level 1 of scientific inductive logic. We knew its premises to be true, and we took its conclusion as a prediction. The 100th jub-jub bird can thus be correctly described as purple—or as the color that makes the conclusion of argument jj-99 true—or as the color that results in a successful prediction by the rules of level 1 of scientific inductive logic. Let us also suppose that similar arguments had been advanced in the past: jj-98, jj-97, etc. Each of these arguments was an e-argument to which the rules of level 1 assigned high inductive probability. Thus, the observations of jub-jub birds 98 and 99, etc., are also observations of successful outcomes to predictions based on assignments of probabilities to e-arguments by rules of level 1. This gives rise to an argument on level 2:

e-arguments on level 1, which are assigned high inductive probability by rules of level 1, have had their conclusions predicted 98 times and all those predictions were successful.

Predicting the conclusion of the next e-argument on level 1 which is assigned high inductive probability will also lead to success.

This argument is assigned high inductive probability by rules of level 2. If the next jub-jub bird to be observed is purple, it makes this level 2 argument successful in addition to making the appropriate level 1 argument successful. A string of such successes gives rise to a similar argument on level 3 and so on, up the ladder, as indicated in Table III.2.

If someone were to object that what is wanted is a justification of scientific induction as a whole and that this has not been given, the proponent of the inductive justification of induction would reply that for every level of rules of scientific inductive logic he has a justification (on a higher level), and that certainly if every level of rules is justified, then the whole system is justified. He would maintain that it makes no sense to ask for a justification for the system *over and above* a justification for each of its parts. This position, it must be admitted, has a good deal of plausibility; a final evaluation of its merits, however, must await some further developments.

The position held by the proponent of the inductive justification of induction contrasts with the position held by Hume in that it sets different

requirements for the rational justification of a system of inductive logic. The following is implicit in the inductive justification of induction:

Rational Justification

Suggestion II: A system of inductive logic is rationally justified if, for every level (k) of rules of that system, there is an e-argument on the next highest level ($k + 1$) which:

- i. Is judged inductively strong by its own system's rules (these will be rules of level $k + 1$).
- ii. Has as its conclusion the statement that the system's rules on the original level (k) will work well next time.

It is important to see that *whether a system of induction meets these conditions depends not only on the system of induction itself but also on the facts, on the way that the world is.* We can imagine a situation in which scientific induction would indeed not meet these conditions. Imagine a world which has been so chaotic that scientific induction on level 1 has not worked well; that is, suppose that the e-arguments on level 1, which according to the rules of level 1 of scientific inductive logic are inductively strong and which have been used to make predictions in the past, have given us *false* conclusions from true premises most of the time. In such a situation the inductive justification of induction could not be carried through. For although the argument on level 2 used to justify the rules of level 1 of scientific induction, that is:

Rules of level 1 of scientific inductive logic have worked well
in the past.

They will work well next time.

would still be judged inductively strong by the rules of level 2 of scientific inductive logic, its premise would not be true. Indeed, in the situation under consideration the following argument on level 2 *would* have a premise that was known to be true and would also be judged inductively strong by the rules of level 2 of scientific inductive logic:

Rules of level 1 of scientific inductive logic have not worked well
in the past.

They will not work well next time.

Thus, we can conceive of situations in which level 2 of scientific induction, instead of justifying level 1 of scientific induction, would tell us that level 1 is unreliable.

We are not, in fact, in such a situation. Level 1 of scientific induction has served us quite well, and it is upon this fact that the inductive justification of induction capitalizes. This is indeed an important fact, but it remains to be seen whether it is sufficient to rationally justify a system of scientific inductive logic.

The proponent of the inductive justification of scientific inductive logic has done us a service in distinguishing the various levels of induction. He has also made an important contribution by pointing out that there are possible situations in which the higher levels of scientific induction do not always support the lower levels and that we are, in fact, not in such a situation. But as a justification of the system of scientific induction his reasoning is not totally satisfactory. While he has not technically begged the question, he has come very close to it. Although he has an argument to justify every level of scientific induction, and although none of his arguments presuppose exactly what they are trying to prove, the justification of each level presupposes the correctness of the level above it. Lower levels are justified by higher levels, but always higher levels of scientific induction. No matter how far we go in the justifying process, we are always within the system of scientific induction. Now, isn't this loading the dice? Couldn't someone with a completely different system of inductive logic execute the same maneuver? Couldn't he justify each level of *his* logic by appeal to higher levels of *his* logic? Indeed he could. Given the same factual situation in which the inductive justification of scientific induction is carried out, an entirely different system of inductive logic could also meet the conditions laid down under Rational Justification, Suggestion II. Let us take a closer look at such a contrasting system of inductive logic.

We said that scientific induction assumes that, in some sense, nature is uniform and the future will be like the past. Some such assumption is to be found backing the rules on each level of scientific inductive logic. The assumptions are not exactly the same on each level; they must be different because we can imagine a situation in which scientific induction on level 2 would tell us that scientific induction on level 1 will not work well. Thus, different principles of the uniformity of nature are presupposed on different levels of scientific inductive logic. But although they are not exactly the same, they are similar; they are all principles of the uniformity of nature. Thus, each level of scientific inductive logic presupposes that, in some sense, nature is uniform and the future will be like the past. A system of inductive logic that would be *diametrically opposed* to scientific inductive logic would be one which presupposed on all levels that the future will not be like the past. We shall call this system a system of *counterinductive logic*.

Let us see how counterinductive logic would work on level 1. Scientific inductive logic, which assumes that the future will be like the past, would assign the following argument a high inductive probability:

Many jub-jub birds have been observed and they have all been purple.

The next jub-jub bird to be observed will be purple.

Counterinductive logic, which assumes that the future will *not* be like the past, would assign it a low inductive probability and would instead assign a high inductive probability to the following argument:

Many jub-jub birds have been observed and they have all been purple.

The next jub-jub bird to be observed will not be purple.

In general, counterinductive logic assigns low inductive probabilities to arguments that are assigned high inductive probabilities by scientific inductive logic, and high inductive probabilities to arguments that are assigned low inductive probabilities by scientific inductive logic.

Now suppose that a counterinductivist decided to give an inductive justification of counterinductive logic. The scientific inductivist would justify his rules of level 1 by the following level 2 argument:

Rules of level 1 of scientific induction have worked well in the past.

They will work well next time.

The counterinductivist, on the other hand, would justify his rules of level 1 by another kind of level 2 argument:

Rules of level 1 of counterinductive logic have *not* worked well in the past.

They will work well next time.

By the counterinductivist's rules, this is an inductively strong argument, for on level 2 he also assumes that the future will be unlike the past. Thus, the counterinductivist is not at all bothered by the fact that his level 1 rules have been failures; indeed he takes this as evidence that they will be successful in the future. Granted his argument appears absurd to us, for we are all at heart scientific inductivists. But if the scientific inductivist is allowed to use his own rules on level 2 to justify his rules on level 1, how can we deny the same right to the counterinductivist? If asked to justify his rules on level 2, the counterinductivist will advance a similar argument on level

3, and so on. If an inductive justification of scientific inductive logic can be carried through, then a parallel inductive justification of counterinductive logic can be carried through. Table III.3 summarizes how this would be done.

Table III.3

Level of argument	Justifying arguments of the scientific inductivist	Justifying arguments of the counter-inductivist
⋮	⋮	⋮
3:	Rules of level 2 of scientific inductive logic have worked well in the past.	Rules of level 2 of counter-inductive logic have not worked well in the past.
	They will work well next time.	They will work well next time.
2:	Rules of level 1 of scientific inductive logic worked well in the past.	Rules of level 1 of counter-inductive logic have not worked well in the past.
	They will work well next time.	They will work well next time.

The counterinductivist is, of course, a fictitious character. No one goes through life consistently adhering to the canons of counterinductive logic, although some of us do occasionally slip into counterinductive reasoning. The poor poker player who thinks that his luck is due to change because he has been losing so heavily is a prime example. But aside from a description of gamblers' rationalizations, counterinductive logic has little practical significance.

It does, however, have great theoretical significance. For what we have shown is that if scientific inductive logic meets the conditions laid down under Rational Justification, Suggestion II, so does counterinductive logic. This is sufficient to show that Suggestion II is inadequate as a definition for rational justification. A rational justification of a system of inductive logic must provide reasons for using that system rather than any other. Thus, if two inconsistent systems, scientific induction and counterinduction, can meet the conditions of Suggestion II, then Suggestion II cannot be an adequate definition of rational justification. The arguments examined in this section do show that scientific inductive logic meets the conditions of Suggestion II, but these arguments do not rationally justify scientific induction.

This is not to say that what has been pointed out is not both important and interesting. Let us say that any system of inductive logic that meets the

conditions of Suggestion II is *inductively coherent with the facts*. It may be true that for a system of inductive logic to be rationally justified it must be inductively coherent with the facts; that is, that inductive coherence with the facts may be a necessary condition for rational justification. But the example of the counterinductivist shows conclusively that inductive coherence with the facts is not by itself sufficient to rationally justify a system of inductive logic. Consequently, the inductive justification of scientific inductive logic fails.

We may summarize our discussion of the inductive justification of induction as follows:

1. The proponent of the inductive justification of scientific induction points out that scientific inductive logic is inductively coherent with the facts.
2. He claims that this is sufficient to rationally justify scientific inductive logic.
3. But it is not sufficient since counterinductive logic is also inductively coherent with the facts.
4. Nevertheless it is important and informative since we can imagine circumstances in which scientific inductive logic would not be inductively coherent with the facts.
5. The proponent of the inductive justification of scientific induction has also succeeded in calling to our attention the fact that there are various levels of induction.

Suggested readings

John Stuart Mill, "The Ground of Induction," reprinted in *A Modern Introduction to Philosophy*, Paul Edwards and Arthur Pap, Eds. (New York: The Free Press, 1973), pp. 133–41.

F. L. Will, "Will the Future Be Like the Past?" reprinted in *A Modern Introduction to Philosophy* (rev. ed.), Paul Edwards and Arthur Pap, Eds. (New York: The Free Press, 1973), pp. 148–58.

Max Black, "Inductive Support of Inductive Rules," *Problems of Analysis* (Ithaca, New York: Cornell University Press, 1954), pp. 191–208.

All of these authors are arguing for some type of inductive justification of induction, although none of them holds the exact position outlined in this section, which is a synthesis of several viewpoints.

III.4. THE PRAGMATIC JUSTIFICATION OF INDUCTION.

Remember that the traditional problem of induction can be formulated as a dilemma: If the reasoning we use to rationally justify scientific inductive logic

is to have any strength at all it must be either deductively valid or inductively strong. But if we try to justify scientific inductive logic by means of a deductively valid argument with premises that are known to be true, our conclusion will be too weak. And if we try to use an inductively strong argument, we are reduced to begging the question. Whereas the proponent of the *inductive* justification of scientific induction attempts to go over the second horn of the dilemma, the proponent of the *pragmatic* justification of induction attacks the first horn; he attempts to justify scientific inductive logic by means of a deductively valid argument.

The pragmatic justification of induction was proposed by Herbert Feigl and elaborated by Hans Reichenbach, both founders of the logical empiricist movement. Reichenbach's pragmatic justification of induction is quite complicated, for it depends on what he believes are the details (at least the basic details) of scientific inductive logic. Thus, no one can fully understand Reichenbach's arguments until he has studied Reichenbach's definition of probability and the method he prescribes for discovering probabilities. We shall return to these questions later; at this point we will discuss a simplified version of the pragmatic justification of induction. This version is correct as far as it goes. Only bear in mind that there is more to be learned.

Reichenbach wishes to justify scientific inductive logic by a deductively valid argument. Yet he agrees with Hume that no deductive valid argument with premises that are known to be true can give us the conclusion that scientific induction will give us true conclusions most of the time. He agrees with Hume that the conditions of Rational Justification, Suggestion I, cannot be met. Since he fully intends to rationally justify scientific inductive logic, the only path open to him is to argue that the conditions of Rational Justification, Suggestion I, need not be met in order to justify a system of inductive logic. He proceeds to advance his own suggestion as to what is required for rational justification and to attempt to justify scientific inductive logic in these terms.

If Hume's arguments are correct, there is no way of showing that scientific induction will give us true conclusions from true premises most of the time. But since Hume's arguments apply equally well to any system of inductive logic there is no way of showing that any competing system of inductive logic will give us true conclusions from true premises most of the time either. Thus, scientific inductive logic has the same status as all other systems of inductive logic in this matter. No other system of inductive logic can be demonstrated to be superior to scientific inductive logic in the sense of showing that it gives true conclusions from true premises more often than scientific inductive logic.

Reichenbach claims that although it is impossible to show that any inductive method will be successful, it can be shown that scientific induction will be successful, if any method of induction will be successful. In other words, it is possible that no inductive logic will guide us to e-arguments that give us true

conclusions most of the time, but if any method will then scientific inductive logic will also. If this can be shown, then it would seem fair to say that scientific induction has been rationally justified. After all, we must make some sort of judgments, conscious or unconscious, as to the inductive strength of arguments if we are to live at all. We must base our decisions on our expectations of the future, and we base our expectations of the future on our knowledge of the past and present. We are all gamblers, with the stakes being the success or failure of our plans of action. Life is an exploration of the unknown, and every human action presumes a wager with nature.

But if our decisions are a gamble and if no method is guaranteed to be successful, then it would seem rational to bet on that method which will be successful, if any method will. Suppose that you were forcibly taken into a locked room and told that whether or not you will be allowed to live depends on whether you win or lose a wager. The object of the wager is a box with red, blue, yellow, and orange lights on it. You know nothing about the construction of the box but are told that either all of the lights, some of them, or none of them will come on. You are to bet on one of the colors. If the colored light you choose comes on, you live; if not, you die. But before you make your choice you are also told that neither the blue, nor the yellow, nor the orange light can come on without the red light also coming on. If this is the only information you have, then you will surely bet on red. For although you have no guarantee that your bet on red will be successful (after all, all the lights might remain dark) you know that if any bet will be successful, a bet on red will be successful. Reichenbach claims that scientific inductive logic is in the same privileged position vis-à-vis other systems of inductive logic as is the red light vis-à-vis the other lights.

This leads us to a new proposal as to what is required to rationally justify a system of inductive logic:

Rational Justification

Suggestion III: A system of inductive logic is rationally justified if we can show that the e-arguments that it judges inductively strong will give us true conclusions most of the time, if e-arguments judged inductively strong by any method will.

Reichenbach attempts to show that scientific inductive logic meets the conditions of Rational Justification, Suggestion III, by a deductively valid argument. The argument goes roughly like this:

Either nature is uniform or it is not.

If nature is uniform, scientific induction will be successful.

If nature is not uniform, then no method will be successful.

If any method of induction will be successful, then scientific induction will be successful.

There is no question that this argument is deductively valid, and the first and second premises are surely known to be true. But how do we know that the third premise is true? Couldn't there be some strange inductive method that would be successful even if nature were not uniform? How do we know that for any method to be successful nature must be uniform?

Reichenbach has a response ready for this challenge. Suppose that in a completely chaotic universe, some method, call it method X, were successful. Then there is still at least one outstanding uniformity in nature: the uniformity of method X's success. And scientific induction would discover *that* uniformity. That is, if method X is successful on the whole, if it gives us true predictions most of the time, then sooner or later the statement "Method X has been reliable in the past" will be true, and the following argument would be judged inductively strong by scientific inductive logic:

Method X has been reliable in the past.

Method X will be reliable in the future.

Thus, if method X is successful, scientific induction will also be successful in that it will discover method X's reliability, and, so to speak, license method X as a subsidiary method of prediction. This completes the proof that scientific induction will be successful if any method will.

The job may appear to be done, but in fact there is a great deal more to be said. In order to analyze just what has been proved and what has not, we shall use the idea of levels of inductive logic, which was developed in the last section. When we talk about a method, we are really talking about a system of inductive logic, while glossing over the fact that a system of inductive logic is composed of distinct levels of rules. Let us now pay attention to this fact. Since a system of inductive logic is composed of distinct levels of rules, in order to justify that system we would have to justify each level of its rules. Thus, to justify scientific inductive logic we would have to justify level 1 rules of scientific inductive logic, level 2 rules of scientific inductive logic, level 3 rules of scientific inductive logic, and so on. If each of these levels of rules is to be justified in accordance with the principle "It is rational to rely on a method that is successful if any method is successful," then the pragmatic justification of induction must establish the following:

- 1: Level 1 rules of scientific induction will be successful if level 1 rules of any system of inductive logic will be successful.

- 2: Level 2 rules of scientific induction will be successful if level 2 rules of any system of inductive logic will be successful.
- ⋮
- k : Level k rules of scientific induction will be successful if level k rules of any system of inductive logic will be successful.

But if we look closely at the pragmatic justification of induction, we see that it does not establish this but rather something quite different.

Suppose that system X of inductive logic is successful on level 1. That is, the arguments that it judges to be inductively strong give us true conclusions from true premises most of the time. Then sooner or later an argument on level 2 which is judged inductively strong by scientific inductive logic, that is:

Rules of level 1 of system X have been reliable in the past.

Rules of level 1 of system X will be reliable in the future.

will come to have a premise that is known to be true. If the rules on level 1 of system X give true predictions most of the time, then sooner or later it will be true that they have given us true predictions most of the time *in the past*. And once we have this premise, scientific induction on level 2 leads us to the conclusion that they will be reliable in the future.

Thus, what has been shown is that if any system of inductive logic has successful rules on level 1, then scientific induction provides a justifying argument for these rules on level 2. Indeed, we can generalize this principle and say that if a system of inductive logic has successful rules on a given level, then scientific induction provides a justifying argument on the next highest level. More precisely, the pragmatist has demonstrated the following: If system X of inductive logic has rules on level k which pick out, as inductively strong arguments of level k , those which give true predictions most of the time, then there is an argument on level $k + 1$, which is judged inductively strong by the rules of level $k + 1$ of scientific inductive logic, which has as its conclusion the statement that the rules of system X on level k are reliable, and which has a premise that will sooner or later be known to be true.

Now this is quite different from showing that if any method works on any level then scientific induction will also work on *that* level, or even from showing that if any method works on level 1 then scientific induction will work on level 1. Instead what has been shown is that if any other method is generally successful on level 1 then scientific induction will have at least one notable success on level 2: it will eventually predict the continued success of that other method on level 1.

Although this is an interesting and important conclusion, it is not sufficient for the task at hand. Suppose we wish to choose a set of rules for level 1. In order to be in a position analogous to the wager about the box with the colored

lights, we would have to know that scientific induction would be successful on level 1 if any method were successful on level 1. But we do not know this. For all we know, scientific induction might fail on level 1 and another method might be quite successful. If this were the case, scientific induction on level 2 would eventually tell us so, but this is quite a different matter.

In summary, the attempt at a pragmatic justification of induction has made us realize that a deductively valid justification of scientific induction would be acceptable if it could establish that: if any system of inductive logic has successful rules on a given level, then scientific inductive logic will have successful rules on that level. But the arguments advanced in the pragmatic justification fail to establish this conclusion. Instead, they show that if any system of inductive logic has successful rules on a given level, then scientific inductive logic will license a justifying argument for those rules on the next higher level.

Both the attempt at a pragmatic justification and the attempt at an inductive justification have failed to provide an absolute justification of scientific induction. Nevertheless, both of them have brought forth useful facts. For instance, the pragmatic justification of induction shows one clear advantage of scientific induction over counterinduction. The counterinductivist cannot prove that if any method is successful on level 1, counterinduction on level 2 will eventually predict its continued success. In fact some care is required to even give a logically consistent formulation of counterinduction as a general policy.

It seems, then, that there is still room for constructive thought on the problem, and that we can learn much from previous attempts to solve it.

Suggested reading

Hans Reichenbach, *Experience and Prediction: An Analysis of the Foundations and the Structure of Knowledge*. (Chicago: University of Chicago Press, 1938).

III.5. SUMMARY. We have developed the traditional problem of induction and discussed several answers to it. We found that each position we discussed had a different set of standards for rational justification of a system of inductive logic.

- I. *Position*: The original presentation of the traditional problem of induction.
Standard for Rational Justification: A system of inductive logic is rationally justified if and only if it is shown that the e-arguments that it judges inductively strong yield true conclusions most of the time.
- II. *Position*: The inductive justification of induction.
Standard for Rational Justification: A system of inductive logic is

rationally justified if for every level (k) of rules of that system there is an e-argument on the next highest level ($k + 1$) which:

- i. Is judged inductively strong by its own system's rules.
- ii. Has as its conclusion the statement that the system's rules on the original level (k) will work well next time.

III. *Position*: The pragmatic justification of induction.

Standard for Rational Justification: A system of inductive logic is rationally justified if it is shown that the e-arguments that it judges inductively strong yield true conclusions most of the time, if e-arguments judged inductively strong by any method will.

The attempt at an inductive justification of scientific inductive logic taught us to recognize different levels of arguments and corresponding levels of inductive rules. It also showed that scientific inductive logic meets the standards for Rational Justification, Suggestion II. However, we saw that Suggestion II is really not a sense of rational justification at all, for both scientific inductive logic and counterinductive logic can meet its conditions. Thus, it cannot justify the choice of one over the other.

The attempt at a pragmatic justification of scientific inductive logic showed us that Suggestion III, properly interpreted in terms of levels of induction, would be an acceptable sense of rational justification, although it would be a weaker sense than that proposed in Suggestion I. However, the pragmatic justification fails to demonstrate that scientific induction meets the conditions of Suggestion III.

It seems that we cannot make more progress in justifying inductive logic until we make some progress in saying exactly what scientific inductive logic is. The puzzles to be discussed in the next chapter show that we have to be careful in specifying the nature of scientific inductive logic.

IV

The Goodman Paradox and The New Riddle of Induction

IV.1. INTRODUCTION. In Chapter III we presented some general specifications for a system of scientific inductive logic. We said it should be a system of rules for assigning inductive probabilities to arguments, with different levels of rules corresponding to the different levels of arguments. This system must accord fairly well with common sense and scientific practice. It must on each level presuppose, in some sense, that nature is uniform and that the future will resemble the past. These general specifications were sufficient to give us a foundation for surveying the traditional problem of induction and the major attempts to solve or dissolve it.

However, to be able to apply scientific inductive logic, as a rigorous discipline, we must know precisely what its rules are. Unfortunately no one has yet produced an adequate formulation of the rules of scientific inductive logic. In fact, inductive logic is in much the same state as deductive logic was before Aristotle. This unhappy state of affairs is not due to a scarcity of brainpower in the field of inductive logic. Some of the great minds of history have attacked its problems. The distance by which they have fallen short of their goals is a measure of the difficulty of the subject. Formulating the rules of inductive logic, in fact, appears to be a more difficult enterprise than doing the same for deductive logic. Deductive logic is a “yes or no” affair; an argument is either deductively valid or it is not. But inductive strength is a matter of degree. Thus, while deductive logic must *classify* arguments as valid or not, inductive logic must *measure* the inductive strength of arguments.

Setting up such rules of measurement is not an easy task. It is in fact beset with so many problems that some philosophers have been convinced it is impossible. They maintain that a system of scientific induction cannot be constructed; that prediction of the future is an art, not a science; and that we must rely on the intuitions of experts, rather than on scientific inductive logic, to predict the future. We can only hope that this gloomy doctrine is as mistaken as the view of those early Greeks who believed deductive logic could never be reduced to a precise system of rules and must forever remain the domain of professional experts on reasoning.

If constructing a system of scientific inductive logic were totally impossible, we would be left with an intellectual vacuum, which could not be filled by appeal to “experts.” For, to decide whether someone is an expert predictor or a charlatan, we must assess the evidence that his predictions will be correct.

And to assess this evidence, we must appeal to the second level of scientific inductive logic.

Fortunately there are grounds for hope. Those who have tried to construct a system of scientific inductive logic have made some solid advances. Although the intellectual jigsaw puzzle has not been put together, we at least know what some of the pieces look like. Later we shall examine some of these “building blocks” of inductive logic, but first we shall try to put the problem of constructing a system of scientific induction in perspective by examining one of the main obstacles to this goal.

IV.2. REGULARITIES AND PROJECTION. At this point you may be puzzled as to why the construction of a system of scientific inductive logic is so difficult. After all, we know that scientific induction assumes that nature is uniform and that the future will be like the past, so if, for example, all observed emeralds have been green, the premise embodying this information confers high probability on the conclusion that the next emerald to be observed will be green. We say that scientific inductive logic *projects an observed regularity* into the future because it assigns high inductive probability to the argument:

All observed emeralds have been green.

The next emerald to be observed will be green.

In contrast, counterinduction would assume that the observed regular connection between being an emerald and being green would not hold in the future, and thus would assign high inductive probability to the argument:

All observed emeralds have been green.

The next emerald to be observed will not be green.

So it seems that scientific induction, in a quite straightforward manner, takes observed patterns or regularities in nature and assumes that they will hold in the future. Along these same lines, the premise that 99 percent of the observed emeralds have been green would confer a slightly lower probability on the conclusion that the next emerald to be observed would be green. Why can we not simply say, then, that arguments of the form

All observed X's have been Y's.

The next observed X will be a Y.

have an inductive probability of 1, and that all arguments of the form

Ninety-nine percent of the observed X 's have been Y 's.

The next observed X will be a Y .

have an inductive probability of 99/100?

That is, why can we not simply construct a system of scientific induction by giving the following rule on each level?

Rule S: An argument of the form

N percent of the observed X 's have been Y 's.

The next observed X will be a Y .

is to be assigned the inductive probability $N/100$.

Rule S does project observed regularities into the future, but there are several reasons why it cannot constitute a system of scientific inductive logic.

The most obvious inadequacy of Rule S is that it only applies to arguments of a specific form, and we are interested in assessing the inductive strength of arguments of different forms. Consider arguments which, in addition to a premise stating the percentage of observed X 's that have been Y 's, have another premise stating how many X 's have been observed. Here the rule does not apply, for the arguments are not of the required form. For example, Rule S does not tell us how to assign inductive probabilities to the following arguments:

I

Ten emeralds have been observed.

Ninety percent of the observed emeralds have been green.

The next emerald to be observed will be green.

II

One million emeralds have been observed.

Ninety percent of the observed emeralds have been green.

The next emerald to be observed will be green.

Obviously scientific inductive logic should tell us how to assign inductive probabilities to these arguments, and in assigning these probabilities it should take into account that the premises of Argument II bring a much greater amount of evidence to bear than the premises of Argument I.

Another type of argument that Rule S does not tell us how to evaluate is one that includes a premise stating in what variety of circumstances the regularity has been found to hold. That is, Rule S does not tell us how to assign inductive probabilities to the following arguments:

III

Every person who has taken drug *X* has exhibited no adverse side reactions.

Drug *X* has only been administered to persons between 20 and 25 years of age who are in good health.

The next person to take drug *X* will have no adverse side reactions.

IV

Every person who has taken drug *X* has exhibited no adverse side reactions.

Drug *X* has been administered to persons of all ages and varying degrees of health.

The next person to take drug *X* will have no adverse side reactions.

Again, scientific inductive logic should tell us how to assign inductive probabilities to these arguments, and in doing so it should take into account the fact that the premises of Argument IV tell us that the regularity has been found to hold in a great variety of circumstances, whereas the premises of Argument III inform us that the regularity has been found to hold in only a limited area.

There are many other types of argument that Rule S does not tell us how to evaluate, including most of the arguments advanced as examples in Chapter I. We can now appreciate why an adequate system of rules for scientific inductive logic must be a fairly complex structure. But there is another shortcoming of Rule S which has to do with arguments to which it does apply, that is, arguments of the form:

N percent of the observed *X*'s have been *Y*'s.

The next observed *X* will be a *Y*.

The following two arguments are of that form, so we can apply Rule S to evaluate them:

V

One hundred percent of the observed samples of pure water have had a freezing point of +32 degrees Fahrenheit.

The next observed sample of pure water will have a freezing point of + 32 degrees Fahrenheit.

VI

One hundred percent of the recorded economic depressions have occurred at the same time as large sunspots.

The next economic depression will occur at the same time as a large sunspot.

If we apply Rule S we find that it assigns an inductive probability of 1 to each of these arguments. But surely Argument V has a much higher degree of inductive strength than Argument VI! We feel perfectly justified in projecting into the future the observed regular connection between a certain type of chemical compound and its freezing point. But we feel that the observed regular connection between economic cycles and sunspots is a coincidence, an accidental regularity or spurious correlation, which should not be projected into the future. We shall say that the observed regularity reported in the premise of Argument V is *projectible*, while the regularity reported in the premise of Argument VI is not. We must now sophisticate our conception of scientific inductive logic still further. Scientific inductive logic does project observed regularities into the future, but only projectible regularities. It does assume that nature is uniform and that the future will resemble the past, but only in certain respects. It does assume that observed patterns in nature will be repeated, but only certain types of patterns. Thus, Rule S is not adequate for scientific inductive logic because it is incapable of taking into account differences in projectibility of regularities.

Exercises

1. Construct five inductively strong arguments to which Rule S does not apply.
2. Give two new examples of projectible regularities and two new examples of unprojectible regularities.
3. For each of the following arguments, state whether Rule S is applicable. If it is applicable, what inductive probability does it assign to the argument?
 - a. One hundred percent of the crows observed have been black.
The next crow to be observed will be black.
 - b. One hundred percent of the crows observed have been black.
All crows are black.
 - c. Every time I have looked at a calendar, the date has been before January 1, 2010.
The next time I look at a calendar the date will be before January 1, 2010.
 - d. Every time fire has been observed, it has continued to burn according to the laws of nature until extinguished.
All unobserved fires continue to burn according to the laws of nature until extinguished.

- e. Eighty-five percent of the time when I have dropped a piece of silverware, company has subsequently arrived.

The next time I drop a piece of silverware company will subsequently arrive.

IV.3. THE GOODMAN PARADOX. If one tries to construct various examples of projectible and unprojectible regularities, he will soon come to the conclusion that projectibility is not simply a “yes or no” affair, but rather a matter of degree. Some regularities are highly projectible, some have a middling degree of projectibility, and some are quite unprojectible. Just how unprojectible a regularity can be has been demonstrated by Nelson Goodman in his famous “grue-bleen” paradox.

Goodman invites us to consider a new color word, “grue.” It is to have the general logical features of our old color words such as “green,” “blue,” and “red.” That is, we can speak of things being a certain color at a certain time—for example, “John’s face is red now”—and we can speak of things either remaining the same color or changing colors. The new color word “grue” is defined in terms of the familiar color words “green” and “blue” as follows:

Definition 6: A certain thing, X , is said to be *grue* at a certain time t if and only if:

X is green at t and t is before the year 2100

or

X is blue at t and t is during or after the year 2100.

Let us see how this definition works. If you see a green grasshopper today, you can correctly maintain that you have seen a grue grasshopper today. Today is before the year 2100, and before the year 2100 something is grue just when it is green. But if you or one of your descendants sees a green grasshopper during or after the year 2100, it would then be incorrect to maintain that a grue grasshopper had been seen. During and after the year 2100, something is grue just when it is blue. Thus, after the year 2100, a blue sky would also be a grue sky.

Suppose now that a chameleon were kept on a green cloth until the beginning of the year 2100 and then transferred to a blue cloth. In terms of green and blue we would say that the chameleon changed color from green to blue. But in terms of the new color word “grue” we would say that it remained the same color: “grue.” The other side of the coin is that when something remains the same color in terms of the old color words, it will change color in terms of the new one. Suppose we have a piece of glass that is green now and that will remain green during and after the year 2100. Then we would have to say that it was grue before the year 2100 but was not grue during and after the year

2100. At the beginning of the year 2100 it changed color from grue to some other color. To name the color that it changed to we introduce the new color word “bleen.” “Bleen” is defined in terms of “green” and “blue” as follows:

Definition 7: A certain thing, X , is said to be *bleen* at a certain time t if and only if:

X is blue at t and t is before the year 2100

or

X is green at t and t is during or after the year 2100.

Thus, before the year 2100 something is grue just when it is green and bleen just when it is blue. In or after the year 2100 something is grue just when it is blue and bleen just when it is green. In terms of the old color words the piece of glass remains the same color (green), but in terms of the new color words the piece of glass changes color (from grue to bleen).

Imagine a tribe of people speaking a language that had “grue” and “bleen” as basic color words rather than the more familiar ones that we use. Suppose we describe a situation in our language—for example, the piece of glass being green before the year 2100 and remaining green afterward—in which we would say that there is no change in color. But if they correctly describe the same situation in their language, then, *in their terms*, there is a change. This leads to the important and rather startling conclusion that whether a certain situation involves change or not may depend on the descriptive machinery of the language used to discuss that situation.

One might object that “grue” and “bleen” are not acceptable color words because they have reference to a specific date in their definitions. It is quite true that *in our language*, in which blue and green are the basic color words, grue and bleen must be defined not only in terms of blue and green but also in terms of the date “2100 A.D.” But a speaker of the grue-bleen language could maintain that definitions of our color words in his language must also have reference to a specific date. In the grue-bleen language, “grue” and “bleen” are basic, and “blue” and “green” are defined as follows:

Definition 8: A certain thing, X , is said to be *green* at a certain time t if and only if:

X is grue at t and t is before the year 2100

or

X is bleen at t and t is during or after the year 2100.

Definition 9: A certain thing X is said to be *blue* at a certain time t if and only if:

X is bleen at t and t is before the year 2100

or

X is grue at t and t is during or after the year 2100.

Defining the old color words in terms of the new requires reference to a specific date as much as defining the new words in terms of the old. So the formal structure of their definitions gives no reason to believe that “grue” and “bleen” are not legitimate, although unfamiliar, color words.

Let us see what can be learned about regularities and projectibility from these new color words. We have already shown that whether there is change in a given situation may depend on what linguistic machinery is used to describe that situation. We shall now show that what regularities we find in a given situation also may depend on our descriptive machinery. Suppose that at one minute to midnight on December 31, 2099, a gem expert is asked to predict what the color of a certain emerald will be after midnight. He knows that all observed emeralds have been green. He projects this regularity into the future and predicts that the emerald will remain green. Notice that this is in accordance with Rule S, which assigns an inductive probability of 1 to the argument:

One hundred percent of the times that emeralds have been observed they have been green.

The next time that an emerald is observed it will be green.

But if the gem expert were a speaker of the grue-bleen language, he would find a different regularity in the color of observed emeralds. He would notice that every time an emerald had been observed it had been grue. (Remember that before the year 2100 everything that is green is also grue.) Now if he followed Rule S he would project *this* regularity into the future, for Rule S also assigns an inductive probability of 1 to the argument:

One hundred percent of the times emeralds have been observed they have been “grue.”

The next time an emerald is observed it will be “grue.”

And if he projected the regularity that all observed emeralds have been grue into the future, he would predict that the emerald will remain grue. But during the year 2100 a thing is “grue” only if it is blue. So by projecting this regularity he is in effect predicting that the emerald will change from green to blue.

Now, we will all agree that this is a ridiculous prediction to make on the basis of the evidence. And no one is really claiming that it should be made. But it cannot be denied that this prediction results from the projection into the

future of an observed regularity in accordance with Rule S. The point is that the regularity of every observed emerald having been grue is a totally unprojectible regularity. And the prediction of our hypothetical grue-bleen-speaking gem expert is an extreme case of the trouble we get into when we try to project, via some rule such as Rule S, regularities that are in fact unprojectible.

The trouble we get into is indeed deep, for the prediction so arrived at will conflict with the prediction arrived at by projecting a projectible regularity. If we project the projectible regularity that every time an emerald has been observed it has been green, then we arrive at the prediction that the emerald will remain green. If we project the unprojectible regularity that every time an emerald has been observed, it has been grue, then we arrive at the prediction that the emerald will change from green to blue. These two predictions clearly are in conflict.¹

Thus, the mistake of projecting an unprojectible regularity may not only lead to a ridiculous prediction. It may, furthermore, lead to a prediction that conflicts with a legitimate prediction which results from projecting a projectible regularity discovered in *the same set of data*. An acceptable system of scientific inductive logic must provide some means to escape this conflict. It must incorporate rules that tell us which regularities are projectible. From the discussion of accidental regularities and the sunspot theory of economic cycles, we already know that scientific inductive logic must have rules for determining projectibility. But the Goodman paradox gives this point new urgency by demonstrating how unprojectible a regularity can be and how serious are the consequences of projecting a totally unprojectible regularity.

Let us summarize what is to be learned from the discussion of “grue” and “bleen”:

1. Whether we find change or not in a certain situation may depend on the linguistic machinery we use to describe that situation.
2. What regularities we find in a sequence of occurrences may depend on the linguistic machinery used to describe that sequence.
3. We may find two regularities in a sequence of occurrences, one projectible and one unprojectible, such that the predictions that arise from projecting them both are in conflict.

¹ Actually they are inconsistent only under the assumption that the emerald will not be destroyed before 2100 A.D., but presumably we will have independent inductive evidence for this assumption.

Exercise:

Define “grue” in terms of “blue,” “green,” and “bleen” without mentioning the year 2100. You can use “and,” “or,” and “not.”

IV.4. THE GOODMAN PARADOX, REGULARITY, AND THE PRINCIPLE OF THE UNIFORMITY OF NATURE. We saw, in the last section, that projecting observed regularities into the future is not as simple as it first appears. The regularities found in a certain sequence of events may depend on the language used to describe that sequence of events. The Goodman paradox showed that if we try to project all regularities that can be found by using any language, our predictions may conflict with one another. This is a startling result, and it dramatizes the need for rules for determining projectibility in scientific induction. (This might be accomplished through the specification of the most fruitful language for scientific description of events.)

This need is further dramatized by the following, even more startling result: For any prediction whatsoever, we can find a regularity whose projection licenses that prediction. Of course, most of these regularities will be unprojectible. The point is that we need rules to eliminate those predictions based on unprojectible regularities. I shall illustrate this principle in three ways: (1) in an example that closely resembles Goodman’s “grue-bleen” paradox, (2) with reference to the extrapolation of curves on graphs, (3) with reference to the problem, often encountered on intelligence tests, of continuing a sequence of numbers. The knowledge gained from this discussion will then be applied to a reexamination of the principle of the uniformity of nature.

Example 1

Suppose you are presented with four boxes, each labeled “Excelsior!” In the first box you discover a green insect; in the second, a yellow ball of wax; in the third, a purple feather. You are now told that the fourth box contains a mask and are asked to predict its color. You must look for a regularity in this sequence of discoveries, whose projection will license a prediction as to the color of the mask. Although on the face of it, this seems impossible, with a little ingenuity a regularity can be found. What is more, for any prediction you wish to make, there is a regularity whose projection will license that prediction. Suppose you want to predict that the mask will be red. The regularity is found in the following manner.

Let us define a new word, “snarf.” A snarf is something presented to you in a box labeled “Excelsior!” and is either an insect, a ball of wax, a feather, or a mask. Now you have observed three snarfs and are about to observe a fourth. This is a step toward regularity, but there is still the problem that the three

observed snarfs have been different colors. One more definition is required in order to find regularity in apparent chaos. A thing X is said to be “murkle” just when:

X is an insect *and* X is green

or

X is a ball of wax *and* X is yellow

or

X is a feather *and* X is purple

or

X is some other type of thing *and* X is red.

Now we have found the regularity: all observed snarfs have been murkle. If we project this regularity into the future, assuming that the next snarf to be observed will be murkle, we obtain the required prediction.² The next snarf to be observed will be a mask, and for a mask to be murkle it must be red. Needless to say, this regularity is quite unprojectible. But it is important to see that we could discover an unprojectible regularity that, if it were projected, would lead to the prediction that the mask is red. And it is easy to see that, if we wanted to discover a regularity that would lead to a prediction that the mask will be a different color, a few alterations to the definition of “murkle” would accomplish this aim. This sort of thing can always be done and, as we shall see, in some areas we need not even resort to such exotic words as “snarf,” “murkle,” “grue,” and “bleen.”

Example 2

When basing predictions on statistical data we often make use of graphs, which help summarize the evidence and guide us in making our predictions. To illustrate, suppose a certain small country takes a census every 10 years, and has taken three so far. The population was 11 million at the time of the first census, 12 million at the second census, and 13 million at the third. This information is represented on a graph in Figure IV.1. Each dot represents the information as to population size gained from one census. For example, the middle dot represents the second census, taken in the year 10, and showing a population of 12 million. Thus, it is placed at the intersection of the vertical

²This projection is in accordance with Rule S, which assigns an inductive probability of 1 to the argument:

All observed snarfs have been murkle.

The next snarf to be observed will be murkle.




Figure IV. 1

line drawn from the year 10 and the horizontal line drawn from the population of 12 million.

Suppose now you are asked to predict the population of this country at the time of the fourth census, that is, in the year 30. You would have to look for a regularity that could be projected into the future. In the absence of any further information, you would probably proceed as follows: First you would notice that the points representing the first three census all fall on the straight line labeled A in Figure IV.2, and would then project this regularity into the future. This is in accordance with Rule S, which assigns an inductive probability of 1 to the following argument:

All points representing census so far taken have fallen on line A.

The point representing the next census to be taken will fall on line A.

This projection would lead you to the prediction that the population at the time of the fourth census will be 14 million, as shown by the dotted lines in Figure IV.2. The process by which you would arrive at your prediction is called *extrapolation*. If you had used similar reasoning to estimate the population during the year 15 at 12.5 million, the process would be called *interpolation*. Interpolation is estimating the position of a point that lies *between* the points representing the data. Extrapolation is estimating the position of a point that lies *outside* the points representing the data. So your prediction would be




Figure IV.2

obtained by extrapolation, and your extrapolation would be a projection of the regularity that all the points plotted so far fell on line *A*.

But it is obvious that there are quite a few other regularities to be found in the data which you did not choose to project. As shown in Figure IV.3 there is the regularity that all the points plotted so far fall on curve *B*, and the regularity that all the points plotted so far fall on curve *C*. The projection of one of these regularities will lead to a different prediction.

If you extrapolate along curve *B*, you can predict that the population in the year 30 will be back to 11 million. If you extrapolate along curve *C*, you can predict that the population will leap to 17 million. There are indeed an infinite number of curves that pass through all the points and thus an infinite number of regularities in the data. Whatever prediction you wish to make, a regularity can be found whose projection will license that prediction.

Example 3

Often intelligence and aptitude tests contain problems where one is given a sequence of numbers and asked to continue the sequence; for example:

- i. 1, 2, 3, 4, 5, ;
- ii. 2, 4, 6, 8, 10, ;
- iii. 1, 3, 5, 7, 9,




Figure IV.3

The natural way in which to continue sequence (i) is to add 6 to the end, for sequence (ii) to add 12, and for sequence (iii) to add 11. These problems are really problems of inductive logic on the intuitive level; one is asked to discover a regularity in the segment of the series given and to project that regularity in order to find the next number of the series.

Let us make this reasoning explicit for the three series given. In example (i) the first member of the series is 1, the second member is 2, the third member is 3, and, in general, for all the members given, the k th member is k . If we project this regularity to find the next member of the series, we will reason that the sixth member is 6, which is the answer intuitively arrived at before. In example (ii) the first member is twice 1, the second is twice 2, and, in general, for all the members given, the k th member is twice k . If we project this regularity, we will reason that the sixth member is twice 6, or 12, which is the answer intuitively arrived at before. In example (iii) the first member is twice 1 less 1, the second member is twice 2 less 1, and the third member is twice 3 less 1. In general, for all the members given, the k th member is twice k less 1. If we project this regularity, we will reason that the sixth member of the series is twice 6 less 1, or 11, which is the result intuitively arrived at. We say that k is a *generating function* for the first series, $2k$ a generating function for the second series, and $2k - 1$ a generating function for the third series. Although

“generating function” may sound like a very technical term, its meaning is quite simple. It is a formula with k in it, such that if 1 is substituted for k it gives the first member of the series, if 2 is substituted for k it gives the second member, and so on.

Thus, the regularity we found in each of these series is that a certain generating function yielded all the given members of the series. This regularity was projected by assuming that the same generating function would yield the next member of the series, and so we were able to fill in the ends of the series. For example, the prediction that the sixth member of series (iii) is 11 implicitly rests on the following argument:

For every given member of series (iii) the k th member of that series was $2k - 1$.

For the next member of series (iii) the k th member will be $2k - 1$.

But, as you may expect, there is a fly in the ointment. If we look more closely at these examples, we can find *other* regularities in the given members of the various series. And the projection of these other regularities conflicts with the projection of the regularities we have already noted. The generating function $(k - 1)(k - 2)(k - 3)(k - 4)(k - 5) + k$ also yields the five given members of series (i). (This can be checked by substituting 1 for k , which gives 1; 2 for k , which gives 2; and so on, up through 5.) But if we project this regularity, the result is that the sixth member of the series is 126!

Indeed, whatever number we wish to predict for the sixth member of the series, there is a generating function that will fit the given members of the series and that will yield the prediction we want. It is a mathematical fact that in general this is true. For any finite string of numbers which begins a series, there are generating functions that fit that string of given numbers and yield whatever next member is desired. Whatever prediction we wish to make, we can find a regularity whose projection will license that prediction.

Thus, if the intelligence tests were simply looking for the projection of a regularity, any number at the end of the series would be correct. What they are looking for is not simply the projection of a regularity but the projection of an intuitively projectible regularity.

If we have perhaps belabored the point in Examples (1), (2), and (3) we have done so because the principle they illustrate is so hard to accept. Any prediction whatsoever can be obtained by projecting regularities. As Goodman puts it, “To say that valid predictions are those based on past regularities, without being able to say *which* regularities, is thus quite pointless. Regularities are where you find them, and you can find them anywhere.” An acceptable scientific inductive logic must have rules for determining the projectibility of regularities.

It remains to be shown how this discussion of regularities and projectibility bears on the principle of the uniformity of nature. Just as we saw that the naïve characterization of scientific inductive logic as a system that projects observed regularities into the future was pointless unless we can say which regularities it projects, so we shall see that the statement that scientific inductive logic presupposes the uniformity of nature is equally pointless unless we are able to say *in what respects* nature is presupposed to be uniform. For it is self-contradictory to say that nature is uniform in all respects, and trivial to say it is uniform in some respects.

In the original statement of the Goodman paradox, the gem expert, who spoke our ordinary language, assumed nature to be uniform with respect to the blueness or greenness of emeralds. Since observed emeralds had always been green, and since he was assuming that nature is uniform and that the future would resemble the past in this respect, he predicted that the emerald would remain green. But the hypothetical gem expert who spoke the grue-bleen language assumed nature to be uniform *with respect to the grueness or bleeness of emeralds*. Since observed emeralds had always been grue and since he was assuming that nature is uniform and that the future would resemble the past in this respect, he predicted that the emerald would remain grue. But we saw that these two predictions were in conflict. The future cannot resemble the past in both these ways. As we have seen, such conflicts can be multiplied *ad infinitum*. The future cannot resemble the past in all respects. It is self-contradictory to say that nature is uniform in all respects.

We might try to retreat to the claim that scientific induction presupposes that nature is uniform in some respects. But this claim is so weak as to be no claim at all. To say that nature is uniform in some respects is to say that it exhibits some patterns, that there are some regularities in nature taken as a whole (in both the observed and unobserved parts of nature). But as we have seen in this section, in any sequence of observations, no matter how chaotic the data may seem, there are always regularities. This holds not only for sequences of observations but also for nature as a whole. No matter how chaotic nature might be, it would always exhibit some patterns; it would always be uniform in some respects. These uniformities might seem highly artificial, such as a uniformity in terms of “grue” and “bleen” or “snarf” and “murkle.” They might be fiendishly complex. But no matter how nature might behave, there would always be some uniformity, “natural” or “artificial,” simple or complex. It is therefore trivial to say that nature is uniform in some respects. Thus, if the statement that scientific induction presupposes that nature is uniform is to convey any information at all, it must specify in what respects scientific induction presupposes that nature is uniform.

The points about regularities and projectibility and the uniformity of nature are really two sides of the same coin. There are so many regularities in any sequence of observations and so many ways for nature to be uniform that the statements “Scientific induction projects observed regularities into the future” and “Scientific induction presupposes the uniformity of nature” lose all meaning. They can, however, be reinvested with meaning if we can formulate *rules of projectibility* for scientific inductive logic. Then we could say that scientific inductive logic projects regularities that meet these standards. And that would be saying something informative. We could reformulate the principle of the uniformity of nature to mean: Nature is such that projecting regularities that meet these standards will lead to correct predictions most of the time. Thus, the whole concept of scientific inductive logic rests on the idea of projectibility. The problem of formulating precise rules for determining projectibility is the new riddle of induction.

Exercise:

In the example of the four boxes labeled “Excelsior!” find a regularity in the observations whose projection would lead to the prediction that the mask will be blue.

IV.5. SUMMARY. This chapter described the scope of the problem of constructing a system of scientific inductive logic. We began with the supposition that scientific inductive logic could be simply characterized as the projection of observed regularities into the future in accordance with some rule, such as Rule S. We saw that this characterization of scientific inductive logic is inadequate for several reasons, the most important being that too many regularities are to be found in any given set of data. In one set of data we can find regularities whose projection leads to conflicting predictions. In fact, for any prediction we choose, there will be a regularity whose projection licenses that prediction.

Scientific inductive logic must select from the multitude of regularities present in any sequence of observations, for indiscriminate projection leads to paradox. Thus, in order to characterize scientific inductive logic we must specify the rules used to determine which regularities it considers to be projectible. The problem of formulating these rules is called the new riddle of induction.

Essentially the same problem reappears if we try to characterize scientific inductive logic as a system that presupposes that nature is uniform. To say that nature is uniform in *some* respects is trivial. To say that nature is uniform in *all* respects is not only false but self-contradictory. Thus, if we are to characterize scientific inductive logic in terms of some principle of the uniformity of nature

which it presupposes, we must say in what respects nature is presupposed to be uniform, which in turn determines what regularities scientific inductive logic takes to be projectible. So the problem about the uniformity of nature is just a different facet of the new riddle of induction.

The problem of constructing a system of scientific inductive logic will not be solved until the new riddle of induction and other problems have been solved. Although these solutions have not yet been found, there have been developments in the history of inductive logic which constitute progress towards a system.

In the next chapter we shall pursue an analysis of causality which casts some light on well-known features of the experimental method. Then we will discuss the major achievement of the field, the probability calculus.

Suggested reading

Nelson Goodman, *Fact, Fiction and Forecast* (Cambridge, MA: Harvard University Press, 1983).

V

Mill's Methods of Experimental Inquiry and the Nature of Causality

V.1. INTRODUCTION. One of the purposes of scientific inductive logic is to assess the evidential warrant for statements of cause and effect. But what exactly do statements claiming causal connection *mean*, and what is their relation to statements describing *de facto* regularities? These are old and deep questions and we can give only partial answers here.

In his *System of Logic*, published in 1843, John Stuart Mill discussed five “methods of experimental inquiry” that he found used in the work of contemporary scientists. When we make some simple distinctions between different senses of “cause,” we will find that we can use the basics of logic introduced in Chapter I to give a logical analysis of Mill’s methods.

V.2. CAUSALITY AND NECESSARY AND SUFFICIENT CONDITIONS. Many of the inquiries of both scientific research and practical affairs may be characterized as the search for the causes of certain effects. The practical application of knowledge of causes consists either in producing the cause in order to produce the effect or in removing the cause in order to prevent the effect. Knowledge of causes is the key to control of effects. Thus, physicians search for the cause of certain diseases so that they may remove the cause and prevent the effect. On the other hand, advertising men engage in motivational research into the causes of consumer demand so that they can produce the cause and thus produce the effect of consumer demand for their products.

However, the word “cause” is used in English to mean several different things. For this reason, it is more useful to talk about *necessary conditions* and *sufficient conditions* rather than about causes.

Definition 10: A property F is a *sufficient condition* for a property G if and only if *whenever F is present, G is present.*

Definition 11: A property H is a *necessary condition* for a property I if and only if *whenever I is present, H is present.*

Being run over by a steamroller is a sufficient condition for death, but it is not a necessary condition. Whenever someone has been run over by a

steamroller, he is dead. But it is not the case that anyone who is dead has been run over by a steamroller. On the other hand, the presence of oxygen is a necessary condition, but not a sufficient condition for combustion. Whenever combustion takes place, oxygen is present. But happily it is not true that whenever oxygen is present, combustion takes place. When we say that A causes B we sometimes mean that A is a sufficient condition for B , sometimes that A is a necessary condition for B , sometimes that A is both necessary and sufficient for B , and sometimes none of these things.

If we are looking for causes in order to produce an effect, it is reasonable to look for sufficient conditions for that effect. If we can manipulate circumstances so that the sufficient condition is present, the effect will also be present. If we are looking for causes in order to prevent an effect, it is reasonable to look for necessary conditions for that effect. If we prevent a necessary condition from materializing, we can prevent the effect. The eradication of yellow fever is a striking illustration of this strategy. Doctors discovered that being bitten by a certain type of mosquito was a necessary condition for contracting yellow fever. It was not a sufficient condition, for some people who were bitten by these mosquitos did not contract yellow fever. Consequently, a campaign was instituted to destroy these mosquitos through the widespread use of insecticide and thus to prevent yellow fever.

From the definitions of necessary and sufficient conditions, we can prove several important principles. It follows immediately from the definitions that:

1. If A is a sufficient condition for B , then B is a necessary condition for A .
2. If C is a necessary condition for D , then D is a sufficient condition for C .

To say that A is a sufficient condition for B is, by definition, to say that whenever A is present, B is present. But to say that B is a necessary condition for A is, by definition, to say the same thing.

Let us look at some illustrations of these principles. Since the presence of oxygen is a *necessary* condition for combustion, then by principle 2 combustion is *sufficient* to ensure the presence of oxygen. Thus, suppose someone lowers a burning candle into a deep mine shaft he proposes to explore. If the candle continues to burn, he will know that the shaft contains sufficient oxygen to breathe. To illustrate principle 1, let us suppose that a professor has constructed a test such that a high grade on the test is *sufficient* to guarantee that the student has studied the material. Then studying the material is a *necessary* condition for doing well on the test.

Two additional principles require a little more thought:

3. If A is a sufficient condition for B , then $\sim B$ is a sufficient condition for $\sim A$.

4. If C is a necessary condition for D , then $\sim D$ is a necessary condition for $\sim C$.

Using the definition of sufficient condition, principle 3 becomes:

3'. If whenever A is present B is present, then whenever $\sim B$ is present $\sim A$ is present.

Now remember from the presence table for negation that $\sim B$ is present just when B is absent and $\sim A$ is present just when A is absent. So principle 3 can be rewritten again as:

3''. If whenever A is present B is present, then whenever B is absent A is absent.

We can now see why this principle is correct. Suppose that whenever A is present, B is present. Suppose further that B is absent in a certain case. Then A must also be absent, for if A were present, B would be present, and it is not. Let us see how this works in a concrete case. Suppose that a certain infection is a sufficient condition for a high fever; that is, everyone who suffers from this infection runs a high fever. Then the absence of a high fever is sufficient to guarantee that a person is not suffering from this infection.

That principle 4 is correct can be demonstrated in the same way. Using the definition of necessary condition, we can rewrite principle 4 as:

4'. If whenever D is present C is present, then whenever $\sim C$ is present $\sim D$ is present.

And using the presence table for negation, we can rewrite it again as:

4''. If whenever D is present C is present, then whenever C is absent D is absent.

And this is simply a restatement of principle 3'' using different letters.

We can use the same example to illustrate principle 4. Since suffering from the infection is a sufficient condition for running a high fever, running the fever is a necessary condition for having the infection (principle 1). By principle 4, since running the fever is a necessary condition for having the infection, not having the infection is a necessary condition for not running a fever. (It is not a sufficient condition since other diseases might result in a fever.)

Two more principles will complete this survey of the basic principles governing necessary and sufficient conditions:

5. If A is a sufficient condition for B , then $\sim A$ is a necessary condition for $\sim B$.
6. If C is a necessary condition for D , then $\sim C$ is a sufficient condition for $\sim D$.

Using the definitions and the presence table for negation we can rewrite principle 5 as:

- 5'. If whenever A is present B is present, then whenever B is absent A is absent.

But this is exactly what we established in 3". In the same manner, we can rewrite principle 6 as:

- 6'. If whenever D is present C is present, then whenever C is absent, D is absent.

But this is exactly what we established in principle 4". A concrete illustration of principle 5 is that, since being run over by a steamroller is a sufficient condition for death, not being run over by a steamroller is a necessary condition for staying alive. And principle 6 can be illustrated by the observation that, if studying is a necessary condition for passing a test, not studying is a sufficient condition for failing it.

When we speak of the cause of an effect in ordinary language, we sometimes mean a sufficient condition, as when we say that the infection was the cause of the fever or that being run over by a steamroller was the cause of death. Sometimes we mean a necessary condition, as when we say that yellow fever was caused by the bite of the mosquito or a high score on the test was due to diligent study. On the other hand, necessary and sufficient conditions are sometimes not causes at all but rather *symptoms* or *signs*. The continued burning of the candle was a sign of the presence of oxygen. The high fever was a *symptom* of the infection. When we analyze Mill's methods, it will be seen that the precise language of necessary and sufficient conditions is much more useful than the vague language of cause and effect, sign and symptom.

Exercises

Show that the following principles are correct and give a concrete illustration of each:

1. If $\sim B$ is a sufficient condition for $\sim A$, then A is a sufficient condition for B .
2. If $\sim D$ is a necessary condition for $\sim C$, then C is a necessary condition for D .
3. If $\sim A$ is a necessary condition for $\sim B$, then A is a sufficient condition for B .
4. If $\sim C$ is a sufficient condition for $\sim D$, then C is a necessary condition for D .

5. If A is a necessary condition for E and B is a necessary condition for E , then $A \& B$ is a necessary condition for E .
6. If $A \& B$ is a necessary condition for E , then A is a necessary condition for E and B is a necessary condition for E .
7. If A is a sufficient condition for E and B is a sufficient condition for E , then $A \vee B$ is a sufficient condition for E .
8. If $A \vee B$ is a sufficient condition for E , then A is a sufficient condition for E and B is a sufficient condition for E .
9. If A is a necessary condition for E , then whatever the property F , $A \vee F$ is a necessary condition for E .
10. If A is a sufficient condition for E , then whatever the property F , $A \& F$ is a sufficient condition for E .

V.3. MILL'S METHODS. Mill presented five methods designed to guide the experimenter in his search for causes. They are the *method of agreement*, the *method of difference*, the *joint method*, the *method of concomitant variation*, and the *method of residues*. However, Mill did not actually originate these methods, nor did he fully understand them.

The theoretical basis of Mill's methods has only recently been fully explored by the philosopher G. H. von Wright. Following von Wright, we will present Mill's methods a little differently than Mill did. We will be able to uncover the theoretical basis of Mill's methods in a discussion of the method of agreement, the method of difference, and the joint method. Since there is nothing essentially new in the method of concomitant variation and the method of residues, we shall not discuss them. However, we shall still be left with five methods, for there are two variations of the method of agreement and two variations of the joint method.

These methods are to be viewed as methods of finding the necessary or sufficient conditions of a given property. The property whose necessary or sufficient conditions are being sought is called the *conditioned property*. A conditioned property may have more than one sufficient condition. If the conditioned property is death, being run over by a steamroller is one sufficient condition for it, but there are many others. A conditioned property may also have more than one necessary condition. If the conditioned property is the occurrence of combustion, the presence of oxygen is a necessary condition for it, but so is the presence of an oxidizable substance. Those properties suspected of being necessary or sufficient conditions for a given conditioned property are called *possible conditioning properties*. The general problem is, "How is the information gained from observing various occurrences used to pick out the necessary and sufficient conditions from the

possible conditioning properties?" The following methods are attempts to answer this question.

V.4. THE DIRECT METHOD OF AGREEMENT. Suppose that one of the possible conditioning properties A , B , C , or D is suspected of being a necessary condition for the conditioned property E , but which one is not known. Suppose further that, either by experimental manipulation or simply by studious observation, a wide variety of occurrences are observed in which E is present, and that the only possible conditioning property that is present on all these occasions is C . The set of observations shown in Example 1 corre-

Example 1

	Possible conditioning properties				Conditioned property
	A	B	C	D	E
Occurrence 1:	P	P	P	A	P
Occurrence 2:	P	A	P	P	P
Occurrence 3:	A	P	P	A	P

sponds to this description. Occurrence 1 shows that D cannot be a necessary condition for E . The definition of necessary condition tells us that a necessary condition for E must be present whenever E is present. But in 1, E is present while D is absent. Thus, occurrence 1 eliminates D from the list of possible necessary conditions. In the same manner, occurrence 2 shows that B cannot be a necessary condition for E , since E is present while B is absent. Occurrence 3 eliminates A and eliminates D once more. The only candidate left for the office of necessary condition for E is C . The observations show that if one of the possible conditioning properties is in fact a necessary condition for E , then C must be that necessary condition.

In Example 1 three occurrences were required before A , B , and D could be eliminated as possible necessary conditions for E . Actually, we might have done without occurrence 1 since occurrence 3 also eliminated D . However, in the occurrence shown in Example 2 all could be eliminated at one stroke. The principle of elimination is the same: Any property that is absent when E is present cannot be a necessary condition for E .

Suppose someone were to object that the absence of D might be necessary for the presence of E ; that is, that $\sim D$ might be a necessary condition for E , and that the data in Example 2 have not eliminated that possibility. This is correct, but it shows no defect in the argument. Only the simple properties A , B , C , and D were included in the possible conditioning properties; the complex property $\sim D$ was not. And all that was claimed was that *if* one of the possible

Example 2

	Possible conditioning properties				Conditioned property
	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	<i>E</i>
Occurrence 1:	A	A	P	A	P

Example 3

	Possible conditioning properties								Conditioned property
	Simple				Complex				<i>E</i>
	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	$\sim A$	$\sim B$	$\sim C$	$\sim D$	
Occurrence 1:	A	A	P	A	P	P	A	P	P

conditioning properties is a necessary condition for *E*, then *C* is that necessary condition. But if we were to add the negations of *A*, *B*, *C*, and *D* to our list of possible conditioning properties, then occurrence 1 of Example 2 would not suffice to eliminate all the alternatives but *C*. This is readily shown in Example 3.

We can tell whether a complex property is present or absent in a given occurrence from the information as to whether its constituent simple properties are present or absent. This information will be found in the presence table for that complex property. Here it need only be remembered, from the presence table for negation, that the negation of a property is absent when that property is present, and present when that property is absent. Now in Example 3, occurrence 1 shows that *A*, *B*, *D*, and $\sim C$ cannot be necessary conditions for *E*. This leaves *C*, $\sim A$, $\sim B$, and $\sim D$ as likely candidates. If the field is to be narrowed down, some more occurrences must be observed. These might give the results shown in Example 4. Again occurrence 1 eliminates *A*, *B*, *D*, and $\sim C$. Occurrence 2 further eliminates $\sim A$, occurrence 3 eliminates $\sim B$, and occurrence 4 eliminates $\sim D$. Thus, the only possible conditioning property left is *C*. If any one of the possible conditioning properties is a necessary condition for *E*, then *C* is that necessary condition.

In Example 4 it took four occurrences to eliminate all the possible conditioning properties but one. However, the two occurrences observed in Example 5 would have done the job. Occurrence 1 shows that *A*, $\sim B$, $\sim C$, and $\sim D$ cannot be necessary conditions for *E* since they are absent when *E* is present. Occurrence 2 further eliminates *B*, *D*, $\sim A$, and $\sim C$, leaving only *C*. Thus, in this example if one of the possible conditioning properties is a necessary condition for *E*, then *C* is that necessary condition. It is true, in gen-

Example 4

	Possible conditioning properties								Conditioned property <i>E</i>
	Simple				Complex				
	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	$\sim A$	$\sim B$	$\sim C$	$\sim D$	
Occurrence 1:	A	A	P	A	P	P	A	P	P
Occurrence 2:	P	A	P	A	A	P	A	P	P
Occurrence 3:	A	P	P	A	P	A	A	P	P
Occurrence 4:	A	A	P	P	P	P	A	A	P

Example 5

	Possible conditioning properties								Conditioned property <i>E</i>
	Simple				Complex				
	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	$\sim A$	$\sim B$	$\sim C$	$\sim D$	
Occurrence 1:	A	P	P	P	P	A	A	A	P
Occurrence 2:	P	A	P	A	A	P	A	P	P

eral, that if we admit both simple properties and their negation as possible conditioning properties, then the minimum number of occurrences that can eliminate all but one of the possible conditioning properties is 2. As we saw before, when only simple properties are admitted as possible conditioning properties, the minimum number of occurrences that can eliminate all but one of them is 1. But the basic principle of elimination remains the same in both cases: A property that is absent when *E* is present cannot be a necessary condition for *E*.

We were able to extend Mill's method of agreement to cover negative possible conditioning properties, and this makes sense, for negative properties are quite often necessary conditions. Not being run over by a steamroller is a necessary condition for remaining alive and not letting one's grade average fall below a certain point may be a necessary condition for remaining in college. We are interested in negative necessary conditions because they tell us what we must avoid in order to attain our goals. But negations of simple properties are not the only complex properties that may be important necessary conditions.

Let us consider disjunctions of simple properties as necessary conditions. Either having high grades in high school *or* scoring well on the entrance examination might be a necessary condition for getting into college. It might not be a sufficient condition since someone who meets this qualification might still be

rejected on the grounds that he is criminally insane. To take another example, in football either making a touchdown *or* a field goal *or* a conversion *or* a

Example 6

	Possible conditioning properties					Conditioned property <i>E</i>
	Simple				Complex	
	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	<i>BvC</i>	
Occurrence 1:	A	P	A	A	P	P
Occurrence 2:	A	A	P	A	P	P

safety is a necessary condition for scoring. In this case the necessary condition is also a sufficient condition. We are interested in disjunctive necessary conditions because they lay out a field of alternatives, one of which must be realized if we are to achieve certain ends.

The question of what happens when disjunctions (alternations) of simple properties are allowed into a set of possible conditioning properties is too involved to be treated fully here. But the principle of elimination remains the same. We can see how this principle operates in two simplified examples that allow only simple properties and one disjunction as possible conditioning properties. In Example 6 the complex property *BvC* is the only property that is always present when *E* is present. Occurrences 1 and 2 eliminate all the simple properties as necessary conditions. Thus, if one of the possible conditioning properties is a necessary condition for *E*, *BvC* is that necessary condition.

In Example 6 the disjunction was the property left after all the others had been eliminated. Let us now look at Example 7, where the disjunction itself is eliminated. Occurrence 1 eliminates *A* and *C* as necessary conditions, and

Example 7

	Possible conditioning properties					Conditioned property <i>E</i>
	Simple				Complex	
	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	<i>BvC</i>	
Occurrence 1:	A	P	A	P	P	P
Occurrence 2:	P	A	A	P	A	P

occurrence 2 shows that neither *B* nor *C* nor *BvC* can be a necessary condition for *E*. This leaves only *D*, so if one of the possible conditioning properties is a necessary condition for *E*, then *D* is that necessary condition. We shall not

explore further the treatment of complex possible conditioning properties by the direct method of agreement.¹ But you cannot go wrong if you remember that the principle of elimination in the direct method of agreement is: A property that is absent when E is present cannot be a necessary condition for E .

Use of the direct method of agreement requires looking for occurrences of the conditioned property in circumstances as varied as possible. If these circumstances are so varied that only one of the possible conditioning properties is present whenever the conditioned property is present, it may be suspected that that property is a necessary condition for the conditioned property. It has been shown that the logic behind this method is the same as the logic behind the method of the master detective who eliminates suspects one by one in order to find the murderer. If only one of the possible conditioning properties is present whenever the conditioned property is present, then all the other possible conditioning properties are eliminated as necessary conditions since they are each absent in at least one occurrence in which the conditioned property is present.

But the method of agreement resembles the method of the master detective in two further ways. When starting on a murder case, the detective cannot be sure that he will be able to eliminate all the suspects but one. After all, the murder might have been done by two people working together. The same is true of the method of agreement, for a conditioned property can have more than one necessary condition. Moreover, the master detective may not have the murderer or murderers in his initial list of suspects and may end up eliminating all the possibilities. In this case he will have to go back and look for more suspects to include in a more comprehensive list. In a similar manner the scientist may not have included the necessary condition or conditions for a conditioned property in his initial list of possible conditioning properties. Thus, his observations might eliminate all his possible conditioning properties.

Exercises

1. In Example 1 which of the following complex properties are eliminated as necessary conditions for E by occurrences 1, 2, and 3?

¹ The direct method of agreement can be expanded to include simple properties, negations of simple properties, and disjunctions of simple properties and their negations as possible conditioning properties. There is no need to worry about conjunctions since a conjunction, that is, $F \& G$, is a necessary condition for E if and only if F is a necessary condition for E and G is a necessary condition for E . Thus, if we can discover all the individual necessary conditions, we automatically have all the conjunctive necessary conditions.

- a. $\sim A$. d. $\sim D$.
 b. $\sim B$. e. $A \vee D$.
 c. $\sim C$. f. $B \vee C$.

2. In the following example, for each occurrence find whether the complex properties are present or absent and which of the possible conditioning properties are eliminated as necessary conditions for E :

	Possible conditioning properties							Conditioned property E	
	Simple			Complex					
	A	B	C	$\sim A$	$\sim B$	$\sim C$	$A \vee C$		$\sim B \vee C$
Occurrence 1:	P	P	P						P
Occurrence 2:	P	P	A						P
Occurrence 3:	P	A	P						P
Occurrence 4:	P	P	A						P
Occurrence 5:	A	A	P						P
Occurrence 6:	A	A	P						P

3. In Exercise 2 one of the possible conditioning properties was not eliminated. Describe an occurrence which would eliminate it.

V.5. THE INVERSE METHOD OF AGREEMENT. The inverse method of agreement is a method for finding sufficient conditions. To find a sufficient condition for a given property, E , we look for a property that is absent whenever E is absent. This is illustrated in Example 8. D is the only possible conditioning property that is absent whenever the conditioned prop-

Example 8

	Possible conditioning properties				Conditioned property E
	A	B	C	D	
Occurrence 1:	P	A	A	A	A
Occurrence 2:	A	P	A	A	A
Occurrence 3:	P	A	P	A	A

erty is absent. Thus, by the inverse method of agreement, if one of the possible conditioning properties is a sufficient condition for E , then D is that sufficient condition.

The inverse method of agreement operates in the following manner: We know from the definition of sufficient condition that a sufficient condition for E cannot be present when E is absent. To say that a certain property is a sufficient condition for E means that whenever that property is present, E is also present. Thus, in Example 8 occurrence 1 shows that A cannot be a sufficient condition for E since it is present when E is absent. Occurrence 2 shows that B cannot be a sufficient condition for E for the same reason, and occurrence 3 does the same for C and A once again. D is therefore the only property left that can be a sufficient condition for E . In this way the inverse method of agreement, like the direct method, works by eliminating possible candidates one by one.

Example 9

	Possible sufficient conditions for E				Possible necessary conditions for $\sim E$				E	$\sim E$
	A	B	C	D	$\sim A$	$\sim B$	$\sim C$	$\sim D$		
Occurrence 1:	A	A	A	A	P	P	P	P	A	P
Occurrence 2:	A	P	P	A	P	A	A	P	A	P
Occurrence 3:	P	A	P	A	A	P	A	P	A	P

The inverse method of agreement may be viewed as an application of the direct method to negative properties. This is possible in light of the principle: if $\sim A$ is a necessary condition for $\sim E$, then A is a sufficient condition for E . Example 9 illustrates this method in action. The only possible necessary condition for $\sim E$ that is present whenever $\sim E$ is present is $\sim D$. Notice that this comes to the same thing as saying that the only one of the possible sufficient conditions for E that is absent whenever E is absent is D . Thus, by the direct method of agreement, if one of the possible necessary conditions for $\sim E$ is actually a necessary condition for $\sim E$, then $\sim D$ is that necessary condition. But by the principle connecting negative necessary conditions for $\sim E$ and positive sufficient conditions for E , this is the same as saying if one of the possible sufficient conditions for E is actually a sufficient condition for E , then D is that sufficient condition. Thus, we arrive at the inverse method of agreement.

At this point it may be useful to compare the direct and inverse methods of agreement. The direct method is a method of finding *necessary conditions*. To find a necessary condition for E , we look for a property that is present whenever E is present. The direct method depends on the following principle of elimination: A property that is absent when E is present cannot be a necessary condition for E . The inverse method is a method for finding sufficient conditions. To find a sufficient condition for E , we look for a property that is absent whenever E is absent. The inverse method depends on the following principle

of elimination: A property that is present when E is absent cannot be a sufficient condition for E .

In Example 8 it required three occurrences to narrow down the field to D . However, the occurrence shown in Example 10 would alone eliminate A , B , and C . In the inverse method of agreement, as in the direct method, if we only admit simple properties as possible conditioning properties, then the least number of occurrences that can eliminate all but one of the possible conditioning properties is 1.

Example 10

	Possible conditioning properties				Conditioned property
	A	B	C	D	E
Occurrence 1:	P	P	P	A	A

Suppose, however, we wish to admit negative properties as possible conditioning properties. This is a reasonable step to take, for negative sufficient conditions can be quite important. Not staying awake while driving may be a sufficient condition for having an accident. Not being able to see may be a sufficient condition for not being called for military service. (By the principle that if $\sim F$ is sufficient for $\sim G$, then F is necessary for G , this would mean that being able to see would be a necessary condition for being called for military service.) We will introduce negative possible conditioning properties as in the section on the direct method of agreement. But this time we will rely on the principle of elimination of the inverse method of agreement: A property that is present when E is absent cannot be a sufficient condition for E . In Example 11 the only possible conditioning property that is not eliminated is a negative one.

Example 11

	Possible conditioning properties								Conditioned property
	Simple				Complex				E
	A	B	C	D	$\sim A$	$\sim B$	$\sim C$	$\sim D$	
Occurrence 1:	A	P	A	P	P	A	P	A	A
Occurrence 2:	A	P	A	A	P	A	P	P	A
Occurrence 3:	P	P	A	A	A	A	P	P	A
Occurrence 4:	A	P	P	A	P	A	A	P	A

$\sim B$ is the only possible conditioning property that is absent in every occurrence in which E is absent, so if one of the possible conditioning properties is a sufficient condition for E , then $\sim B$ is that sufficient condition.

It need not take as many occurrences as in Example 11 to eliminate all the possible conditioning properties but one. Two occurrences of the right kind could do the job, as shown in Example 12. In this example, if one of the possi-

Example 12

	Possible conditioning properties								Conditioned property <i>E</i>
	Simple				Complex				
	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	$\sim A$	$\sim B$	$\sim C$	$\sim D$	
Occurrence 1:	P	P	P	A	A	A	A	P	A
Occurrence 2:	A	P	A	P	P	A	P	A	A

ble conditioning properties is a sufficient condition for *E*, then $\sim B$ is that sufficient condition. In the inverse method of agreement, as in the direct method, if only simple properties and their negations are admitted as possible conditioning properties, then the least number of occurrences that can eliminate all but one of the possible conditioning properties is 2.

We may further extend the inverse method of agreement to allow conjunctions of simple properties as possible conditioning properties. For example, suppose we are told that eating good food *and* getting plenty of rest *and* getting a moderate amount of exercise is a sufficient condition for good health. The inverse method of agreement would *support* this contention if we found that whenever good health was absent, this complex condition was also absent (that is, if everyone who was in poor health had not eaten good food *or* had not gotten enough rest *or* had not exercised). The inverse method of agreement would *disprove* this contention if an occurrence was found where good health was absent and the complex condition was present (that is, if someone was found in poor health who had eaten good food, and gotten plenty of rest, and gotten a moderate amount of exercise).

Let us look at two examples of the inverse method of agreement where a conjunction is admitted as a possible conditioning property. In Example 13 all the possible conditioning properties except the conjunction are eliminated. The only possible conditioning property that is absent whenever *E* is

Example 13

	Possible conditioning properties					Conditioned property <i>E</i>
	Simple				Complex	
	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	<i>C&D</i>	
Occurrence 1:	P	P	A	P	A	A
Occurrence 2:	P	A	P	A	A	A

absent is the complex property $C\&D$. If one of the possible conditioning properties is a sufficient condition for E , $C\&D$ is that sufficient condition.

In Example 14 the conjunction itself is eliminated. If one of the possible conditioning properties is a sufficient condition for E , then D is that sufficient condition. We shall not explore further the treatment of complex possible conditioning properties by the inverse method of agreement.² But you cannot go wrong if you remember that the principle of elimination in the inverse method

Example 14

	Possible conditioning properties					Conditioned property E
	Simple				Complex	
	A	B	C	D	$B\&C$	
Occurrence 1:	P	P	A	A	A	A
Occurrence 2:	A	P	P	A	P	A

of agreement is: A property that is present when E is absent cannot be a sufficient condition for E .

The parallels drawn between the method of the master detective and the direct method of agreement hold also for the inverse method of agreement. It should not be thought that the field of possible sufficient conditions can always be narrowed down to one, for a property can have several sufficient conditions. We should also be prepared for the eventuality that the observed occurrences will eliminate all possible conditioning properties in the list. After all, a sufficient condition may not have been included in the list of possible conditioning properties. In such a case we would have to construct a more comprehensive list of possible conditioning properties. In some cases this more comprehensive list might be constructed by considering complex properties that were not included in the original list.

In Example 15 the five occurrences show that none of the possible conditioning properties can be a sufficient condition for E . But they suggest that the complex property $B\&C$ might be added to the list of possible conditioning properties. This property is always absent when E is absent.

² The inverse method of agreement can be expanded to include simple properties, negations of simple properties, and conjunctions of simple properties and their negations as possible conditioning properties. There is no need to worry about disjunctions, since a disjunction, $F\vee G$, is a sufficient condition for E if and only if F is a sufficient condition for E and G is a sufficient condition for E . For this reason if one can discover all the sufficient conditions that are not disjunctions, he or she will automatically have all the sufficient conditions that are disjunctions.

Example 15

	Possible conditioning properties				Conditioned property
	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	<i>E</i>
Occurrence 1:	A	P	A	P	A
Occurrence 2:	P	A	P	P	A
Occurrence 3:	P	A	A	P	A
Occurrence 4:	A	A	A	P	A
Occurrence 5:	P	P	A	A	A

The situation, however, might be more problematic. The observed occurrences might not only eliminate all the simple properties in the list but also all the (contingent) complex properties that can be constructed out of them. Such is the case in Example 16. It is impossible to discover in this list any

Example 16

	Possible conditioning properties			Conditioned property
	<i>A</i>	<i>B</i>	<i>C</i>	<i>E</i>
Occurrence 1:	P	P	P	A
Occurrence 2:	P	P	A	A
Occurrence 3:	P	A	P	A
Occurrence 4:	P	A	A	A
Occurrence 5:	A	P	P	A
Occurrence 6:	A	P	A	A
Occurrence 7:	A	A	P	A
Occurrence 8:	A	A	A	A

(contingent) complex property constructed out of *A*, *B*, and *C* which is not eliminated as a sufficient condition for *E* by these eight occurrences. In such a case some new simple properties would have to be added to the list of possible conditioning properties.

Exercises

1. In Example 9 which of the following properties are eliminated as sufficient conditions for *E* by occurrences 1, 2, and 3?

a. <i>A</i> .	c. <i>C</i> .	e. <i>A</i> & <i>B</i> .
b. <i>B</i> .	d. <i>D</i> .	f. <i>B</i> & <i>C</i> .
		g. <i>A</i> & <i>D</i> .

2. In Example 10 which of the following properties are eliminated by occurrence 1?
- | | |
|---------------|---------------------|
| a. <i>A</i> . | e. <i>A&C</i> . |
| b. <i>B</i> . | f. <i>B&C</i> . |
| c. <i>C</i> . | g. <i>A&B</i> . |
| d. <i>D</i> . | h. <i>A&D</i> . |
3. In Example 11 which of the following properties are eliminated by the four occurrences?
- | | |
|---------------------|---------------------|
| a. <i>A&B</i> . | d. <i>A&D</i> . |
| b. <i>B&C</i> . | e. <i>A&C</i> . |
| c. <i>B&D</i> . | f. <i>C&D</i> . |
4. In Example 13 are there any conjunctions of the simple properties listed other than *C&D* which are not eliminated by occurrences 1 and 2?

V.6. THE METHOD OF DIFFERENCE. The direct method of agreement was a method for finding the necessary conditions of a given property. The inverse method of agreement was a method for finding the sufficient conditions of a given property. Suppose, however, that our inquiry had a more restricted goal. Suppose that we wanted to find out which of the properties present in a certain occurrence of the conditioned property are sufficient conditions for it. To illustrate, let us suppose we find a dead man with no marks of violence on his body. In trying to determine the cause of death, we are looking for a sufficient condition for death. But we are not looking for *any* sufficient condition for death. Being run over by a steamroller is a sufficient condition for death, but that fact is irrelevant to our inquiry since this particular man was not run over by a steamroller. The conditioning property "being run over by a steamroller" is absent in this particular occurrence. What we are looking for is a sufficient condition for death among the properties that are present in this particular occurrence in which death is present. It is this sort of inquiry for which the method of difference is designed.

It is important to note why no analogous question can be raised for necessary conditions. It follows from the definition of necessary condition that all the necessary conditions for a given property must be present whenever that property is present. If loss of consciousness is a necessary condition for death, it will be present in every case of death. The questions "What properties are necessary conditions for *E*?" and "Which of the properties that are present in this particular occurrence of *E* are necessary conditions for *E*?" have exactly the same answer. In contrast, when a given property is present, some of its sufficient conditions may be absent. Many people who die have not been run over by a steamroller nor been decapitated on the guillotine. The question

“Which of the properties that are present in this particular occurrence of E are sufficient conditions for E ?” will, in general, have a shorter list of properties as its answer than the question “Which properties are sufficient conditions for E ?”

In Example 17 occurrence * does not eliminate any of the possible conditioning properties as sufficient conditions for E . But if the question of interest is “Which of the properties that are present in occurrence * are sufficient

Example 17

	Possible conditioning properties				Conditioned property E
	A	B	C	D	
Occurrence *:	P	A	P	P	P

conditions for E ?” then the candidates are limited to A , C , and D . Let us now look for other occurrences that will narrow down the field. The principle of elimination is the same as that employed in the inverse method of agreement: A property that is present when E is absent cannot be a sufficient condition for E . Therefore, let us look for additional occurrences when E is absent. Suppose that the results of our investigation are as shown in Example 18. In this example, occurrences 1 and 2 eliminate A and D as sufficient conditions for E . Of the possible conditioning properties that were present in occurrence *, only C is left. Thus, if one of the possible conditioning properties that was present in occurrence * is a sufficient condition for E , then C is that sufficient condition. Note that B might also be a sufficient condition for E , but it is not one we would be

Example 18

	Possible conditioning properties				Conditioned property E
	A	B	C	D	
Occurrence *:	P	A	P	P	P
Occurrence 1:	P	A	A	A	A
Occurrence 2:	A	A	A	P	A

interested in, since we are looking for a sufficient condition which was present in occurrence *. Occurrences 1 and 2 eliminate candidates in exactly the same way as in the inverse method of agreement. In the inverse method of agreement, however, we started with all the possible conditioning properties as candidates, while in the method of difference, we start with the possible conditioning properties that are present in a particular occurrence in which

the conditioned property is present. (We shall always call the occurrence that *defines* the candidates "occurrence *" and will number as before the occurrences that *eliminate* some of the candidates.)

If only simple properties are admitted as possible conditioning properties, then occurrence * might only leave one candidate, as shown in Example 19. In this example the only possible conditioning property present is *D*. Thus, with-

Example 19

	Possible conditioning properties				Conditioned property
	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	
Occurrence *:	A	A	A	P	P

out looking for eliminating occurrences it may be concluded that if one of the possible conditioning properties that is present in occurrence * is a sufficient condition for *E*, then *D* is that sufficient condition. In fact, if only simple properties are admitted as possible conditioning properties, then occurrence * might leave no candidates whatsoever.

But there is no reason why negations of the simple properties cannot be admitted as possible conditioning properties, as was done in the treatment of the direct and inverse methods of agreement. If both simple properties and their negations are allowed as possible conditioning properties, then in an occurrence *, exactly half of the possible conditioning properties will be left as candidates, since exactly half of them must be present in any occurrence. Further occurrences must be sought then in order to eliminate some of these properties, as in the inverse method of agreement.

Let us determine which of the properties present in occurrence * are sufficient conditions for *E*. In Example 20, *D*, $\sim A$, $\sim B$, and $\sim C$ are the possible

Example 20

	Possible conditioning properties								Conditioned property
	Simple				Complex				
	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	$\sim A$	$\sim B$	$\sim C$	$\sim D$	
Occurrence *:	A	A	A	P	P	P	P	A	P
Occurrence 1:	P	P	P	P	A	A	A	A	A
Occurrence 2:	P	A	A	A	P	A	P	P	A

conditioning properties present in occurrence * and are the candidates. Occurrences 1 and 2 eliminate *D*, $\sim A$, and $\sim C$. The only one of the candidates

which remains is $\sim B$, leading to the conclusion that if one of the possible conditioning properties that is present in occurrence * is a sufficient condition for E , then $\sim B$ is that sufficient condition.

As has been shown, if simple properties and their negations are allowed as possible conditioning properties, then occurrence * will leave exactly half of the possible conditioning properties as candidates. The least number of additional occurrences needed to eliminate all these candidates but one is one, if that one occurrence is of the right kind, as shown in Example 21. The possible conditioning properties that are present in occurrence * constitute the candidates, and they are B , C , $\sim A$, and $\sim D$. Occurrence 1 eliminates B , $\sim A$, and $\sim D$ since they are present when E is absent. Thus, if one of the possible con-

Example 21

	Possible conditioning properties								Conditioned property E
	Simple				Complex				
	A	B	C	D	$\sim A$	$\sim B$	$\sim C$	$\sim D$	
Occurrence *:	A	P	P	A	P	A	A	P	P
Occurrence 1:	A	P	A	A	P	A	P	P	A

ditioning properties that is present in occurrence * is a sufficient condition for E , then C is that sufficient condition.

If you look closely at Example 21, you will notice that the reason all candidates but C were eliminated is that C is the only possible conditioning property that was both present in occurrence * (where E was present) and absent in occurrence 1 (where E was absent). All other possible conditioning properties present in occurrence * were also present in occurrence 1, where E was absent, and thus were eliminated. It follows that all the possible conditioning properties that were absent in occurrence * were also absent in occurrence 1 (except for $\sim C$). In other words, there was only one change in the presence or absence of the possible conditioning properties from occurrence * to occurrence 1: the change from C being present and $\sim C$ being absent in occurrence * to C being absent and $\sim C$ being present in occurrence 1. This one change in the possible conditioning properties corresponds to the change in the conditioned property: E is present in occurrence * and absent in occurrence 1. When both simple properties and their negations are allowed to be possible conditioning properties in the method of difference, this is the only way in which one eliminating occurrence can eliminate all but one of the possible conditioning properties. This rather special case of the method of difference is what Mill describes as "the method of difference." However, Mill's view of the method of difference was too narrow, for, as has been shown, the

method has application when several eliminating occurrences, rather than just one, narrow down the field.

The method of difference may be expanded, in exactly the same way as in the inverse method of agreement, by allowing conjunctions of simple properties as possible conditioning properties. Simply remember that we start with a particular occurrence, occurrence *, in which *E* is present. The candidates will then be all the possible conditioning properties that are present in occurrence *. We then look for occurrences where *E* is absent, so that some of the candidates can be eliminated. A candidate is eliminated if it is present in an occurrence where *E* is absent, since a sufficient condition for *E* cannot be present when *E* is absent. If all candidates but one are eliminated, we can conclude that, if one of the possible conditioning properties present in occurrence * is a sufficient condition for *E*, then the remaining candidate is that sufficient condition. But, as in the direct and inverse methods of agreement, it is not always possible to narrow down the field to one candidate. More than one sufficient condition for *E* may be present in occurrence *. When a man is simultaneously beheaded, shot through the heart, and exposed to a lethal dose of nerve gas, several sufficient conditions for death are present. On the other hand, the eliminating occurrences might eliminate all the candidates. This would show that the list of possible conditioning properties did not include a property that was both present in occurrence * and a sufficient condition for *E*. In such a case other factors that were present in occurrence * must be sought and included in a new, expanded list of possible conditioning properties.

Exercises

Consider the following example:

	Possible conditioning properties								Conditioned property <i>E</i>
	Simple				Complex				
	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	$\sim A$	$\sim B$	$\sim C$	$\sim D$	
Occurrence *:	P	A	A	P	A	P	P	A	P

1. What are the candidates?
2. Describe an eliminating occurrence that would eliminate all the candidates but one.
3. Describe an eliminating occurrence that would eliminate all the candidates.
4. Describe three eliminating occurrences, each of which would eliminate exactly one of the candidates.
5. What would you conclude if you observed the occurrence that you described in Exercise 2?

6. What would you conclude if you observed the occurrence you described in Exercise 3?
7. What would you conclude if you observed the three occurrences you described in Exercise 4?
8. What would you conclude if you observed all the occurrences that you described in Exercises 2 and 4? There are several correct answers to Exercises 2 and 4, and the answer to this question will depend on which ones you chose.

V.7. THE COMBINED METHODS. Sometimes a property is both a necessary and sufficient condition for another property. It has already been pointed out that in football the complex property “making a touchdown or making a field goal or making a conversion or making a safety” is both a necessary and sufficient condition for scoring. Medical authorities thought until recently that stoppage of the heart for more than a few minutes was both a necessary and sufficient condition for death. In elementary physics being acted on by a net force is both a necessary and sufficient condition for a change in a body’s velocity. Since there is a method for finding necessary conditions—the direct method of agreement—and two methods for finding sufficient conditions—the inverse method of agreement and the method of difference—they may be combined in order to find conditions that are both necessary and sufficient.

In Example 22 the direct and inverse methods of agreement are combined

Example 22

	Possible conditioning properties								Conditioned property <i>E</i>
	Simple				Complex				
	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	$\sim A$	$\sim B$	$\sim C$	$\sim D$	
Occurrence 1:	P	A	P	A	A	P	A	P	P
Occurrence 2:	A	P	P	P	P	A	A	A	P
Occurrence 3:	A	P	A	P	P	A	P	A	A
Occurrence 4:	P	A	A	A	A	P	P	P	A

into what is called the *double method of agreement*. Occurrence 1 eliminates *B*, *D*, $\sim A$, and $\sim C$ and occurrence 2 eliminates *A*, $\sim B$, $\sim C$, and $\sim D$ as necessary conditions for *E* in accordance with the *direct* method of agreement, for they are absent when *E* is present. We can conclude then from occurrences 1 and 2, by the direct method of agreement, that if one of the possible conditioning properties is a necessary condition for *E*, then *C* is that necessary condition. In accordance with the *inverse* method of agreement, occurrence 3

eliminates B , D , $\sim A$, and $\sim C$ and occurrence 4 eliminates A , $\sim B$, $\sim C$, and $\sim D$ as sufficient conditions for E , since they are present when E is absent. This again leaves only C . We can conclude from occurrences 3 and 4, by the inverse method of agreement, that if one of the possible conditioning properties is a sufficient condition for E , then C is that sufficient condition. Putting these results together leads to the conclusion that if one of the possible conditioning properties is both a necessary and a sufficient condition for E , then C is that property. However, a stronger conclusion may be drawn: If one of the possible conditioning properties is a necessary condition for E , and one of the possible conditioning properties is a sufficient condition for E , then one and the same possible conditioning property is both a necessary and sufficient condition for E , and that property is C .

The *joint method of agreement and difference*, which is illustrated in Example 23, combines the direct method of agreement and the method of difference. The first step is to apply the method of difference to Example 23. Occurrence * sets up as candidates for the sufficient condition for E those properties that are present in occurrence *, namely, A , C , $\sim B$, and $\sim D$. But occurrence 1 shows that neither A nor $\sim B$ nor $\sim D$ can be a sufficient condition for E , since they are all present when E is absent. This leaves only C .

Example 23

	Possible conditioning properties								Conditioned property E
	Simple				Complex				
	A	B	C	D	$\sim A$	$\sim B$	$\sim C$	$\sim D$	
Occurrence *:	P	A	P	A	A	P	A	P	P
Occurrence 1:	P	A	A	A	A	P	P	P	A
Occurrence 2:	A	P	P	P	P	A	A	A	P

Thus, we can conclude from occurrence * and occurrence 1 that if one of the possible conditioning properties present in occurrence * is a sufficient condition for E , then C is that sufficient condition.

Now let us apply the direct method of agreement to Example 23. Occurrence * may be used again since, in accordance with the direct method of agreement, it eliminates B , D , $\sim A$, and $\sim C$ as necessary conditions for E . Occurrence 2 further eliminates A , $\sim B$, and $\sim D$ as necessary conditions for E since they also are absent in an occurrence where E is present. This leaves only C . So from occurrence * and occurrence 2, by the direct method of agreement, we can conclude that if one of the possible conditioning properties is a necessary condition for E , then C is that necessary condition. Putting the

results of the method of difference and the direct method of agreement together leads to the conclusion that: If one of the possible conditioning properties present in occurrence * is a sufficient condition for E and if one of the possible conditioning properties is a necessary condition for E , then one and the same possible conditioning property that is present in occurrence * is both a necessary and sufficient condition for E , and that property is C .

In comparing the example of the joint method of agreement and difference with the previous example of the double method of agreement, note that occurrences *, 1, and 2 of Example 23 are the same, respectively, as occurrences 1, 4, and 2 of Example 22. Notice also that Example 22, using the double method of agreement, takes four occurrences to narrow down the field to C , while Example 23, using the joint method of agreement and difference, takes only three occurrences. Does this mean that the joint method of agreement and difference is, in some way, a more efficient method than the double method of agreement? Not at all. Less occurrences are needed in Example 23 than in Example 22 because the conclusion drawn from Example 23 is weaker than that drawn from Example 22. From Example 23 we may conclude that if one of the possible conditioning properties *which is present in occurrence ** is a sufficient condition for E and one of the possible conditioning properties is a necessary condition for E , then C is both the necessary and the sufficient condition. If we want to remove the restriction “which is present in occurrence *” then the extra occurrence that appears in Example 22 is needed, and the double method of agreement must be used. Consequently, from Example 22 the stronger conclusion may be drawn that if one of the possible conditioning properties is a sufficient condition for E and one of the possible conditioning properties is a necessary condition for E , then C is both the necessary and the sufficient condition.

Whether the joint method of agreement and difference or the double method of agreement is chosen depends on what previous knowledge we have. Suppose we have observed an occurrence and have good reason to believe that one of the possible conditioning properties which is present in that occurrence is a sufficient condition for E . We would then designate that occurrence as occurrence * and proceed with the joint method of agreement and difference. If, however, we had good reason to believe only that one or another of the possible conditioning properties is a sufficient condition for E , we would have to rely on the double method of agreement. The combined methods are equally efficient, but they are appropriate in different circumstances.

The combined methods may be expanded to include other complex properties (disjunctions and conjunctions of simple properties and the negations of simple properties), but a discussion of these more involved forms of

Mill's methods belongs in more advanced texts. Remember, however, that everything that has been said about Mill's methods, and everything that can be said about their more involved forms, rests on two simple principles of elimination:

- i. A necessary condition for *E* cannot be absent when *E* is present.
- ii. A sufficient condition for *E* cannot be present when *E* is absent.

These two principles are more important to remember than Mill's methods themselves, and they should always be borne in mind when a mass of data is being analyzed.

Exercises

Suppose you have observed the following occurrences:

	Possible conditioning properties								Conditioned property <i>E</i>
	Simple				Complex				
	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	$\sim A$	$\sim B$	$\sim C$	$\sim D$	
Occurrence 1:	P	P	A	A	A	A	P	P	P
Occurrence 2:	P	A	A	A	A	P	P	P	A
Occurrence 3:	A	P	P	P	P	A	A	A	P

1. Suppose you know that one of the possible conditioning properties is a necessary condition for *E*. Which one is it? What occurrences did you use and which of Mill's methods did you apply?
2. Suppose you know that one of the possible conditioning properties which is present in occurrence 1 is a sufficient condition for *E*. Which one is it? What occurrences did you use and which one of Mill's methods did you apply?
3. Suppose you know that one of the possible conditioning properties is a necessary condition for *E* and that one of the possible conditioning properties which is present in occurrence 1 is a sufficient condition for *E*. Do you know whether one possible conditioning property is both a necessary and sufficient condition for *E*? If so, which one is it and which one of Mill's methods did you use?
4. Suppose you know that one of the possible conditioning properties is a necessary condition for *E*. You also know that one of the possible conditioning properties is a sufficient condition for *E*, but you do not know whether it is a property that is present in occurrence 1. Furthermore, you have observed an additional occurrence:

	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	$\sim A$	$\sim B$	$\sim C$	$\sim D$	<i>E</i>
Occurrence 4:	A	A	P	P	P	P	A	A	A

Do you know whether one possible conditioning property is both a necessary and a sufficient condition for E ? If so, which one is it and which one of Mill's methods did you use?

5. Suppose you had only observed occurrences 1 and 2 but you knew that one of the possible conditioning properties was both a necessary and a sufficient condition for E . Using the two principles of elimination, can you tell which one it is?

V.8. THE APPLICATION OF MILL'S METHODS. The conclusions we drew from various applications of Mill's methods always began with phrases such as "If one of the possible conditioning properties is a necessary condition for E . . . ," or "If one of the possible conditioning properties which is present in occurrence * is a sufficient condition for E . . . ," and so on. It would seem that our confidence that Mill's methods have found a necessary condition, or a sufficient condition, or a necessary and sufficient condition depends on our confidence that the list of possible conditioning properties contains the requisite kind of condition. But how can we be sure that this list does contain the type of condition being sought?

One suggestion might be to include *all* properties as possible conditioning properties and to rely on some principle of the uniformity of nature to guarantee that each conditioned property has some necessary and some sufficient conditions. There are many things wrong with this suggestion, but the most practical objection is that there are simply too many properties to take into account. Even if we are interested only in the properties that are present in a given occurrence, as in the method of difference, not all properties that are present can be considered. In any occurrence there are countless properties present. When you sneeze, there are hundreds of chemical reactions going on within your body: various electrical currents are circulating in your nerve fibers; you are being bombarded by various types of electromagnetic radiation; diverse happenings, great and small, surround you. It would be an impossible task to measure and catalog all these things and eliminate them, one by one, by Mill's methods in order to find a sufficient condition for sneezing.

For Mill's methods to be of any use, there must be some way of ascertaining what factors are likely to be relevant to the conditioned property in which we are interested; there must be some way of setting up a list of reasonable length of possible conditioning properties which probably contains the necessary or sufficient conditions being sought. The only way to do this is to apply inductive logic to a previously acquired body of evidence. Mill's methods are of no use unless we already have some inductive knowledge to guide us in setting up the list of possible conditioning properties.

Mill's methods are useful in science, but their usefulness depends on inductively based judgments as to what factors are likely to be relevant to a

given conditioned property. Of course, inductively based judgments are not infallible. We may be mistaken in believing that the list of possible conditioning properties contains a necessary or a sufficient condition. The occurrences observed may eliminate all of the possible conditioning properties. If this happens, all the evidence at our disposal must be reexamined, and perhaps new evidence must be sought, in order to find new properties that are probably relevant to the conditioned property under investigation. Once the inductive judgment has been made as to which additional properties must be considered, Mill's methods may then be reapplied. In the search for necessary and sufficient conditions, Mill's methods are part of the picture, but they are not the whole picture. The most basic, and least understood, part of the process is the setting up of lists of possible conditioning properties.

Imagine the following scientific experimental situation in which Mill's methods might be applied. Suppose we have three new drugs that hold promise for the cure of a hitherto incurable disease: drug *A*, drug *B*, and drug *C*. We administer various combinations of these drugs and note whether the patient is cured or not. The results are tabulated in Example 24, where "*A*" means drug *A* has been administered, "*B*" means drug *B* has been administered, and "*C*" means drug *C* has been administered. "*E*" means the patient

Example 24

	Possible conditioning properties			Conditioned property
	<i>A</i>	<i>B</i>	<i>C</i>	<i>E</i>
Occurrence 1:	P	P	A	P
Occurrence 2:	P	A	P	P
Occurrence 3:	A	P	A	A
Occurrence 4:	A	A	P	A
Occurrence 5:	A	A	A	A

has been cured. Occurrence 5 represents the cases of all the previous patients who had not taken any of these drugs and who had not been cured. The cases of patients to whom various combinations of the new drugs have been administered are tabulated under occurrences 1 through 4. Example 24 constitutes a case of the double method of agreement and it warrants the conclusion that if one of the possible conditioning properties is a necessary condition for *E*, and one of the possible conditioning properties is a sufficient condition for *E*, then *A* is both the necessary and the sufficient condition for *E*. Thus, these results lead to the conclusion that, to the best of our knowledge, the administration of drug *A* is both a necessary and a sufficient condition for a cure of the disease.

But suppose that someone tries other combinations of the drugs and gets the results shown in Example 25. Occurrence 7 shows that *A* is not a sufficient condition for *E*. Therefore, if we wished to find a sufficient condition for *E*, we would have to expand our list of possible conditioning properties. Suppose now that biochemical theory suggests that there may be a chemical interaction

Example 25

	Possible conditioning properties			Conditioned property
	<i>A</i>	<i>B</i>	<i>C</i>	<i>E</i>
Occurrence 6:	P	A	A	P
Occurrence 7:	P	P	P	A
Occurrence 8:	A	P	P	A

if the three drugs are administered simultaneously, and that such a chemical interaction might cancel out their effectiveness against the disease. The sum of our observations would then suggest that what may be happening in occurrence 7 is that drugs *B* and *C* are interacting and preventing drug *A* from curing the disease. (We could imagine different occurrences that would suggest that two drugs are effective only in combination.) Drug *A* seems to be effective when taken alone (occurrence 6), or when taken with *B* but without *C* (occurrence 1), or when taken with *C* but without *B* (occurrence 2). This suggests that the complex property, $A \& \sim(B \& C)$, that is, taking drug *A* but not in conjunction with both drug *B* and drug *C*, is really the sufficient condition for *E*. If we were to add this complex property to our list of possible conditioning properties, and use all eight occurrences, we would find that it is then the only possible conditioning property that is present whenever *E* is present and absent whenever *E* is absent. In this way we can reapply the double method of agreement to an enlarged set of possible conditioning properties, in the face of additional occurrences, in order to revise our conclusion and make it more sophisticated.

But we might not be finished even at this point. Suppose that another researcher were to point out that all our tests have been made on patients in whom the disease was at an early stage, and that many diseases are more easily cured in their early stages than in their advanced stages. This would suggest that our complex property only appears to be a sufficient condition for *E*, because we have not tested our drugs on advanced cases of the disease. What we would now have to do is to take this additional factor into account in our list of possible conditioning properties. We could introduce a new property, *D*, which is said to be present when the disease is in its advanced stages and

absent otherwise. In all the occurrences where drugs have been administered so far, D has been absent. Now we would have to find various occurrences where D was present. That is, we should administer various combinations of drugs to patients in advanced stages of the disease and note the results. If the treatment that effected a cure before were to still effect a cure, then we would not have to revise our belief that $A \& \sim(B \& C)$ is a sufficient condition for E . But if our treatment failed in advanced cases, then we might have to say that the sufficient condition for being cured is having the disease in an early stage and receiving the correct combination of drugs. That is, we would have to say that $A \& \sim(B \& C)$ is not a sufficient condition for E , but that $\sim D \& A \& \sim(B \& C)$ is a sufficient condition for E .

We could imagine an endless stream of developments which might force us to add more and more complex and simple properties to our list of possible conditioning properties and to continually reevaluate our results. Someone might develop a new drug that effects a cure in the absence of drug A and thus show that A is not a necessary condition for E . Additional research might suggest other factors that might be relevant and whose relationship to E we might wish to examine. It is by such a process that Mill's methods, in conjunction with a continual search for new occurrences, and new relevant possible conditioning properties, contribute to the growth of scientific knowledge.

Suggested reading

Georg Henrik von Wright, *A Treatise on Induction and Probability* (Patterson, NJ: Littlefield, Adams & Co., 1960).

V.9. SUFFICIENT CONDITIONS AND FUNCTIONAL RELATIONSHIPS. The preceding treatment of Mill's methods in terms of necessary and sufficient conditions proceeded entirely in qualitative terms. One may wonder what relevance, if any, that discussion has for sciences which have moved from qualitative to quantitative language. Here, ascriptions of cause or statements of necessary and sufficient conditions have been replaced by functional relationships expressed by mathematical equations. The basic logic of the situation, however, is not as different as it may seem. An equation expressing a functional relationship between physical quantities is tantamount to not one but an infinite number of statements to the effect that one physical property is a sufficient condition for another.

To understand this, we must look first at the relation between properties and physical quantities. Consider a physical quantity, for example, temperature, as measured on a given scale (e.g., degrees Kelvin). We make a factual claim about a state of a physical system when we say that its temperature

(in degrees Kelvin) has a certain value. Temperature (so measured) is thus a *relation* between states of physical systems and (non-negative real) numbers. This is to say no more than:

For every non-negative real number, x , there is associated a unique *physical property, having the temperature x in degrees Kelvin.*

A physical quantity can thus be seen as not one but rather an infinite family of physical properties. The properties in such a family are *mutually exclusive* (a physical system cannot have two different temperatures at the same time) and *jointly exhaustive* (a physical system in a given state must have some temperature or other) over the states of the appropriate type of physical system (the concept of temperature has no meaning when applied, for instance, to the nucleus of an atom). An appropriate set of real numbers serves as a *fruitful filing system* for the physical properties in such a family. We can thus say that:

A *physical quantity* is a family of physical qualities, mutually exclusive and jointly exhaustive over the states of the intended class of physical systems, *indexed* by some set of real numbers.

We said that the indexing of the physical quantities by the index set of real numbers forms a *fruitful* filing system. It is fruitful just in that it, together with the filing systems of other physical quantities, enables us to formulate physical laws in terms of mathematical equations. To see how this works, let us consider a few simple equations. First the equation $x = 2y$. This equation gives concise expression to an infinite number of statements, of which a few are:

If y is 0, x is 0.

If y is 1, x is 2.

If y is $3\frac{1}{2}$, x is 7.

In general, for each value of y , the equation correlates a unique value of x . We give expression to this fact by saying that here x is a *function* of y . This equation also makes y a function of x , since for every value of x it correlates a unique value for y (i.e., $\frac{1}{2}x$). It does not always follow, however, that if x is a function of y , y is a function of x . Consider the equation $x = y^2$. Here, each value of y determines a unique value of x , but the converse is not true. If x is +4, y may be either +2 or -2. Thus, x is a function of y , but y is not a function of x .

What does this mean in physical terms when the variables of the equation represent physical quantities? If the variables represent physical quantities measured on fixed scales (e.g., temperature Kelvin) then, as we have seen, each numerical value of a variable represents a physical quality (e.g., having a

temperature of 10 degrees Kelvin). If a physical quantity, Q_2 , is a *function* of another, Q_1 , then for every value v_i of Q_1 there is a uniquely determined value v_j of Q_2 such that whenever a physical system has v_i of Q_1 it has v_j of Q_2 . That is, having v_i of Q_1 is a *sufficient condition* for having v_j of Q_2 . Thus, in terms of our view of physical quantities as families of physical qualities, we may say that:

If Q_2 is a function of Q_1 , then for every member of the family Q_1 , there is some member of the family Q_2 for which it is a sufficient condition.

This generalizes in a straightforward way to functions of several variables. For example, consider the ideal gas law:

$$V = k \frac{T}{P}$$

where V is volume, T is temperature, and P is pressure. k is called a *system-dependent constant* since it varies from system to system (e.g., different balloons filled with gas) but remains constant over different states of the same system (e.g., heating a balloon, or submerging it to 200 fathoms). The equation establishes V as a function of k , T , and P in that each triple of values for k , T , and P uniquely determines a value for V . Thus, in conjunction:

$$k = 1 \text{ and } T = 100 \text{ and } P = 50$$

is a sufficient condition for $V = 2$.

In general, we can say that if a quantity, Q_0 , is a function of several others, $Q_1 \dots Q_n$, then for every conjunctive physical property which contains as conjuncts just one member from each of the families $Q_1 \dots Q_n$, there is a member of the family Q_0 for which it is a sufficient condition. Mathematical equations establishing functional relations between physical quantities thus allow succinct expression of extremely rich claims about sufficient conditions.

Since we are still, at basis, dealing with sufficient conditions, the fundamental principles that we used to analyze Mill's methods must still apply, although in a slightly more complicated way. Instead of one conditioned property, we have the family of properties comprising a physical quantity (the *dependent variable*). We must find a way for establishing, for each member of this family, a list of possible conditioning properties. This is a two-stage process. The first stage is to construct a list of physical quantities whose values are likely to be relevant in determining the value of the dependent variable. We can call this our list of *independent variables*. The second stage is to construct a list of likely looking *functions* which make the dependent variable a function of our independent variables. For each conditioned property (value of the dependent variable), each of these functions determines one or more complex properties which are possible sufficient conditions for it. For example, consider the two

functions:

$$(i) V = k \frac{T}{P} \text{ and } (ii) V = k \frac{T^2}{P}$$

Function (i) would make the following, among others, sufficient conditions for $V = 1$:

$$k = 1 \text{ and } T = 1 \text{ and } P = 1$$

$$k = 1 \text{ and } T = 2 \text{ and } P = 2$$

Function (ii) also makes

$$k = 1 \text{ and } T = 1 \text{ and } P = 1$$

a sufficient condition for $V = 1$, but disagrees with function (i) in making

$$k = 1 \text{ and } T = 2 \text{ and } P = 2$$

a sufficient condition for $V = 2$ rather than for $V = 1$. Since $V = 1$ and $V = 2$ are mutually exclusive physical properties, in any occurrence at least one of them must be absent. We have, then, only to look for an occurrence where $k = 1$ and $T = 2$ and $P = 2$ in order to eliminate either the hypothesis that function (i) gives the correct sufficient conditions or the hypothesis that condition (ii) does. The method operative here is thus a straightforward application of the inverse method of agreement. The only new twist is that we have a family of conditioned properties which are mutually exclusive, so that if one is present, the rest must be absent.

Although the mechanism for the elimination of the proposed functions is quite clear here, the process for setting up the list of possible functions is, as before, quite murky. Such is to be expected, since the first process is really deductive whereas the latter is genuinely inductive. Nevertheless, the analysis given so far enables us to shed some light on the process of isolating relevant independent variables.

Remember that if P is a sufficient condition for C , then so is P and Q (and, of course, P and not- Q). There are, then, some very cumbersome sufficient conditions around, but obviously the most knowledge is gained by finding the shortest ones. In terms of functions, this means that if a quantity, Q_0 , is a function of another, Q_1 , it is also a function of Q_1 and Q_2 for any quantity Q_2 . Again, the most interesting functions are stated in terms of the minimum number of variables needed to do the job. Suppose we start with a list of likely conditioning quantities Q_1, Q_2, \dots, Q_n , vary them independently and find that one of them, say Q_1 , doesn't make any difference in the conditioned quantity. That is, for different fixed combinations of values for $Q_2 \dots Q_n$, the value of the conditioned

quantity remains the same when the value of Q_1 is varied. Then we have good reason to believe that values of Q_1 would be excess fat in statements of sufficient conditions for the conditioned quantity. In other words, we have good (inductive) reason for believing that Q_1 is not a *relevant* variable. This process of reducing our list of conditioning quantities is called *isolating the relevant variables*.

The second stage in setting up the conditioning properties was to select a likely list of functions which make the dependent variable a function of the remaining independent variables. It is difficult to say anything very informative about the selection process. Sometimes we are guided by the sorts of functional relationships which have already been found to hold in similar physical situations, but it is difficult to say what "*similar*" means here. Sometimes we seem to be guided by considerations of *simplicity* of the expressions which designate the function. But, except in certain special cases, simplicity is a highly elusive concept.

To the question, "Why do we have this two-stage method of formulating statements of sufficient conditions?" we have already seen a relatively superficial (though correct) answer. That is, this method allows the succinct formulation of statements of such power and scope that they would otherwise exceed the resources of our language. But there is another, more profound, reason: that *the two stages are not inductively independent*. To see what I mean, consider a new physical quantity, T^* , which is just like temperature Kelvin except that when temperature in degrees Kelvin equals 10, $T^* = 90$ and when temperature in degrees Kelvin equals 90, $T^* = 10$. T^* comprehends exactly the same physical qualities as temperature, and uses the same set of numbers to index them, but the filing system is different. Imagine now, formulating the ideal gas law in terms of T^* rather than temperature. Suddenly, the simple becomes more complex (and by the same token, the complex can become more simple). Our basic physical magnitudes come to us, then, not simply as artless vehicles for the expression of factual claims, but rather as bearers of inductive wisdom. It is they, rather than other families of physical qualities, which have found expression in our language precisely because they have been found to enter into simply expressible functional relationships in a wide variety of physical contexts. Language comes to us with inductive commitments—commitments so deeply ingrained that it is easy to overlook them. But, as we learned from Goodman in the last chapter, overlooking them leads to an excessively simple-minded view of the nature of the inductive process.

V.10. LAWLIKE AND ACCIDENTAL CONDITIONS. In Section V.5 we defined a sufficient condition as follows:

A property, F , is a sufficient condition for a property, G , if and only if *whenever F is present, G is present*.

Hence, the following are all legitimate statements of sufficient conditions:

1. *Being a brother* is a sufficient condition for *being male*.
2. *Being over six feet tall* is a sufficient condition for *being over five feet tall*.
3. *Being pure water at a pressure of one atmosphere and a temperature of 100 degrees centigrade* is a sufficient condition for *boiling*.
4. *Having an inertial mass of one kilogram* is a sufficient condition for *having a gravitational mass of one kilogram*.
5. *Eating dinner at my house on January 12, 1999*, is a sufficient condition for *being under seven feet tall*.
6. Nelson Goodman had only dimes, quarters, and half-dollars in his pocket on VE day, so *being a coin in Nelson Goodman's pocket on VE day* is a sufficient condition for *being made of silver*.

It is obvious from these examples that there are strikingly different *grades* of sufficiency. The sufficiency of the condition in Example 1 is due, in a most transparent way, to the meanings of the terms involved. We may call it an *analytic* sufficient condition. Example 2 also depends on the concepts involved, rather than the way we find the world to be, so we shall also call it an analytic sufficient condition. However, Example 2 should remind us that an account of analyticity is not always so easy to give as in the case of Example 1. It is not clear whether Example 3 is analytic or not. Is having a certain boiling point part of what we mean by being pure water? Is the boiling point of pure water involved in the definition of the centigrade scale? If the answer is yes to either of these questions, then we may have an analytic sufficient condition. If we have independent definitions of pure water and temperature centigrade then Example 3 states a sufficient condition which is informative about the way the world operates. Actual practice tends to shift from one set of meanings to another depending on what is most convenient to the occasion. Thus, actual practice does not provide an unambiguous answer as to whether Example 3 is an analytic sufficient condition or not. Such a semantically muddled state of affairs is common in human language, ordinary and scientific, and in such cases an unambiguous answer is only to be had by making a *decision* as to how the words are to be used on a particular occasion. Examples 4, 5, and 6 are all synthetic, but 4 is clearly different in kind from 5 and 6. Example 4 states a condition which is sufficient *by virtue of physical law*. In this respect it resembles Example 3, when Example 3 is interpreted as a synthetic statement. Examples 5 and 6, however, state conditions which are sufficient *simply by happenstance*. It simply *happened* that no one over six feet tall came to my house to eat dinner that

day. It just *happened* that Nelson Goodman had no pennies in his pocket on VE day. We say that these truths are *accidental* rather than *lawlike*.

Although each of the distinctions between grades of sufficiency raises important and interesting questions, we shall focus here on the last one: *the distinction between accidental and lawlike sufficient conditions*. This is an important distinction for inductive logic. The establishment of either sort of sufficient condition may be an inductive affair, but the roles they play are so different that one would suspect that inductive logic should treat them differently. It is lawlike sufficient conditions which make up the body of science. Statements of accidental sufficient conditions like 5 and 6 may, like any other factual statement, set up the application of a scientific theory, but they never form part of such a theory itself. The examples that we used to illustrate Mill's methods were, accordingly, all examples of lawlike sufficient conditions (or necessary conditions—all the distinctions being made here obviously apply to necessary conditions also).

Now, does this imperil our analysis of Mill's methods? Not at all. Our analysis depended only on the principle that a *sufficient condition cannot be present when the conditioned property is absent*. This principle of elimination follows from the definition of *sufficient condition* and thus holds for *all* sufficient conditions; accidental, lawlike, or analytic (likewise for necessary conditions). The story that we have told about Mill's methods is nothing but the truth. It is, however, far from the whole truth.

Mill's methods apply to analytic sufficient conditions, but eliminating other conditioning properties is surely not the most efficient way to arrive at *being a brother* as a sufficient condition for *being male*. It would be a hopelessly incompetent mathematician or logician who relied on the experimental method for his theorems. A scientist, on the other hand, would not even want an analytic sufficient condition in his list of conditioning properties, for knowing an analytic sufficient condition gives us no information about the way the world behaves. Neither would a scientist want a property like *being a coin in Nelson Goodman's pocket on VE day* on his list of possible sufficient conditions for *being composed of silver*. He knows that the overwhelming likelihood is that if this turns out to be a sufficient condition, it will turn out to be an accidental sufficient condition.

This raises two questions:

- (A) How do we distinguish lawlike from accidental sufficient conditions?
- (B) Why is it that lawlike conditions find a place in the body of science whereas accidental ones do not?

The answer to the second question, if it is to have any philosophical importance, must flow from considerations of the function of scientific law. And

anything better than an *ad hoc* answer to the first question must flow from a satisfactory answer to the second.

Looking at Examples 1 through 6 it is easy to conjecture that the difference between accident and law is the difference between part and whole; that laws are truths about the whole universe, throughout space and time, whereas truths which are about restricted parts of it (e.g., Nelson Goodman's pocket for the specified period of time) may be accidental.

Such a view has its attractions. Surely the most striking examples of accidental conditions stem from generalizations of spatio-temporally limited scope. The preoccupation of science with *lawlike* sufficient conditions is neatly explained by the *universality* of science. Science is concerned with patterns which recur throughout the universe, rather than with gossip about a particular spatio-temporal region. This concern flows from the essential pursuits of science: *explanation* and *prediction*. Science always explains an event by showing it, in some way, to be an instance of a general pattern, rather than just a freak occurrence. As for prediction, our generalizations about Nelson Goodman's pocket are obviously not very powerful predictive instruments, because they don't cover much territory and typically *we don't know about the sufficient conditions until we have already covered the territory!* The contents of Goodman's pocket at the time in question had to be completely surveyed before confidence could be placed in our statements of sufficient conditions. Given such a complete survey, there is nothing left for them to predict. Since no complete survey of the universe is possible, generalizations about it must be known, if at all, while there is still predictive life left in them.

Suppose we accordingly try to define a law as a true generalization which does not name specific times, places, or individuals. Isn't it possible that even the general descriptive machinery we have left may pick out a small finite class of objects? For instance, isn't it possible that a description of Nelson Goodman's pocket on VE day down to the finest detail, down to the trajectories of subatomic particles, could be so specific without containing names for times, places, or individuals, that the only thing in the whole universe which would answer to it would be Nelson Goodman's pocket on VE day? Then, according to our definition, *being a coin enclosed in such a structure* would have to be a lawlike sufficient condition for *being composed of silver*. But it is clearly accidental. In fact, it is *doubly* accidental, for it would be something of an accident that Goodman's pocket would be the *only* structure in the universe answering the description in question.

What has gone wrong? One natural line of thought is to conjecture that the trouble lies in defining spatio-temporal limitation of scope *via* the terms in which the generalization is couched, rather than by the objects to which it refers. Why not say that a law is a true generalization which does not refer to

any spatio-temporally limited (or alternatively, to any finite) class of objects?
Then

All coins enclosed in a structure of type I are composed of silver.

would fail to be a law, even if true, if Nelson Goodman's pocket on VE day constituted the only structure of type I.

But wait! Why do we assume that this generalization is *only about* coins enclosed in structures of type I? To be sure, if we know that a certain object is such a coin, we know that it is *crucial* to the generalization. It is crucial in that, if it turns out not to be silver, it falsifies the generalization. But if we know of another object that it is *not silver*, then similarly we know it to be crucial to the generalization. It is crucial in that if it turns out to be a coin enclosed in a structure of type I, it will falsify the generalization. In all fairness, then, we ought to allow that our generalization is *also about* objects not composed of silver. Another way to put the same point is to note that *being a coin enclosed in a structure of type I* is a sufficient condition for *being composed of silver* just in case *not being composed of silver* is a sufficient condition for *not being a coin enclosed in a structure of type I*. Thus, our generalization refers both to coins enclosed in a structure of type I and to objects not made of silver. A little further discussion might convince us that it refers to everything else as well. But we have already gone far enough to see that we are on the wrong track. The class of objects referred to by our generalization is no longer spatio-temporally limited or finite.

The attempt to locate the dividing line between accidental and lawlike sufficient conditions in considerations of spatio-temporal limitation of scope seems to have come to a dead end. And if the problems so far raised for this approach are not enough, consider the following example (due to Professor Carl Hempel):

7. It seems likely that there is no body of pure gold in the universe whose mass equals or exceeds 100,000 kilograms. If so, *being a body composed of pure gold* is a sufficient condition for *having a mass of less than 100,000 kilograms*.

Note that our belief in the foregoing is quite compatible with the belief in an infinite universe strewn with an infinite number of bodies composed of pure gold *and* an infinite number of bodies having a mass of more than 100,000 kilograms.

Yet, for all that, we would consider such a sufficient condition not a matter of law but rather an accident—a “global accident,” if you please. A world might obey the same physical laws as ours, and yet contain huge masses of

gold just as a world with the same laws might have particles moving with different velocities. What then is the difference between such global accidents and true laws?

A major difference seems to be that laws are crucial to the structure of our whole view of the world in a way that accidental generalizations are not. If astronomers announced the discovery of a large interstellar body of pure gold, we would find it surprising, but not disturbing. It would arouse our curiosity and our desire for an explanation. The falsification of a physical law, on the other hand, would call for revision throughout a whole system of beliefs and would destroy a whole tissue of explanations. Tranquility is restored only when a new law reorders the chaos.

One way of viewing this difference is to regard laws not merely as beliefs about the world but, in addition, as contingent rules for changing our beliefs under the pressure of new evidence. I now believe my harpsichord to be safely at rest in Goleta, California. If I would learn that a huge net force were being applied to it, say by a hurricane, I would revise that belief in accordance with the laws of physics and fear for its safety. It is not surprising that our system of beliefs should suffer a greater disturbance when rules normally used for changing beliefs must themselves be revised in comparison to situations in which they remain intact.

It is of course true that an accidental generalization, or indeed any statement we believe to be true, plays a role in determining how we change our beliefs under the pressure of new evidence. But the role appears to be different from that played by laws. Let us compare. If I am told on good authority that a new heavenly body of mass greater than 100,000 kilograms has been discovered, I will assume that this is not a mass of pure gold. But if later investigations convince me that it is, in fact, pure gold I will not (as in the case of the harpsichord) revise my previous belief and conclude that it must really weigh less than 100,000 kilograms. Rather, I will give up my belief that all bodies composed of pure gold have a mass of less than 100,000 kilograms.

But consider the coins in Goodman's pocket on VE day. You may have extremely strong grounds for believing that all these coins are silver; say you were present at the time, observed Goodman turning his pockets inside out yielding just three coins, that you tested them chemically and found them to be silver, and so on. Now if someone convinces you that he has one of the coins in Goodman's pocket on VE day you will assume that it is silver. If he then fishes out a copper cent, exclaiming "This is it!" you will revise your opinions both as to the origin of the coin and the veracity of its possessor. Thus, the fact that a belief is held with extreme tenacity does not guarantee that it is functioning as a law, even though laws are typically more stable and central pieces of our intellectual equipment than mere factual judgments.

Perhaps the matter can be clarified if we consider the farfetched sort of circumstances under which *being a coin in Goodman's pocket on VE day* would be considered a lawlike sufficient condition for *being made of silver*. Suppose that our grounds for believing that all these coins are silver is that we know Goodman's pockets had a certain physical structure; that this structure sets up a force field which allows only silver articles to enter (or, more fancifully, one which transmutes all other elements to silver). If the suggestion that laws have a special place as rules for changing beliefs has any currency then we should be able to find differences between such application of this sufficient condition in the lawlike and accidental cases.

Human observation is fallible, and there is some likelihood, however small, that we missed a coin when examining Goodman's pockets. (Perhaps it stuck in that little corner of the pocket that doesn't turn inside out; perhaps this pair of pants had *two* watch pockets; etc.) Suppose that there is such a coin. If we are suddenly informed of its existence, what are we to think of its composition? *In the accidental case* we have no clue what to think and if it turns out to be copper we will not find this disturbing over and above our initial disturbance at having missed it. *In the lawlike case*, the inference rule still applies and we will be quite confident that it is silver. If it turns out to be copper, we will hasten to reexamine the structure of Goodman's pocket and if we find no fault in our previous beliefs about it, we will be forced to seek for some revised physical theory to account for these facts.

Laws then, *do* seem to have a special status as rules for revising our beliefs. This special status is perhaps most easily seen in our reasoning about *what might have been*. We will say, of a glass of pure, cold water (at a pressure of one atmosphere):

(A) If this water *had been* heated to 100 degrees centigrade it *would have* boiled.

because we believe that:

(B) All pure water at a pressure of one atmosphere and a temperature of 100 degrees centigrade boils.

is a *law*. (A) is said to be a *counterfactual* conditional since, as the water has *not* been heated, its if-clause is contrary to fact. The law (B) is said to *support* the counterfactual condition (A). If we review our examples, we will find that laws support counterfactuals in a way that accidental generalizations do not. Suppose I have a box with an inertial mass of $\frac{3}{5}$ of a kilogram. I say without trepidation that if this box *had* had an inertial mass of 1 kilogram, it *would have* had a gravitational mass of 1 kilogram. But if a certain man is seven and a half feet tall I will certainly *not* say that if he had eaten dinner at my house on

January 12, 1999, he would have been under seven feet tall. I will say that if a net force had been applied to my harpsichord, it would have moved. But I will not say that if this penny had been in Goodman's pocket on VE day it would have been silver nor will I say that if Jupiter were made of pure gold it would have a mass of less than 100,000 kilograms.

Some metaphysicians have held that statements of what might have been are objective statements about parallel worlds or branches of time. Other thinkers hold that correct counterfactuals are fables constructed according to our contingent rules for changing beliefs. Reasoning about what might have been has value for them only as practice for reasoning about what might be.

It should be now clear that lawlike and accidental conditions are different, and you have some general indication of how they are different, but the specification of differences has not been precise. How exactly do laws function as contingent rules of inference? What are the rules for changing our beliefs about laws? Just what is needed for a law to support a given counterfactual? Despite an enormous amount of work there is, as yet, no generally satisfactory solution to these and related problems. They remain a major area of concern for the philosophy of science.

Suggested readings

Nelson Goodman, *Fact, Fiction and Forecast* (4th. ed.). (Cambridge, MA: Harvard University Press, 1983).

Douglas Stalker (ed.), *Grue! the new riddle of induction* (Chicago Open Court, 1994).

VI

The Probability Calculus

VI.1. INTRODUCTION. The theory of probability resulted from the cooperation of two eminent seventeenth-century mathematicians and a gambler. The gambler, Chevalier de Méré, had some theoretical problems with practical consequences at the dice tables. He took his problems to Blaise Pascal who in turn entered into correspondence with Pierre de Fermat, in order to discuss them. The mathematical theory of probability was born in the Pascal-Fermat correspondence.

We have used the word “probability” rather freely in the discussion so far, with only a rough, intuitive grasp of its meaning. In this chapter we will learn the mathematical rules that a quantity must satisfy in order to qualify as a probability.

VI.2. PROBABILITY, ARGUMENTS, STATEMENTS, AND PROPERTIES. The word “probability” is used for a number of distinct concepts. Earlier I pointed out the difference between inductive probability, which applies to arguments, and epistemic probability, which applies to statements. There is yet another type of probability, which applies to properties. When we speak of the probability *of* throwing a “natural” in dice, or the probability *of* living to age 65, we are ascribing probabilities to properties. When we speak of the probability *that* John Q. Jones will live to age 65, or the probability *that* the next throw of the dice will come up a natural, we are ascribing probabilities to statements. Thus, there are at least three different types of probability which apply to three different types of things: arguments, statements, and properties.

Luckily, there is a common core to these various concepts of probability: Each of these various types of probability obeys the rules of the mathematical theory of probability. Furthermore, the different types of probability are inter-related in other ways, some of which were brought out in the discussion of inductive and epistemic probability. In Chapter VI it will be shown how these different concepts of probability put flesh on the skeleton of the mathematical theory of probability. Here, however, we shall restrict ourselves to developing the mathematical theory.

The mathematical theory is often called the *probability calculus*. In order to facilitate the framing of examples we shall develop the probability calculus as it applies to *statements*. But we shall see later how it can also accommodate arguments and properties.

Remember that the truth tables for “ \sim ,” “ $\&$,” and “ \vee ” enable us to find out whether a complex statement is true or false if we know whether its simple constituent statements are true or false. However, truth tables tell us nothing about the truth or falsity of the simple constituent statements. In a similar manner, the rules of the probability calculus tell us how the probability of a complex statement is related to the probability of its simple constituent statements, but they do not tell us how to determine the probabilities of simple statements. The problem of determining the probability of simple statements (or properties or arguments) is a problem of inductive logic, but it is a problem that is not solved by the probability calculus.

Probability values assigned to complex statements range from 0 to 1. Although the probability calculus does not tell us how to determine the probabilities of simple statements, it does assign the extreme values of 0 and 1 to special kinds of complex statements. Previously we discussed complex statements that are *true* no matter what the facts are. These statements were called tautologies. Since a tautology is guaranteed to be true, no matter what the facts are, it is assigned the highest possible probability value.

Rule 1: If a statement is a tautology, then its probability is equal to 1.

Thus, just as the complex statement $s\vee\sim s$ is true no matter whether its simple constituent statement, s , is true or false, so its probability is 1 regardless of the probability of the simple constituent statement.

We also discussed another type of statement that is *false* no matter what the facts are. This type of statement, called the self-contradiction, is assigned the lowest possible probability value.

Rule 2: If a statement is a self-contradiction, then its probability is equal to 0.

Thus, just as the complex statement $s\&\sim s$ is false no matter whether its simple constituent statement, s , is true or false, so its probability is 0 regardless of the simple constituent statement.

When two statements make the same factual claim, that is, when they are true in exactly the same circumstances, they are logically equivalent. Now if a statement that makes a factual claim has a certain probability, another statement that makes exactly the same claim in different words should be equally probable. The statement “My next throw of the dice will come up a natural” should have the same probability as “It is *not* the case that my next throw of the dice will *not* come up a natural.” This fact is reflected in the following rule:

Rule 3: If two statements are logically equivalent, then they have the same probability.

By the truth table method it is easy to show that the simple statement p is logically equivalent to the complex statement that is its double negation, $\sim\sim p$, since they are true in exactly the same cases.

	p	$\sim p$	$\sim\sim p$
Case 1:	T	F	T
Case 2:	F	T	F

Thus, the simple statement “My next throw of the dice will come up a natural” has, according to Rule 3, the same probability as its double negation, “It is *not* the case that my next throw of the dice will *not* come up a natural.”

The first two rules cover certain special cases. They tell us the probability of a complex statement if it is either a tautology or a contradiction. The third rule tells us how to find the probability of a complex contingent statement from its simple constituent statements, if that complex statement is logically equivalent to one of its simple constituent statements. But there are many complex contingent statements that are not logically equivalent to any of their simple constituent statements, and more rules shall be introduced to cover them. The next two sections present rules for each of the logical connectives.

Exercises

Instead of writing “The probability of p is $\frac{1}{2}$,” we shall write, for short “ $\Pr(p) = \frac{1}{2}$.” Now suppose that $\Pr(p) = \frac{1}{2}$ and $\Pr(q) = \frac{1}{4}$. Find the probabilities of the following complex statements, using Rules 1 through 3 and the method of truth tables:

1. $p \vee p$.
2. $q \& q$.
3. $q \& \sim q$.
4. $\sim(q \& \sim q)$.
5. $\sim(p \vee \sim p)$.
6. $\sim\sim(p \vee \sim p)$.
7. $p \vee (q \& \sim q)$.
8. $q \& (p \vee \sim p)$.

VI.3. DISJUNCTION AND NEGATION RULES. The probability of a disjunction $p \vee q$ is most easily calculated when its disjuncts, p and q , are *mutually exclusive* or inconsistent with each other. In such a case the probability of the disjunction can be calculated from the probabilities of the disjuncts by means of the *special disjunction rule*. We shall use the notation introducing the exercises at the end of the previous section writing “The probability of p is x ” as: “ $\Pr(p) = x$.”

Rule 4: If p and q are mutually exclusive, then $\Pr(p \vee q) = \Pr(p) + \Pr(q)$.

For example, the statements “Socrates is both bald and wise” and “Socrates is neither bald nor wise” are mutually exclusive. Thus, if the probability that Socrates is both bald and wise is $\frac{1}{2}$ and the probability that Socrates is neither bald nor wise is $\frac{1}{4}$, then the probability that Socrates is either both bald and wise *or* neither bald nor wise is $\frac{1}{2} + \frac{1}{4}$, or $\frac{3}{4}$.

We can do a little more with the special alternation rule in the following case: Suppose you are about to throw a single six-sided die and that each of the six outcomes is equally probable; that is:

$$\Pr(\text{the die will come up a 1}) = \frac{1}{6}$$

$$\Pr(\text{the die will come up a 2}) = \frac{1}{6}$$

$$\Pr(\text{the die will come up a 3}) = \frac{1}{6}$$

$$\Pr(\text{the die will come up a 4}) = \frac{1}{6}$$

$$\Pr(\text{the die will come up a 5}) = \frac{1}{6}$$

$$\Pr(\text{the die will come up a 6}) = \frac{1}{6}$$

Since the die can show only one face at a time, these six statements may be treated as being mutually exclusive.¹ Thus, the probability of getting a 1 or a 6 may be calculated by the special disjunction rule as follows:

$$\Pr(1\vee6) = \Pr(1) + \Pr(6) = \frac{1}{6} + \frac{1}{6} = \frac{1}{3}$$

The probability of getting an even number may be calculated as

$$\Pr(\text{even}) = \Pr(2\vee4\vee6) = \Pr(2) + \Pr(4) + \Pr(6) = \frac{1}{6} + \frac{1}{6} + \frac{1}{6} = \frac{1}{2}$$

The probability of getting an even number that is greater than 3 may be calculated as

$$\Pr(\text{even and greater than 3}) = \Pr(4\vee6) = \Pr(4) + \Pr(6) = \frac{1}{6} + \frac{1}{6} = \frac{1}{3}$$

The probability of getting an even number or a 3 may be calculated as

$$\Pr(\text{even or 3}) = \Pr(2\vee4\vee6\vee3) = \frac{4}{6} = \frac{2}{3}$$

Finally, calculating the probability of getting either a 1, 2, 3, 4, 5, or 6 (that is, the probability that the die will show one face or another) gives $\frac{6}{6}$, or 1.

¹Actually the statements are not mutually exclusive in the logical sense. We cannot show that they are inconsistent with each other by the method of truth tables, and it is logically possible that the die might change shape upon being thrown so as to display two faces simultaneously. To treat this case rigorously, we would have to use the general disjunction rule, along with a battery of assumptions: $\Pr(1\&2) = 0$, $\Pr(2\&3) = 0$, $\Pr(1\&3) = 0$, etc. However, we shall see that the result is the same as when we use the special disjunction rule, and treat these statements as if they were mutually exclusive.

We will now apply the special disjunction rule to a case of more general interest. It can be shown, by the method of truth tables, that any statement p is inconsistent with its negation, $\sim p$. Since p and $\sim p$ are therefore mutually exclusive, the special disjunction rule permits the conclusion that

$$\Pr(p \vee \sim p) = \Pr(p) + \Pr(\sim p)$$

But the statement $p \vee \sim p$ is a tautology, so by Rule 1,

$$\Pr(p \vee \sim p) = 1$$

Putting these two conclusions together gives

$$\Pr(p) + \Pr(\sim p) = 1$$

If the quantity $\Pr(p)$ is subtracted from both sides of the equation, the sides will remain equal, so we may conclude that

$$\Pr(\sim p) = 1 - \Pr(p)$$

This conclusion holds good for any statement, since any statement is inconsistent with its negation, and for any statement p its disjunction with its negation, $p \vee \sim p$, is a tautology. This therefore establishes a general negation rule, which allows us to calculate the probability of a negation from the probability of its constituent statement:

Rule 5: $\Pr(\sim p) = 1 - \Pr(p)$.

Suppose in the example using the die we wanted to know the probability of not getting a 3:

$$\Pr(\sim 3) = 1 - \Pr(3) = 1 - \frac{1}{6} = \frac{5}{6}$$

Note that we get the same answer as we would if we took the long road to solving the problem and confined ourselves to using the special disjunction rule:

$$\begin{aligned} \Pr(\sim 3) &= \Pr(1 \vee 2 \vee 4 \vee 5 \vee 6) \\ &= \Pr(1) + \Pr(2) + \Pr(4) + \Pr(5) + \Pr(6) \\ &= \frac{1}{6} + \frac{1}{6} + \frac{1}{6} + \frac{1}{6} + \frac{1}{6} = \frac{5}{6} \end{aligned}$$

We shall apply the special disjunction rule one more time in order to establish another generally useful rule. For any two statements, p , q , we can show by the truth table method that the complex statements $p \& q$, $p \& \sim q$, and $\sim p \& q$ are inconsistent with each other. As shown in the following table, there is no case in which two of them are true:

	p	q	$\sim p$	$\sim q$	$p \& q$	$p \& \sim q$	$\sim p \& q$
Case 1:	T	T	F	F	T	F	F
Case 2:	T	F	F	T	F	T	F
Case 3:	F	T	T	F	F	F	T
Case 4:	F	F	T	T	F	F	F

Since they are mutually exclusive, we can apply the special disjunction rule and conclude:

- a. $\Pr[(p \& q) \vee (p \& \sim q)] = \Pr(p \& q) + \Pr(p \& \sim q)$
- b. $\Pr[(p \& q) \vee (\sim p \& q)] = \Pr(p \& q) + \Pr(\sim p \& q)$
- c. $\Pr[(p \& q) \vee (p \& \sim q) \vee (\sim p \& q)] = \Pr(p \& q) + \Pr(p \& \sim q) + \Pr(\sim p \& q)$

But the complex statement $(p \& q) \vee (p \& \sim q)$ is logically equivalent to the simple statement p , as is shown by the following truth table:

	p	q	$\sim q$	$p \& q$	$p \& \sim q$	$(p \& q) \vee (p \& \sim q)$
Case 1:	T	T	F	T	F	T
Case 2:	T	F	T	F	T	T
Case 3:	F	T	F	F	F	F
Case 4:	F	F	T	F	F	F

Since, according to Rule 3, logically equivalent statements have the same probability, equation (a) may be rewritten as

$$a'. \Pr(p) = \Pr(p \& q) + \Pr(p \& \sim q)$$

A similar truth table will show that the complex statement $(p \& q) \vee (\sim p \& q)$ is logically equivalent to the simple statement q . Therefore, equation (b) may be rewritten as

$$b'. \Pr(q) = \Pr(p \& q) + \Pr(\sim p \& q)$$

Finally, a truth table will show that the complex statement $(p \& q) \vee (p \& \sim q) \vee (\sim p \& q)$ is logically equivalent to the complex statement $p \vee q$, which enables us to rewrite equation (c) as

$$c'. \Pr(p \vee q) = \Pr(p \& q) + \Pr(p \& \sim q) + \Pr(\sim p \& q)$$

Now let us add equations (a') and (b') together to get

$$d. \Pr(p) + \Pr(q) = 2 \Pr(p \& q) + \Pr(p \& \sim q) + \Pr(\sim p \& q)$$


If we subtract the quantity $\Pr(p \& q)$ from both sides of the preceding equation, we get

$$d'. \Pr(p) + \Pr(q) - \Pr(p \& q) = \Pr(p \& q) + \Pr(p \& \sim q) + \Pr(\sim p \& q)$$

If equation (d') is compared with equation (c') we see that $\Pr(p \vee q)$ is equal to the same thing as $\Pr(p) + \Pr(q) - \Pr(p \& q)$. This establishes a general disjunction rule that is good for all disjunctions, whether the disjuncts are mutually exclusive or not:

Rule 6: $\Pr(p \vee q) = \Pr(p) + \Pr(q) - \Pr(p \& q)$.

If some of the algebra used to establish the general disjunction rule has left you behind, the following diagram may help to make the reasoning clear:



When $\Pr(p)$ is added to $\Pr(q)$, then $\Pr(p \& q)$ is counted twice. But to get $\Pr(p \vee q)$, it should be counted only once. Thus, to get $\Pr(p \vee q)$, we add $\Pr(p)$ and $\Pr(q)$ and then subtract $\Pr(p \& q)$ to make up for having counted it twice. In the case in which p and q are mutually exclusive, this makes no difference, because when p and q are mutually exclusive, $\Pr(p \& q) = 0$. No matter how many times 0 is counted, we will always get the same result. For example, by the general disjunction rule, $\Pr(p \vee \sim p) = \Pr(p) + \Pr(\sim p) - \Pr(p \& \sim p)$. But the statement $p \& \sim p$ is a self-contradiction, so its probability is zero. Thus, we get the same result as if we had used the special disjunction rule. Counting $\Pr(p \& q)$ twice *does* make a difference when p and q are *not* mutually exclusive. Suppose we use the general disjunction rule to calculate the probability of the complex statement $p \vee p$:

$$\Pr(p \vee p) = \Pr(p) + \Pr(p) - \Pr(p \& p)$$

But since the complex statement $p \& p$ is logically equivalent to the simple statement p , $\Pr(p \& p) = \Pr(p)$, we get

$$\Pr(p \vee p) = \Pr(p) + \Pr(p) - \Pr(p) = \Pr(p)$$

We know this is the correct answer, because the complex statement $p \vee p$ is also logically equivalent to the simple statement p .

The example with the die shall be used to give one more illustration of the use of the general disjunction rule. Suppose that we want to know the probability that the die will come up an even number or a number less than 3. There is a way to calculate this probability using only the special disjunction rule:

$$\begin{aligned}\Pr(\text{even } \vee \text{ less than } 3) &= \Pr(1 \vee 2 \vee 4 \vee 6) \\ &= \Pr(1) + \Pr(2) + \Pr(4) + \Pr(6) = \frac{4}{6} = \frac{2}{3}\end{aligned}$$

We may use the special disjunction rule because the outcomes 1, 2, 4, and 6 are mutually exclusive. However, the outcomes “even” and “less than 3” are not mutually exclusive, since the die might come up 2. Thus, we may apply the general disjunction rule as follows:

$$\begin{aligned}\Pr(\text{even } \vee \text{ less than } 3) \\ &= \Pr(\text{even}) + \Pr(\text{less than } 3) - \Pr(\text{even} \&\text{less than } 3)\end{aligned}$$

Now we may calculate $\Pr(\text{even})$ as $\Pr(2 \vee 4 \vee 6)$ by the special disjunction rule; it is equal to $\frac{1}{2}$. We may calculate $\Pr(\text{less than } 3)$ as $\Pr(1 \vee 2)$ by the special disjunction rule; it is equal to $\frac{1}{3}$. And we may calculate $\Pr(\text{even} \&\text{less than } 3)$ as $\Pr(2)$, which is equal to $\frac{1}{6}$. So, by this method,

$$\Pr(\text{even } \vee \text{ less than } 3) = \frac{1}{2} + \frac{1}{3} - \frac{1}{6} = \frac{2}{3}$$

The role of the subtraction term can be seen clearly in this example. What we have done is to calculate $\Pr(\text{even } \vee \text{ less than } 3)$ as

$$\Pr(2 \vee 4 \vee 6) + \Pr(1 \vee 2) - \Pr(2)$$

so the subtraction term compensates for adding in $\Pr(2)$ twice when we add $\Pr(\text{even})$ and $\Pr(\text{less than } 3)$. In this example use of the general disjunction rule was the long way of solving the problem. But in some cases it is necessary to use the general disjunction rule. Suppose you are told that

$$\Pr(p) = \frac{1}{2}$$

$$\Pr(q) = \frac{1}{3}$$

$$\Pr(p \& q) = \frac{1}{4}$$

You are asked to calculate $\Pr(p \vee q)$. Now you cannot use the special disjunction rule since you know that p and q are not mutually exclusive. If they were, $\Pr(p \& q)$ would be 0, and you are told that it is $\frac{1}{4}$. Therefore, you must use the general disjunction rule in the following way:

$$\begin{aligned}\Pr(p \vee q) &= \Pr(p) + \Pr(q) - \Pr(p \& q) \\ &= \frac{1}{2} + \frac{1}{3} - \frac{1}{4} = \frac{7}{12}\end{aligned}$$

In Section VI.2, we compared the rules of the probability calculus to the way in which the truth tables for the logical connectives relate the truth or falsity of a complex statement to the truth or falsity of its simple constituent statements. We are now at the point where we must qualify this comparison. We can always determine the truth or falsity of a complex statement if we know whether its simple constituent statements are true or false. But we cannot always calculate the probability of a complex statement from the probabilities of its simple constituent statements. Sometimes, as in the example above, in order to calculate the probability of the complex statement pvq , we need not only know the probabilities of its simple constituent statements, p and q , we also need to know the probability of another complex statement, $p\&q$. We shall discuss the rules that govern the probabilities of such conjunctions in the next section. However, we shall find that it is not always possible to calculate the probability of a conjunction simply from the probabilities of its constituent statements.

Exercises

1. Suppose you have an ordinary deck of 52 playing cards. You are to draw one card. Assume that each card has a probability of $1/52$ of being drawn. What is the probability that you will draw:
 - a. The ace of spades?
 - b. The queen of hearts?
 - c. The ace of spades or the queen of hearts?
 - d. An ace?
 - e. A heart?
 - f. A face card (king, queen, or jack)?
 - g. A card that is not a face card?
 - h. An ace or a spade?
 - i. A queen or a heart?
 - j. A queen or a non-spade?
2. $\Pr(p) = \frac{1}{2}$, $\Pr(q) = \frac{1}{2}$, $\Pr(p\&q) = \frac{1}{8}$. What is $\Pr(pvq)$?
3. $\Pr(r) = \frac{1}{2}$, $\Pr(s) = \frac{1}{4}$, $\Pr(rvs) = \frac{3}{4}$. What is $\Pr(r\&s)$?
4. $\Pr(u) = \frac{1}{2}$, $\Pr(t) = \frac{3}{4}$, $\Pr(u\&\sim t) = \frac{1}{8}$. What is $\Pr(uv\sim t)$?

VI.4. CONJUNCTION RULES AND CONDITIONAL PROBABILITY. Before the rules that govern the probability of conjunctions are discussed, it is necessary to introduce the notion of *conditional probability*. We may write $\Pr(q \text{ given } p)$ as the probability of q *on the condition that* p .

This probability may or may not be different from $\Pr(q)$. We shall deal with the concept of conditional probability on the intuitive level before a precise definition for it is introduced.

In the example with the die, we found that the probability of throwing an even number was $\frac{1}{2}$. However, the probability of getting an even number *given that* a 2 or a 4 is thrown is not $\frac{1}{2}$ but 1. And the probability of casting an even number *given that* a 1 or a 3 is thrown is 0. To take a little more complicated example, suppose that the die remains unchanged and you are to bet on whether it will come up even, with a special agreement that if it comes up 5 all bets will be off and it will be thrown again. In such a situation you would be interested in the probability that it will come up even *given that* it will be either a 1, 2, 3, 4, or 6. This probability should be greater than $\frac{1}{2}$ since the condition excludes one of the ways in which the die could come up odd. It is, in fact, $\frac{3}{5}$. Thus, the probabilities of “even,” given three different conditions, are each different from the probability of “even” by itself:

- a. $\Pr(\text{even}) = \frac{1}{2}$
- b. $\Pr(\text{even given } 2 \vee 4) = 1$
- c. $\Pr(\text{even given } 1 \vee 3) = 0$
- d. $\Pr(\text{even given } 1 \vee 2 \vee 3 \vee 4 \vee 6) = \frac{3}{5}$

Conditional probabilities allow for the fact that if a certain statement, p , is known to be true, this may affect the probability to be assigned to another statement, q . The most striking cases occur when there is a deductively valid argument from p to q :

$$\begin{array}{l}
 p = \text{The next throw of the die will come up 2} \\
 \vee \\
 \text{the next throw of the die will come up 4.} \\
 \hline
 q = \text{The next throw of the die will come up even.}
 \end{array}$$

In this case, $\Pr(q \text{ given } p) = 1$:²

$$\Pr(\text{even given } 2 \vee 4) = 1$$

Suppose there is a deductively valid argument from p to $\sim q$:

²We must make one qualification to this statement. When p is a self-contradiction, then for any statement q there is a deductively valid argument from p to q and a deductively valid argument from p to $\sim q$. In such a case, $\Pr(q \text{ given } p)$ has no value.

$p =$ The next throw of the die will come up 1

v

the next throw of the die will come up 3.

$\sim q =$ The next throw of the die will *not* come up even.

In this case, $\text{Pr}(q \text{ given } p) = 0$:

$$\text{Pr}(\text{even given } 1v3) = 0.^3$$

There are, however, important cases where neither the argument from p to q nor the argument from p to $\sim q$ is deductively valid and yet $\text{Pr}(q \text{ given } p)$ differs from $\text{Pr}(q)$, as in the previous example with the die:

$$\text{Pr}(\text{even given } 1v2v3v4v6) = \frac{3}{5}$$

$$\text{Pr}(\text{even}) = \frac{1}{2}$$

There are other cases where the knowledge that p is true may be completely irrelevant to the probability to be assigned to q . For example, it was said that the probability that the next throw of the die will come up even is $\frac{1}{2}$. We could say that the probability that the next throw of the die will come up even, given that the President of the United States sneezes simultaneously with our throw, is still $\frac{1}{2}$. The President's sneeze is irrelevant to the probability assigned to "even." Thus, the two statements "The next throw of the die will come up even" and "The President of the United States will sneeze simultaneously with the next throw of the die" are *independent*.⁴

We can now give substance to the intuitive notions of conditional probability and independence by defining them in terms of pure statement probabilities. First we will define conditional probability:

Definition 12: *Conditional probability:*⁵

$$\text{Pr}(q \text{ given } p) = \frac{\text{Pr}(p \& q)}{\text{Pr}(p)}$$

Let us see how this definition works out in the example of the die:

³We must make one qualification to this statement. When p is a self-contradiction, then for any statement q there is a deductively valid argument from p to q and a deductively valid argument from p to $\sim q$. In such a case, $\text{Pr}(q \text{ given } p)$ has no value.

⁴This type of independence is called probabilistic or stochastic independence. It should not be confused with the mutual logical independence discussed in deductive logic. Stochastic independence of two statements is neither a necessary nor a sufficient condition for their mutual logical independence.

⁵When $\text{Pr}(p) = 0$ the quotient is not defined. In this case there is no $\text{Pr}(q \text{ given } p)$.

$$\text{a. } \Pr(\text{even given } 2v4) = \frac{\Pr[\text{even}\&(2v4)]}{\Pr(2v4)} = \frac{\Pr(2v4)}{\Pr(2v4)} = 1$$

$$\text{b. } \Pr(\text{even given } 1v3) = \frac{\Pr[\text{even}\&(1v3)]}{\Pr(1v3)} = \frac{0}{\frac{1}{3}} = 0$$

$$\begin{aligned} \text{c. } \Pr(\text{even given } 1v2v3v4v6) &= \frac{\Pr[\text{even}\&(1v2v3v4v6)]}{\Pr(1v2v3v4v6)} \\ &= \frac{\Pr(2v4v6)}{\Pr(1v2v3v4v6)} = \frac{\frac{3}{6}}{\frac{5}{6}} = \frac{3}{5} \end{aligned}$$

Notice that the conditional probabilities computed by using the definition accord with the intuitive judgments as to conditional probabilities in the die example. We may test the definition in another way. Consider the special case of $\Pr(q \text{ given } p)$, where p is a tautology and q is a contingent statement. Since a tautology makes no factual claim, we would not expect knowledge of its truth to influence the probability that we would assign to the contingent statement, q . The probability that the die will come up even given that it will come up either even or odd should be simply the probability that it will come up even. In general, if we let T stand for an arbitrary tautology, we should expect $\Pr(q \text{ given } T)$ to be equal to $\Pr(q)$. Let us work out $\Pr(q \text{ given } T)$, using the definition of conditional probability:

$$\Pr(q \text{ given } T) = \frac{\Pr(T\&q)}{\Pr(T)}$$

But the probability of a tautology is always equal to 1. This gives

$$\Pr(q \text{ given } T) = \Pr(T\&q)$$

When T is a tautology and q is any statement whatsoever, the complex statement $T\&q$ is logically equivalent to the simple statement q . This can always be shown by truth tables. Since logically equivalent statements have the same probability, $\Pr(q \text{ given } T) = \Pr(q)$.⁶ Again, the definition of conditional probability gives the expected result.

Now that conditional probability has been defined, that concept can be used to define independence:

⁶We could have constructed the probability calculus by taking conditional probabilities as basic, and then defining pure statement probabilities as follows: The probability of a statement is defined as its probability *given* a tautology. Instead we have taken statement probabilities as basic, and defined conditional probabilities. The choice of starting point makes no difference to the system as a whole. The systems are equivalent.

Definition 13: *Independence:* Two statements p and q are independent if and only if $\Pr(q \text{ given } p) = \Pr(q)$.

We talk of two statements p and q being independent, rather than p being independent of q and q being independent of p . We can do this because we can prove that $\Pr(q \text{ given } p) = \Pr(q)$ if and only if $\Pr(p \text{ given } q) = \Pr(p)$. If $\Pr(q \text{ given } p) = \Pr(q)$, then, by the definition of conditional probability,

$$\frac{\Pr(p \& q)}{\Pr(p)} = \Pr(q)$$

Multiplying both sides of the equation by $\Pr(p)$ and dividing both sides by $\Pr(q)$, we have

$$\frac{\Pr(p \& q)}{\Pr(q)} = \Pr(p)$$

But by the definition of conditional probability, this means $\Pr(p \text{ given } q) = \Pr(p)$.

This proof only works if neither of the two statements has 0 probability. Otherwise, one of the relevant quotients would not be defined. To take care of this eventuality, we may add an additional clause to the definition and say that two statements are also independent if at least one of them has probability 0. It is important to realize the difference between independence and mutual exclusiveness. The statement about the outcome of the throw of the die and the statement about the President's sneeze are independent, but they are not mutually exclusive. They can very well be true together. On the other hand, the statements "The next throw of the die will come up an even number" and "The next throw of the die will come up a 5" are mutually exclusive, but they are not independent. $\Pr(\text{even}) = \frac{1}{2}$, but $\Pr(\text{even given } 5) = 0$. $\Pr(5) = \frac{1}{6}$, but $\Pr(5 \text{ given even}) = 0$. In general, if p and q are mutually exclusive they are not independent, and if they are independent they are not mutually exclusive.⁷

Having specified the definitions of conditional probability and independence, the rules for conjunctions can now be introduced. The *general conjunction rule* follows directly from the definition of conditional probability:

Rule 7: $\Pr(p \& q) = \Pr(p) \times \Pr(q \text{ given } p)$.

The proof is simple. Take the definition of conditional probability:

$$\Pr(q \text{ given } p) = \frac{\Pr(p \& q)}{\Pr(p)}$$

⁷The exception is when at least one of the statements is a self-contradiction and thus has probability 0.

Multiply both sides of the equation by $\Pr(p)$ to get

$$\Pr(p) \times \Pr(q \text{ given } p) = \Pr(p \& q)$$

which is the general conjunction rule. When p and q are independent, $\Pr(q \text{ given } p) = \Pr(q)$, and we may substitute $\Pr(q)$ for $\Pr(q \text{ given } p)$ in the general conjunction rule, thus obtaining

$$\Pr(p) \times \Pr(q) = \Pr(p \& q)$$

Of course, the substitution may only be made in the special case when p and q are independent. This result constitutes the *special conjunction rule*:

Rule 8: If p and q are independent, then $\Pr(p \& q) = \Pr(p) \times \Pr(q)$.

The general conjunction rule is more basic than the special conjunction rule. But since the special conjunction rule is simpler, its application will be illustrated first. Suppose that two dice are thrown simultaneously. The basic probabilities are as follows:

Die A	Die B
$\Pr(1) = \frac{1}{6}$	$\Pr(1) = \frac{1}{6}$
$\Pr(2) = \frac{1}{6}$	$\Pr(2) = \frac{1}{6}$
$\Pr(3) = \frac{1}{6}$	$\Pr(3) = \frac{1}{6}$
$\Pr(4) = \frac{1}{6}$	$\Pr(4) = \frac{1}{6}$
$\Pr(5) = \frac{1}{6}$	$\Pr(5) = \frac{1}{6}$
$\Pr(6) = \frac{1}{6}$	$\Pr(6) = \frac{1}{6}$

Since the face shown by die A presumably does not influence the face shown by die B, or vice versa, it shall be assumed that all statements claiming various outcomes for die A are independent of all the statements claiming various outcomes for die B. That is, the statements "Die A will come up a 3" and "Die B will come up a 5" are independent, as are the statements "Die A will come up a 6" and "Die B will come up a 6." The statements "Die A will come up a 5" and "Die A will come up a 3" are not independent; they are mutually exclusive (when made in regard to the same throw).

Now suppose we wish to calculate the probability of throwing a 1 on die A and a 6 on die B. The special conjunction rule can now be used:

$$\begin{aligned} \Pr(1 \text{ on } A \& 6 \text{ on } B) &= \Pr(1 \text{ on } A) \times \Pr(6 \text{ on } B) \\ &= \frac{1}{6} \times \frac{1}{6} = \frac{1}{36} \end{aligned}$$

In the same way, the probability of each of the 36 possible combinations of results of die A and die B may be calculated as $\frac{1}{36}$, as shown in Table VI.1. Note that each of the cases in the table is mutually exclusive of each other case.

Thus, by the special alternation rule, the probability of case 1 v case 3 is equal to the probability of case 1 plus the probability of case 3.

Table VI.1

Possible results when throwing two dice					
Case	Die A	Die B	Case	Die A	Die B
1	1	1	19	4	1
2	1	2	20	4	2
3	1	3	21	4	3
4	1	4	22	4	4
5	1	5	23	4	5
6	1	6	24	4	6
7	2	1	25	5	1
8	2	2	26	5	2
9	2	3	27	5	3
10	2	4	28	5	4
11	2	5	29	5	5
12	2	6	30	5	6
13	3	1	31	6	1
14	3	2	32	6	2
15	3	3	33	6	3
16	3	4	34	6	4
17	3	5	35	6	5
18	3	6	36	6	6

Suppose now that we wish to calculate the probability that the dice will come up showing a 1 and a 6. There are two ways this can happen: a 1 on die A and a 6 on die B (case 31). The probability of this combination appearing may be calculated as follows:

$$\Pr(1 \ \& \ 6) = \Pr[(1 \text{ on } A \ \& \ 6 \text{ on } B) \vee (1 \text{ on } B \ \& \ 6 \text{ on } A)]$$

Since the cases are mutually exclusive, the special disjunction rule may be used to get

$$\begin{aligned} \Pr[(1 \text{ on } A \ \& \ 6 \text{ on } B) \vee (1 \text{ on } B \ \& \ 6 \text{ on } A)] \\ = \Pr(1 \text{ on } A \ \& \ 6 \text{ on } B) + \Pr(1 \text{ on } B \ \& \ 6 \text{ on } A) \end{aligned}$$

But it has already been shown, by the special conjunction rule, that

$$\Pr(1 \text{ on } A \ \& \ 6 \text{ on } B) = \frac{1}{36}$$

$$\Pr(1 \text{ on } B \ \& \ 6 \text{ on } A) = \frac{1}{36}$$

so the answer is $\frac{1}{36} + \frac{1}{36}$, or $\frac{1}{18}$.

The same sort of reasoning can be used to solve more complicated problems. Suppose we want to know the probability that the sum of spots showing on both dice will equal 7. This happens only in cases 6, 11, 16, 21, 26, and 31. Therefore

$$\begin{aligned} \Pr(\text{total of } 7) &= \Pr[(1 \text{ on } A \ \& \ 6 \text{ on } B) \\ &\quad \vee (2 \text{ on } A \ \& \ 5 \text{ on } B) \\ &\quad \vee (3 \text{ on } A \ \& \ 4 \text{ on } B) \\ &\quad \vee (4 \text{ on } A \ \& \ 3 \text{ on } B) \\ &\quad \vee (5 \text{ on } A \ \& \ 2 \text{ on } B) \\ &\quad \vee (6 \text{ on } A \ \& \ 1 \text{ on } B)] \end{aligned}$$

Using the special disjunction rule and the special conjunction rule $\Pr(\text{total of } 7) = \frac{6}{36}$, or $\frac{1}{6}$.

In solving a particular problem, there are often several ways to apply the rules. Suppose we wanted to calculate the probability that both dice will come up even. We could determine in which cases both dice are showing even numbers, and proceed as before, but this is the long way to solve the problem. Instead, we can calculate the probability of getting an even number on die A as $\frac{1}{2}$ by the special disjunction rule:

$$\begin{aligned} \Pr(\text{even on } A) &= \Pr(2 \text{ on } A \ \vee \ 4 \text{ on } A \ \vee \ 6 \text{ on } A) \\ &= \Pr(2 \text{ on } A) + \Pr(4 \text{ on } A) + \Pr(6 \text{ on } A) \\ &= \frac{3}{6} = \frac{1}{2} \end{aligned}$$

and calculate the probability of getting an even number on die B as $\frac{1}{2}$ by the same method. Then, by the special conjunction rule,⁸

$$\begin{aligned} \Pr(\text{even on } A \ \& \ \text{even on } B) &= \Pr(\text{even on } A) \times \Pr(\text{even on } B) \\ &= \frac{1}{2} \times \frac{1}{2} = \frac{1}{4} \end{aligned}$$

We apply the *general conjunction rule* when two statements are not independent. Such is the case in the following example. Suppose you are presented with a bag containing ten gumdrops, five red and five black. You are to shake the bag, close your eyes and draw out a gumdrop, look at it, eat it, and then repeat the process once more. We shall assume that, at the time of each draw, each gumdrop in the bag has an equal probability of being drawn. The problem is to find the probability of drawing two red gumdrops.

⁸It can be shown that the statements “Die A will come up even” and “Die B will come up even” are independent, on the basis of the independence assumptions made in setting up this example.

To solve this problem we must find the probability of the conjunction $\Pr(\text{red on 1} \ \& \ \text{red on 2})$. We will first find $\Pr(\text{red on 1})$. We will designate each of the gumdrops by a letter: $A, B, C, D, E, F, G, H, I, J$. We know that we will draw one of these on the first draw, so

$$\Pr(A \text{ on 1} \vee B \text{ on 1} \vee C \text{ on 1} \vee \dots \vee J \text{ on 1}) = 1$$

Now, by the special disjunction rule,

$$\Pr(A \text{ on 1}) + \Pr(B \text{ on 1}) + \Pr(C \text{ on 1}) + \dots + \Pr(J \text{ on 1}) = 1$$

Since each of the gumdrops has an equal chance of being drawn, and there are 10 gumdrops, therefore

$$\Pr(A \text{ on 1}) = \frac{1}{10}$$

$$\Pr(B \text{ on 1}) = \frac{1}{10}$$

⋮

$$\Pr(J \text{ on 1}) = \frac{1}{10}$$

We said that there were five red ones. We will use the letters A, B, C, D , and E to designate the red gumdrops and the remaining letters to designate the black ones. By the special disjunction rule, the probability of getting a red gumdrop on draw 1 is

$$\begin{aligned} & \Pr(A \text{ on 1} \vee B \text{ on 1} \vee C \text{ on 1} \vee D \text{ on 1} \vee E \text{ on 1}) \\ &= \Pr(A \text{ on 1}) + \Pr(B \text{ on 1}) + \Pr(C \text{ on 1}) + \Pr(D \text{ on 1}) + \Pr(E \text{ on 1}) \\ &= \frac{5}{10} = \frac{1}{2} \end{aligned}$$

We shall have to use the general conjunction rule to find $\Pr(\text{red on 1} \ \& \ \text{red on 2})$, since the statements “A red gumdrop will be drawn the first time” and “A red gumdrop will be drawn the second time” are not independent. If a red gumdrop is drawn the first time, this will leave four red and five black gumdrops in the bag with equal chances of being drawn on the second draw. But if a black gumdrop is drawn the first time, this will leave five red and four black gumdrops awaiting the second draw. Thus, the knowledge that a red one is drawn the first time will influence the probability we assign to a red one being drawn the second time, and the two statements are not independent. Applying the general conjunction rule, we get

$$\Pr(\text{red on 1} \ \& \ \text{red on 2}) = \Pr(\text{red on 1}) \times \Pr(\text{red on 2 given red on 1})$$

We have already found $\Pr(\text{red on 1})$. Now we must calculate $\Pr(\text{red on 2 given red on 1})$. Given that we draw a red gumdrop on the first draw, there will be nine gumdrops remaining: four red and five black. We must draw one of them,

and they each have an equal chance of being drawn. By reasoning similar to that used above, each has a probability of $\frac{1}{9}$ of being drawn, and the probability of drawing a red one is $\frac{4}{9}$. Therefore

$$\Pr(\text{red on 2 given red on 1}) = \frac{4}{9}$$

We can now complete our calculations:

$$\Pr(\text{red on 1 \& red on 2}) = \frac{1}{2} \times \frac{4}{9} = \frac{2}{9}$$

We can calculate $\Pr(\text{black on 1 \& red on 2})$ in the same way:

$$\Pr(\text{black on 1}) = \frac{1}{2}$$

$$\Pr(\text{red on 2 given black on 1}) = \frac{5}{9}$$

Therefore, by the general conjunction rule,

$$\Pr(\text{black on 1 \& red on 2}) = \frac{1}{2} \times \frac{5}{9} = \frac{5}{18}$$

At this point the question arises as to what the $\Pr(\text{red on 2})$ is. We know $\Pr(\text{red on 2 given red on 1}) = \frac{4}{9}$. We know $\Pr(\text{red on 2 given black on 1}) = \frac{5}{9}$. But what we want to know now is the probability of getting a red gumdrop on the second draw before we have made the first draw. We can get the answer if we realize that *red on 2* is logically equivalent to

$$(\text{red on 1 \& red on 2}) \vee (\text{not-red on 1 \& red on 2})$$

Remember that the simple statement q is logically equivalent to the complex statement $(p \& q) \vee (\sim p \& q)$. Therefore

$$\Pr(\text{red on 2}) = \Pr[(\text{red on 1 \& red on 2}) \vee (\text{not-red on 1 \& red on 2})]$$

By the special disjunction rule,

$$\Pr(\text{red on 2}) = \Pr(\text{red on 1 \& red on 2}) + \Pr(\text{not-red on 1 \& red on 2})$$

We have calculated $\Pr(\text{red on 1 \& red on 2})$ as $\frac{2}{9}$. We have also calculated

$$\Pr(\text{not-red on 1 \& red on 2}) = \Pr(\text{black on 1 \& red on 2}) = \frac{5}{18}$$

Therefore

$$\Pr(\text{red on 2}) = \frac{2}{9} + \frac{5}{18} + \frac{4}{18} + \frac{5}{18} = \frac{9}{18} = \frac{1}{2}$$

The same sort of applications of conditional probability and the general conjunction rule would apply to card games where the cards that have

been played are placed in a discard pile rather than being returned to the deck. Such considerations are treated very carefully in manuals on poker and blackjack. In fact, some gambling houses have resorted to using a new deck for each hand of blackjack in order to keep astute students of probability from gaining an advantage over the house.

Exercises

1. $\Pr(p) = \frac{1}{2}$, $\Pr(q) = \frac{1}{2}$, p and q are independent.
 - a. What is $\Pr(p \& q)$?
 - b. Are p and q mutually exclusive?
 - c. What is $\Pr(p \vee q)$?
2. Suppose two dice are rolled, as in the example above.
 - a. What is the probability of both dice showing a 1?
 - b. What is the probability of both dice showing a 6?
 - c. What is the probability that the total number of spots showing on both dice will be either 7 or 11?
3. A coin is flipped three times. Assume that on each toss $\Pr(\text{heads}) = \frac{1}{2}$ and $\Pr(\text{tails}) = \frac{1}{2}$. Assume that the tosses are independent.
 - a. What is $\Pr(3 \text{ heads})$?
 - b. What is $\Pr(2 \text{ heads and } 1 \text{ tail})$?
 - c. What is $\Pr(1 \text{ head and } 2 \text{ tails})$?
 - d. What is $\Pr(\text{head on toss } 1 \& \text{ tail on toss } 2 \& \text{ head on toss } 3)$?
 - e. What is $\Pr(\text{at least } 1 \text{ tail})$?
 - f. What is $\Pr(\text{no heads})$?
 - g. What is $\Pr(\text{either } 3 \text{ heads or } 3 \text{ tails})$?
4. Suppose you have an ordinary deck of 52 cards. A card is drawn and is not replaced, then another card is drawn. Assume that on each draw each of the cards then in the deck has an equal chance of being drawn.
 - a. What is $\Pr(\text{ace on draw } 1)$?
 - b. What is $\Pr(10 \text{ on draw } 2 \text{ given ace on draw } 1)$?
 - c. What is $\Pr(\text{ace on draw } 1 \& 10 \text{ on draw } 2)$?
 - d. What is $\Pr(10 \text{ on draw } 1 \& \text{ ace on draw } 2)$?
 - e. What is $\Pr(\text{an ace and a } 10)$?
 - f. What is $\Pr(2 \text{ aces})$?
5. The probability that George will study for the test is $\frac{4}{5}$. The probability that he will pass the test given that he studies is $\frac{3}{5}$. The probability that he will pass the test given that he does not study is $\frac{1}{10}$. What is the probability that George will pass the test? Hint: The simple statement "George will pass the test" is logically equivalent to the complex statement "Either George will study and pass the test or George will not study and pass the test."

VI.5. EXPECTED VALUE OF A GAMBLE. The attractiveness of a wager depends not only on the probabilities involved, but also on the odds given. The probability of getting a head and a tail on two independent tosses of a fair coin is $\frac{1}{2}$, while the probability of getting two heads is only $\frac{1}{4}$. But if someone were to offer either to bet me even money that I will not get a head and a tail or give 100 to 1 odds against my getting two heads, I would be well advised to take the second wager. The probability that I will win the second wager is less, but this is more than compensated for by the fact that if I win, I will win a great deal, and if I lose, I will lose much less. The attractiveness of a wager can be measured by calculating its *expected value*. To calculate the expected value of a gamble, first list all the possible outcomes, along with their probabilities and the amount won in each case. A loss is listed as a negative amount. Then for each outcome multiply the probability by the amount won or lost. Finally, add these products to obtain the expected value. To illustrate, suppose someone bets me 10 dollars that I will not get a head and a tail on two tosses of a fair coin. The expected value of this wager for me can be calculated as follows:

Possible outcomes				
Toss 1	Toss 2	Probability	Gain	Probability \times Gain
H	H	$\frac{1}{4}$	-\$10	-\$2.50
H	T	$\frac{1}{4}$	10	2.50
T	H	$\frac{1}{4}$	10	2.50
T	T	$\frac{1}{4}$	-10	-2.50
Expected value:				\$0.00

Thus, the expected value of the wager for me is \$0, and since my opponent wins what I lose and loses what I win, the expected value for him is also \$0. Such a wager is called a *fair bet*. Now let us calculate the expected value for me of a wager where my opponent will give me 100 dollars if I get two heads, and I will give him one dollar if I do not.

Possible outcomes				
Toss 1	Toss 2	Probability	Gain	Probability \times Gain
H	H	$\frac{1}{4}$	\$100	\$25.00
H	T	$\frac{1}{4}$	-1	-0.25
T	H	$\frac{1}{4}$	-1	-0.25
T	T	$\frac{1}{4}$	-1	-0.25
Expected value:				\$24.25

The expected value of this wager for me is \$24.25. Since my opponent loses what I win, the expected value for him is $-\$24.25$. This is not a fair bet, since it is favorable to me and unfavorable to him.

The procedure for calculating expected value and the rationale behind it are clear, but let us try to attach some meaning to the numerical answer. This can be done in the following way. Suppose that I make the foregoing wager many times. And suppose that over these many times the distribution of results corresponds to the probabilities; that is, I get two heads one-fourth of the time; a head and then a tail one-fourth of the time; a tail and then a head one-fourth of the time; and two tails one-fourth of the time. Then the expected value will be equal to my average winnings on a wager (that is, my total winnings divided by the number of wagers I have made).

I said that expected value was a measure of the attractiveness of a wager. Generally, it seems reasonable to accept a wager with a positive expected gain and reject a wager with a negative expected gain. Furthermore, if you are offered a choice of wagers, it seems reasonable to choose the wager with the highest expected value. These conclusions, however, are oversimplifications. They assume that there is no positive or negative value associated with risk itself, and that gains or losses of equal amounts of money represent gains or losses of equal amounts of money represent gains or losses of equal amount of value to the individual involved. Let us examine the first assumption.

Suppose that you are compelled to choose an even-money wager either for 1 dollar or for 100 dollars. The expected value of both wagers is 0. But if you wish to avoid risks as much as possible, you would choose the smaller wager. You would, then, assign a negative value to risk itself. However, if you enjoy taking larger risks for their own sake, you would choose the larger wager. Thus, although expected value is a major factor in determining the attractiveness of wagers, it is not the only factor. The positive or negative values assigned to the magnitude of the risk itself must also be taken into account.

We make a second assumption when we calculate expected value in terms of money. We assume that gains or losses of equal amounts of money represent gains or losses of equal amounts of value to the individual involved. In the language of the economist this is said to be the assumption that money has a constant marginal utility. This assumption is quite often false. For a poor man, the loss of 1000 dollars might mean he would starve, while the gain of 1000 dollars might mean he would merely live somewhat more comfortably. In this situation, the real loss accompanying a monetary loss of 1000 dollars is much greater than the real gain accompanying a monetary gain of 1000 dollars. A man in these circumstances would be foolish to accept an even money bet of 1000 dollars on the flip of a coin. In terms of money, the wager has an expected value of 0. But in terms of real value, the wager has a negative expected value.

Suppose you are in a part of the city far from home. You have lost your wallet and only have a quarter in change. Since the bus fare home is 35 cents, it looks as though you will have to walk. Now someone offers to flip you for a dime. If you win, you can ride home. If you lose, you are hardly any worse off than before. Thus, although the expected value of the wager in monetary terms is 0, in terms of real value, the wager has a positive expected value. In assessing the attractiveness of wagers by calculating their expected value, we must always be careful to see whether the monetary gains and losses accurately mirror the real gains and losses to the individual involved.

Exercises

1. What is the expected value of the following gamble? You are to roll a pair of dice. If the dice come up a natural, 7 or 11, you win 10 dollars. If the dice come up snake-eyes, 2, or boxcars, 12, you lose 20 dollars. Otherwise the bet is off.
2. What is the expected value of the following gamble? You are to flip a fair coin. If it comes up heads you win 1 dollar, and the wager is over. If it comes up tails you lose 1 dollar, but you flip again for 2 dollars. If the coin comes up heads this time you win 2 dollars. If it comes up tails you lose 2 dollars, but flip again for 4 dollars. If it comes up heads you win 4 dollars. If it comes up tails you lose 4 dollars. But in either case the wager is over.

Hint: The possible outcomes are:

Toss 1	Toss 2	Toss 3
H	None	None
T	H	None
T	T	H
T	T	T

3. Suppose you extended the doubling strategy of Exercise 2 to four tosses. Would this change the expected value?
4. Suppose that you tripled your stakes instead of doubling them. Would this change the expected value?

VI.6. BAYES' THEOREM. You may wonder what the relation is between a conditional probability $\Pr(q \text{ given } p)$ and its converse $\Pr(p \text{ given } q)$. They need not be equal. The probability that Ezekial is an ape, given that he is a gorilla, is 1. But the probability that Ezekial is a gorilla, given that he is an ape, is less than 1. The value of a conditional probability is not determined by the value of its converse alone. But the value of a conditional probability can

Table VI.2

Step	Justification
1. $\Pr(q \text{ given } p) = \frac{\Pr(p \& q)}{\Pr(p)}$	Definition of conditional probability
2. $\Pr(q \text{ given } p) = \frac{\Pr(p \& q)}{\Pr(p \& q) \vee (p \& \sim q)}$	p is logically equivalent to $(p \& q) \vee (p \& \sim q)$
3. $\Pr(q \text{ given } p) = \frac{\Pr(p \& q)}{\Pr(p \& q) + \Pr(p \& \sim q)}$	Special disjunction rule
4. $\Pr(q \text{ given } p) =$ General conjunction rule	
$\frac{\Pr(q) \times \Pr(p \text{ given } q)}{[\Pr(q) \times \Pr(p \text{ given } q)] + [\Pr(\sim q) \times \Pr(p \text{ given } \sim q)]}$	

be calculated from the value of its converse, together with certain other probability values. The basis of this calculation is set forth in *Bayes' theorem*. A simplified version of a proof of Bayes' theorem is presented in Table VI.2. Step 4 of this table states the simplified version of Bayes' theorem.⁹ Note that it allows us to compute conditional probabilities going in one direction—that is, $\Pr(q \text{ given } p)$ —from conditional probabilities going in the opposite direction—that is, $\Pr(p \text{ given } q)$ and $\Pr(p \text{ given } \sim q)$ —together with certain statement probabilities—that is, $\Pr(q)$ and $\Pr(\sim q)$. Let us see how this theorem is applied in a concrete example.

Suppose we have two urns. Urn 1 contains eight red balls and two black balls. Urn 2 contains two red balls and eight black balls. Someone has selected an urn by flipping a fair coin. He then has drawn a ball from the urn he

⁹ The general form of Bayes' theorem arises as follows: Suppose that instead of simply the two statements q and $\sim q$ we consider a set of n mutually exclusive statements, q_1, q_2, \dots, q_n , which is *exhaustive*. That is, the complex statement, $q_1 \vee q_2 \vee \dots \vee q_n$, is a tautology. Then it can be proven that the simple statement p is logically equivalent to the complex statement $(p \& q_1) \vee (p \& q_2) \vee \dots \vee (p \& q_n)$. This substitution is made in step 2, and the rest of the proof follows the model of the proof given. The result is

$$\Pr(q_1 \text{ given } p) = \frac{\Pr(q_1) \times \Pr(p \text{ given } q_1)}{[(\Pr(q_1) \times \Pr(p \text{ given } q_1))] + [\Pr(q_2) \times \Pr(p \text{ given } q_2)] + \dots + [\Pr(q_n) \times \Pr(p \text{ given } q_n)]}$$

selected. Assume that each ball in the urn he selected had an equal chance of being drawn. What is the probability that he selected urn 1, given that he drew a red ball? Bayes' theorem tells us the $\Pr(\text{urn 1 given red})$ is equal to

$$\frac{\Pr(\text{urn 1}) \times \Pr(\text{red given urn 1})}{[\Pr(\text{urn 1}) \times \Pr(\text{red given urn 1})] + [\Pr(\sim\text{urn 1}) \times \Pr(\text{red given } \sim\text{urn 1})]}$$

The probabilities needed may be calculated from the information given in the problem:

$$\Pr(\text{urn 1}) = \frac{1}{2}$$

$$\Pr(\sim\text{urn 1}) = \Pr(\text{urn 2}) = \frac{1}{2}$$

$$\Pr(\text{red given urn 1}) = \frac{8}{10}$$

$$\Pr(\text{red given } \sim\text{urn 1}) = \Pr(\text{red given urn 2}) = \frac{2}{10}$$

If these values are substituted into the formula, they give

$$\Pr(\text{urn 1 given red}) = \frac{\frac{1}{2} \times \frac{8}{10}}{(\frac{1}{2} \times \frac{8}{10}) + (\frac{1}{2} \times \frac{2}{10})} = \frac{\frac{4}{10}}{\frac{4}{10} + \frac{1}{10}} = \frac{4}{5}$$

A similar calculation will show that $\Pr(\text{urn 2 given red}) = \frac{1}{5}$. Thus, the application of Bayes' theorem confirms our intuition that a red ball is more likely to have come from urn 1 than urn 2, and it tells us how much more likely.

It is important to emphasize the importance of the pure statement probabilities $\Pr(q)$ and $\Pr(\sim q)$ in Bayes' theorem. If we had not known that the urn to be drawn from had been selected by flipping a fair coin, if we had just been told that it was selected some way or other, we could not have computed $\Pr(\text{urn 1 given red})$. Indeed if $\Pr(\text{urn 1})$ and $\Pr(\sim\text{urn 1})$ had been different, then our answer would have been different. Suppose that the urn had been selected by throwing a pair of dice. If the dice came up "snake-eyes" (a 1 on each die), urn 1 would be selected; otherwise urn 2 would be selected. If this were the case, then $\Pr(\text{urn 1}) = \frac{1}{36}$ and $\Pr(\sim\text{urn 1}) = \Pr(\text{urn 2}) = \frac{35}{36}$. Keeping the rest of the example the same, Bayes' theorem gives

$$\Pr(\text{urn 1 given red}) = \frac{\frac{1}{36} \times \frac{8}{10}}{(\frac{1}{36} \times \frac{8}{10}) + (\frac{35}{36} \times \frac{2}{10})} = \frac{\frac{8}{360}}{\frac{8}{360} + \frac{70}{360}} = \frac{8}{78} = \frac{4}{39}$$

This is quite a different answer from the one we got when urns 1 and 2 had an equal chance of being selected. In each case $\Pr(\text{urn 1 given red})$ is higher than $\Pr(\text{urn 1})$. This can be interpreted as saying that in both cases the additional information that a red ball was drawn would raise confidence that urn 1 was selected. But the initial level of confidence that urn 1 was selected is different in the two cases, and consequently the final level is also.

Exercises

1. The probability that George will study for the test is $\frac{4}{10}$. The probability that he will pass, given that he studies, is $\frac{9}{10}$. The probability that he passes, given that he does not study, is $\frac{3}{10}$. What is the probability that he has studied, given that he passes?
2. Suppose there are three urns. Urn 1 contains six red balls and four black balls. Urn two contains nine red balls and one black ball. Urn 3 contains five red balls and five black balls. A ball is drawn at random from urn 1. If it is black a second ball is drawn at random from urn 2, but if it is red the second ball is drawn at random from urn 3.
 - a. What is the probability of the second ball being drawn from urn 2?
 - b. What is the probability of the second ball being drawn from urn 3?
 - c. What is the probability that the second ball drawn is black, given that it is drawn from urn 2?
 - d. What is the probability that the second ball drawn is black, given that it is drawn from urn 3?
 - e. What is the probability that the second ball is black?
 - f. What is the probability that the second ball was drawn from urn 2, given that it is black?
 - g. What is the probability that the second ball was drawn from urn 3, given that it is black?
 - h. What is the probability that the second ball drawn was drawn from urn 2, given that it is red?
 - i. What is the probability that the second ball drawn was drawn from urn 3, given that it is red?
 - j. What is the probability that the first ball drawn was red, given that the second ball drawn is black?
 - k. What is the probability that the first ball is black, given that the second ball is black?
 - l. What is the probability that both balls drawn are black?
3. A fair coin is flipped twice. The two tosses are independent. What is the probability of a heads on the first toss given a heads on the second toss?
4. Their captors have decided that two of three prisoners—Smith, Jones, and Fitch—will be executed tomorrow. The choice has been made at random, but the identity of the unfortunate selectees is to be kept from the prisoners until the final hour. The prisoners, who are held in separate cells, unable to communicate with each other, know this. Fitch asks a guard to tell the name of one of the other prisoners who will be executed. Regardless of whether Fitch was chosen or not, one of the others will be executed, so the guard reasons that he is not giving Fitch any illicit information by answering truthfully. He says: “Jones will be executed.” Fitch is heartened by the news for he reasons that his probability of being the one who escapes execution has risen from $\frac{1}{3}$ to $\frac{1}{2}$. Has Fitch made a mistake? Has the guard? Use Bayes’ theorem to analyze the reasoning involved. (Hint: Calculate

the probability that Fitch will not be executed given that *the guard tells him that Jones will be executed*, not the probability that Fitch will not be executed given that Jones will be. What assumptions are possible about the probability that the guard tells Fitch that Jones will be executed given that Fitch escapes execution?)

VI.7. PROBABILITY AND CAUSALITY. What is meant when it is said that smoking causes lung cancer? Not that smoking is a *sufficient* condition for contraction of lung cancer, for many people smoke and never contract the disease. Not that smoking is a *necessary* condition for lung cancer, for some who never smoke nevertheless develop lung cancer. What is meant is something probabilistic: that smoking increases one's chances of getting lung cancer.

We might say that smoking has a tendency in the direction of sufficientness if $\Pr(\text{cancer given smoking})$ is greater than $\Pr(\text{cancer given } \sim\text{smoking})$ —that is, if smoking is *positively statistically relevant* to cancer. We might say that smoking has a tendency in the direction of necessaryness for lung cancer if $\Pr(\text{having smoked given cancer})$ is greater than $\Pr(\text{having smoked given no cancer})$ —that is, if cancer is positively statistically relevant to smoking. But we can show from the probability calculus that for any two statements, P, Q ,¹⁰ P is positively statistically relevant to Q if and only if Q is positively statistically relevant to P . By Bayes' theorem:

$$\Pr(Q \text{ given } P) = \frac{\Pr(P \text{ given } Q) \Pr(Q)}{\Pr(P)}$$

$$\text{So: } \frac{\Pr(Q \text{ given } P)}{\Pr(Q)} = \frac{\Pr(P \text{ given } Q)}{\Pr(P)}$$

P is positively relevant to Q just in case the left-hand side of the equation is greater than one; Q is positively relevant to P just in case the right-hand side of the equation is greater than one. So the probabilistic notions of being a tendency toward a sufficient condition, and having a tendency toward being a necessary condition come to the same thing! Considerations appear to be simpler in this way in a probabilistic setting than in a deterministic one.

But there is a complication that we must now discuss. Suppose that smoking itself did not cause the cancer, but that desire to smoke and cancer were both effects of some underlying genetically determined biological condition. Then smoking would still be positively statistically relevant to cancer, but as a symptom of having the bad gene rather than as a cause of cancer. If this hy-

¹⁰With positive probability.

hypothesis were correct, we would not say that smoking raised one's *chances of* getting lung cancer. If someone, say *you*, had the bad genes, then your chances of contracting cancer would be already high and smoking would not make them worse; if you didn't have the bad genes, your chances of contracting cancer would be lower and smoking wouldn't make them worse. That is, the positive statistical relevance of smoking to cancer would disappear if we looked at probabilities *conditional* on having the bad genes; likewise if we looked at probabilities conditional on not having the bad genes:

$$\Pr(\text{cancer given smoking and bad genes}) = \Pr(\text{cancer given bad genes})$$

$$\Pr(\text{cancer given smoking and good genes}) = \Pr(\text{cancer given good genes})$$

To support the claim that smoking is a probabilistic cause of lung cancer, the foregoing hypothesis (and others like it) must be ruled out. Perhaps identical twins can be found such that one of each pair is a long-time smoker, and more of the smokers develop cancer. Perhaps subjects who don't want to smoke but are forced to inhale smoke anyway (certain laboratory mice, cocktail waitresses, and so on) have a higher incidence of lung cancer.

If we believe that a certain constellation of factors determines the chance of getting lung cancer, then we consider smoking a probabilistic cause of lung cancer if, when we hold all the other preexisting factors *fixed*, smoking increases the chance of lung cancer. That is, if:

$$\Pr(\text{cancer given background factors and smoking}) \text{ is greater than } \Pr(\text{cancer given background factors and no smoking})$$

Whether X is a probabilistic cause of Y for individual a may depend on just what constellation of background factors is present for a . Some lucky people have a biochemistry such that for them, contact with poison oak is not a probabilistic cause of skin eruptions and intense itching, but for most of us it unfortunately is.

Exercises

1. Discuss the following argument: Most heroin users have previously smoked marijuana. Therefore, marijuana use causes heroin use.
2. How would you go about finding out whether *for you* exposure to ragweed pollen is a cause of a stuffed-up nose, runny eyes, and so on?
3. Some studies have found that, on average, convicted criminals exhibit vitamin deficiencies. This suggests to some researchers that vitamin deficiencies might

affect personality and lead to criminal behavior. An alternative hypothesis might be that, in many cases, there is a common cause—heroin addiction—that leads to both criminal behavior and malnutrition with its attendant vitamin deficiencies. Can you think of any other possible hypotheses? What sort of data would you gather to decide which is correct?

VII

Kinds of Probability

VII.1. INTRODUCTION. Historically, a number of distinct but related concepts have been associated with the word *probability*. These fall into three families: rational degree of belief, relative frequency, and chance. Each of the probability concepts can be thought of as conforming to the mathematical rules of probability calculus, but each carries a different meaning. We have, in one way or another, met each of these probability concepts already in this book. A biased coin has a certain objective *chance* of coming up heads. If we are uncertain as to how the coin is biased and what the objective chance really is, we may have a rational *degree of belief* that the coin will come up heads that is unequal to the true chance. If we flip the coin a number of times, a certain percentage of the tosses will come up heads; that is, the *relative frequency* of heads in the class of tosses will be a number in the interval from 0 to 1. The relative frequency of heads may well differ from both our degree of belief that the coin will come up heads and the objective chance that the coin comes up heads. The concepts are distinct, but they are closely related. Observed *relative frequencies* are important evidence that influences our *rational degrees of belief* about *objective chances*. If, initially, we are unsure whether the coin is biased 2 to 1 in favor of heads or 2 to 1 in favor of tails (degree of belief $\frac{1}{2}$), and then we flip the coin 1000 times and get 670 heads, we will have gotten strong evidence indeed that the coin is biased toward heads. This chapter is devoted to a review of these conceptions of probability and a sketch of their interrelation.

VII.2. RATIONAL DEGREE OF BELIEF. Belief is not really an all or nothing affair; it admits of degrees. You might be reasonably sure that the president was guilty without being absolutely certain. You might be extremely dubious about the plaintiff's supposed whiplash injury without being certain that he is malingering. You might think of it as only slightly more likely than not that the cause of a sore throat is a virus. Degrees of belief can be represented numerically, with larger numbers corresponding to stronger beliefs. What should the mathematics of these numbers be for a rational agent?

Rational degrees of belief are used to make rational decisions on the basis of expected values. If probabilities are to be used to calculate expected value of gambles, there are elementary practical reasons for the mathematical rules of the probability calculus. Let us assume for simplicity that money has constant marginal utility and that there is no value, positive or negative, attached

to risk itself. Then the expected value that you attach to a wager that pays you 1 dollar if p and such that you lose nothing if $\sim p$, is just $\$Pr(p)$.

$$\$1 \times Pr(p) + \$0 \times Pr(\sim p) = \$Pr(p)$$

Then a *tautology* should have probability 1 because the wager: “You get 1 dollar if a tautology is true, nothing otherwise” should have a value of 1 dollar—a payoff that you are certain to get because tautologies are true in all cases. Likewise, a *contradiction* should get probability 0 because the wager “You get 1 dollar if a contradiction is true, nothing otherwise” is obviously worth nothing.

Nothing should get a probability greater than 1 or less than 0. If a statement p got a probability greater than 1, the wager “You get 1 dollar if p , nothing otherwise” would get an expected value greater than 1 dollar, an expected value greater than anything you could possibly win. If a statement p got a probability less than 0, then the wager “You get 1 dollar if p , nothing otherwise,” which you can win but not lose, would get negative expected value; and the wager “You get -1 dollar (that is, you lose 1 dollar) if p , nothing otherwise,” which you can lose but not win, would get positive expected value.

We can also give a gambling rationale for the special disjunction rule. Suppose that $p; q$ are mutually exclusive. Then a bet “1 dollar if p , nothing otherwise,” taken together with a bet “1 dollar if q , nothing otherwise” gives a total payoff of 1 dollar if p is true or if q is true, nothing if both are false (they can’t both be true because they are incompatible). That is, the bet on p and on q , taken together, is tantamount to the bet: “1 dollar if p or q , nothing otherwise.” If I’m willing to pay $\$Pr(p)$ for the first bet and then $\$Pr(q)$ for the second one, I will have paid $\$Pr(p) + Pr(q)$ for the lot. If my evaluations are consistent, I should then give that same value to the bet “1 dollar if p or q , nothing otherwise.” But the expected value of this bet is $Pr(p \vee q)$, so I have the special disjunction rule:

$$Pr(p \vee q) = Pr(p) + Pr(q)$$

(when $p; q$ are mutually exclusive)

All the rules of the probability calculus can be shown to follow from those justified here.

VII.3. UTILITY. We have been operating so far within a set of assumptions that often approximate the truth for monetary gambles at small stakes. It is time to take a more global viewpoint and question these assumptions.

An extra 100 dollars means less to a millionaire than to an ordinary person. But if I win a million, I’m a millionaire. So the difference between winning a

million + 100 dollars and winning a million means less to me than the difference between winning 100 dollars and winning nothing. In the terminology of economics, money has decreasing rather than constant marginal utility for me.

The idea of utility was introduced into the literature on gambling in this connection by Daniel Bernoulli in 1738. Bernoulli was concerned with the St. Petersburg game. In this game, you flip a fair coin until it comes up heads. If it comes up heads on the first toss, you get 2 dollars; if on the second toss, 4 dollars; if on the third toss, 8 dollars; if on the n th toss, 2^n dollars. The expected dollar value of this game is infinite. (*Exercise: check this!*) How much would you pay to get into this game? Bernoulli's idea was that if the marginal utility of money decreased in the right way, the St. Petersburg game could have a reasonable finite expected *utility* even though the monetary expectation is infinite.

When we consider decisions whose payoffs are in real goods rather than money, there is another complication we must take into account. That is, the value of having two goods together may not be equal to the sum of their individual values because of interactions between the goods. If a man wants to start a pig farm, and getting a sow has value b for him and getting a boar has value c , then getting both a sow and a boar may have value greater than $b + c$. The sow and the boar are, for him, *complementary goods*. Interaction between goods can also be negative, as in the case of the prospective chicken farmer who wins two roosters in two lotteries. The presence of an active market reduces, but does not eliminate the effect of complementarities. The second rooster is still of more value to a prospective chicken farmer in Kansas than to Robinson Crusoe. The farmer can, for example, swap it for a hen; or at least sell it and put the money toward a hen. Because of complementarities, we cannot in general assume that if a bettor makes a series of bets each of which he considers to be fair, he will judge the result of making them all together as fair. Where payoffs interact, the right hand may need to know what the left is doing.

The preceding points about how utility works are intuitively easy to grasp. But it is harder to say just what utility is. We know how to count money, pigs, and chickens; but how do we measure utility? Von Neumann and Morgenstern showed how to use the expected utility principle to measure utility if we have some *chance* device (such as a wheel of fortune, a fair coin, a lottery) for which we know the *chances*. We pick the best payoff in our decision problem and give it (by convention) utility 1; likewise, we give the worst payoff utility 0. Then we measure the utility of a payoff, P , in between by judging what sort of a gamble with the worst and the best payoffs as possible outcomes has value equal to P . For instance, farmer Jones wants a horse, a pig, a chicken, and a husband. Her current decision situation is structured so that she will get exactly one of these. She ranks the payoffs:

1. Horse
2. Husband
3. Pig
4. Chicken

She is indifferent between (1) a lottery that gives $\frac{4}{5}$ chance of a horse and $\frac{1}{5}$ chance of a chicken, and one that gives a husband for sure; and (2) a lottery that gives $\frac{1}{2}$ chance of a horse and $\frac{1}{2}$ chance of a chicken, and one that gives a pig for sure. Thus, her utility scale looks like this:

	<i>Utility</i>
Horse	1
Husband	.8
Pig	.5
Chicken	0

If her decision situation were structured so that she might end up with all these goods, and if they didn't interfere with one another, then her utility scale might have a different top: Horse and Husband and Pig and Chicken. If it were structured so that she might end up getting none of these goods, after going to some expense, there might be a different bottom, which would have utility 0.

Utility, as measured by the von Neumann–Morgenstern method, is *subjective* utility, determined by the decision maker's own preferences. There are, no doubt, various senses in which a decision maker can be wrong about what is good for him. However, such questions are not addressed by this theory.

From a decision maker's utilities we can infer his degrees of belief. Farmer Smith has bought two tickets to win for a race at the county fair, one on Stewball and one on Molly. If he holds a ticket on a winning horse, he wins a pig; otherwise he gets nothing. We assume that he does not care about the outcome of the horse race per se; it is important to him only insofar as it does or does not win him a pig. He is indifferent to keeping his ticket on Stewball or exchanging it for an objective lottery ticket with a known 10 percent chance of winning; likewise for Molly or an objective lottery ticket with a 15 percent chance of winning.

Farmer Smith's utility scale looks like this:

	<i>Utility</i>
Pig	1
Ticket on Molly	.15
Ticket on Stewball	.10
Nothing	0

If he maximizes expected utility, his expected utility for his bet (ticket) on Molly is:

$$\text{Degree of Belief (Molly Wins) Utility (Pig) +} \\ \text{Degree of Belief (Molly Loses) Utility (Nothing)}$$

This is just equal to his degree of belief that Molly wins. Then his *degree of belief* that Molly wins is .15; in the same way, his degree of belief that Stewball wins is .10. Subjective degrees of belief are here recovered from subjective utilities in an obvious and simple way. (Things would be more complicated if farmer Smith cared about the outcome of the race over and above the question of the pig, but as we shall see in the next section, his subjective degrees of belief could still be found.)

Exercises

1. A decision maker with declining marginal utility of money is *risk averse* in monetary terms. He will prefer 50 dollars for sure to a wager that gives a chance of $\frac{1}{2}$ of winning 100 dollars and a chance of $\frac{1}{2}$ of winning nothing, because the initial 50 dollars has more utility for him than the second 50 dollars. Suppose that winning 100 dollars is the best thing that can happen to him and winning nothing is the worst.
 - a. What is his utility for winning 100 dollars?
 - b. What is his utility for winning nothing?
 - c. What is his utility for a wager that gives a known objective chance of $\frac{1}{2}$ of winning 100 dollars and $\frac{1}{2}$ of winning nothing?
 - d. What can we say about his utility for getting 50 dollars?
 - e. Draw a graph of utility as against money for a decision maker who is generally risk averse.
2. Suppose farmer Smith has one ticket on each horse running at the county fair, and thus will win a pig no matter which horse wins. Let $U(\text{pig}) = 1$ and $U(\text{nothing}) = 0$. Suppose farmer Smith's preferences go by expected utility.
 - a. Farmer Smith believes that all his tickets taken together are worth one pig for sure. What does this tell you about his degrees of belief about the race?
 - b. Suppose that farmer Smith also believes that each of his tickets has equal utility. What does this tell you about his degrees of belief about the race?

VII.4. RAMSEY. The von Neumann–Morgenstern theory of utility is really a rediscovery of ideas contained in a remarkable essay, “Truth and Probability,” written by F. P. Ramsey in 1926. In the essay, Ramsey goes even deeper into the foundations of utility and probability. The von

Neumann–Morgenstern method requires that the decision maker know some objective chances, which are then used to scale his subjective utilities. From his subjective utilities and preferences, information about his subjective probabilities can be recovered. Ramsey starts without the assumption of knowledge of some chances, and with only the decision maker's preferences.

Ramsey starts by identifying propositions that, like the coin flips, lotteries, and horse races of the previous section, have no value to the decision maker in and of themselves, but only insofar as certain payoffs hang on them. He calls such propositions "ethically neutral." A proposition, p , is ethically neutral for a collection of payoffs, B , if it makes no difference to the agent's preferences, that is, if he is indifferent between B with p true and B with p false. A proposition, p , is ethically neutral if p is ethically neutral for maximal collections of payoffs relevant to the decision problem. The nice thing about ethically neutral propositions is that the expected utility of gambles on them depends only on their probability and the utility of their outcomes. Their own utility is not a complicating factor.

Now we can identify an ethically neutral proposition, H , as having probability $\frac{1}{2}$ for the decision maker if there are two payoffs, $A;B$, such that he prefers A to B but is indifferent between the two gambles: (1) Get A if H is true, B if H is false; (2) get B if H is true, A if H is false. (If he thought H was more likely than $\sim H$, he would prefer gamble 1; if he thought $\sim H$ was more likely than H , he would prefer gamble 2.) For the purpose of scaling the decision maker's utilities, such a proposition is just as good as the proposition that a fair coin comes up heads.

The same procedure works in general to identify surrogates for fair lotteries. Suppose there are 100 ethically neutral propositions, $H_1;H_2; \dots ;H_{100}$, which are pairwise incompatible and jointly exhaustive. Suppose there are 100 payoffs, $G_1;G_2; \dots ;G_{100}$, such that G_1 is preferred to G_2 , G_2 is preferred to G_3 , and so forth up to G_{100} . Suppose the decision maker is indifferent between the complex gamble:

If H_1 get G_1 &
 If H_2 get G_2 &
 If H_i get G_i &
 If H_{100} get G_{100}

and every other complex gamble you can get from it by moving the G_i s around. Then each of the H_i s gets probability .001, and together they are just as good as a fair lottery with 100 tickets for scaling the decision maker's utilities.

A rich enough preference ordering has enough ethically neutral propositions forming equiprobable partitions of the kind just discussed to carry out

the von Neumann–Morgenstern type of scaling of utilities described in the last section to any desired degree of precision. Once the utilities have been determined, the degree of belief probabilities of the remaining ethically neutral propositions can be determined in the simple way we have seen before. The decision maker’s degree of belief in the ethically neutral proposition, p , is just the utility he attaches to the gamble: *Get G if p, B otherwise*, where G has utility 1 and B has utility 0.

With utilities in hand, we can also solve for the decision maker’s degrees of belief in non-ethically neutral propositions, although things are not quite so simple here. Suppose that farmer Smith owned Stewball and wanted his horse to win, as well as wanting to win a pig. Then “Stewball wins” and “Molly wins” are not ethically neutral for him. Now suppose we want to determine his degree of belief in the proposition that Molly wins. Given our conventions, we can’t set up a gamble that gives utility 1 if Molly wins because what farmer Smith desires most and gives utility 1 is: “Get a pig and Stewball wins.” But we know that the expected utility of the wager “Pig if Molly wins, no prize if she loses” is equal to:

$$\Pr(\text{Molly wins}) U(\text{get pig and Molly wins}) + \\ 1 - \Pr(\text{Molly wins}) U(\text{no prize and Molly loses})$$

If we know the utility of the wager, of “Get pig and Molly wins,” and of “No prize and Molly loses,” we can solve for $\Pr(\text{Molly wins})$.

For a rich and coherent preference ordering over gambles, Ramsey has conjured up both a subjective utility assignment and a degree of belief probability assignment such that preference goes by expected utility. This sort of representation theorem shows how deeply the probability concept is rooted in practical reasoning.

Exercises

1. Suppose that the four propositions, $HH;HT;TH;TT$, are pairwise incompatible (at most one of them can be true) and jointly exhaustive (at least one must be true). Describe the preferences you would need to find to conclude that they are ethically neutral and equiprobable.
2. Suppose that farmer Smith owns Stewball and that “Molly wins” is not ethically neutral. His most preferred outcome is “Get pig and Stewball wins;” his least preferred is “No pig and Stewball loses;” therefore, these get utility 1 and 0, respectively. Propositions, $HH;HT;TH;TT$ are as in Exercise 1. Farmer Smith is indifferent between “Get pig and Molly wins” and a hypothetical gamble that would ensure that he would get the pig and Stewball would win if HH or HT or TH and that he would get no pig and Stewball would lose if TT . (What does this tell you about his utility for “Get pig and Molly wins?”) He is indifferent between

“No pig and Molly loses” and the hypothetical gamble that would ensure that he would get the pig and Stewball would win if HH and that he would get no pig and Stewball would lose if HT or TH or HH . He is indifferent between the gamble “Pig if Molly wins; no pig if she loses” and the gamble “Get pig and Stewball wins if HH or HT , but no pig and Stewball loses if TH or TT .”

- a. What are his utilities for “Get pig and Molly wins;” “No pig and Molly loses;” the gamble “Pig if Molly wins; no pig if Molly loses?”
- b. What is his degree of belief probability that Molly will win?

VII.5. RELATIVE FREQUENCY. If I flip a coin ten times and it comes up heads six of those times, the observed relative frequency of heads is .6; six heads out of ten trials. If our language contains the means of reporting the outcomes of single trials and contains standard propositional logic, it already contains the means for reporting the relative frequencies in finite numbers of trials. For example, (supposing that $H1$ means “heads on trial 1,” etc.) we can render “relative frequency of heads equals .1” as the disjunction of those sequences of outcomes with exactly one head and nine tails:

$H1 \& T2 \& T3 \& T4 \& T5 \& T6 \& T7 \& T8 \& T9 \& T10$	or
$T1 \& H2 \& T3 \& T4 \& T5 \& T6 \& T7 \& T8 \& T9 \& T10$	or
$T1 \& T2 \& H3 \& T4 \& T5 \& T6 \& T7 \& T8 \& T9 \& T10$	or
$T1 \& T2 \& T3 \& H4 \& T5 \& T6 \& T7 \& T8 \& T9 \& T10$	or
$T1 \& T2 \& T3 \& T4 \& H5 \& T6 \& T7 \& T8 \& T9 \& T10$	or
$T1 \& T2 \& T3 \& T4 \& T5 \& H6 \& T7 \& T8 \& T9 \& T10$	or
$T1 \& T2 \& T3 \& T4 \& T5 \& T6 \& H7 \& T8 \& T9 \& T10$	or
$T1 \& T2 \& T3 \& T4 \& T5 \& T6 \& T7 \& H8 \& T9 \& T10$	or
$T1 \& T2 \& T3 \& T4 \& T5 \& T6 \& T7 \& T8 \& H9 \& T10$	or
$T1 \& T2 \& T3 \& T4 \& T5 \& T6 \& T7 \& T8 \& T9 \& H10$	

(*Exercise:* Write out the comparable description for “relative frequency of heads equals .2 in ten trials.”) Relative frequencies of outcome types obey the laws of the probability calculus: They are always numbers in the interval from 0 to 1. The relative frequency of a tautological outcome type (for example, either heads or not) is 1 and that of a contradictory outcome type (for example, both heads and not) is 0. The relative frequency of two mutually exclusive outcome types is the sum of their individual relative frequencies.

Statements of observed relative frequencies constitute an especially important kind of evidence for statements about chance. In certain familiar cases they summarize everything that is relevant to the chances in a series of observed trials. Consider again the example of the coin with unknown bias.

Suppose that the coin is flipped ten times with the outcome being:

H1&T2&H3&H4&T5&H6&T7&H8&H9&H10

The relative frequency of heads is .7. Now suppose that we want to find the probability conditional on the outcome that the chance of heads is $\frac{2}{3}$. We use Bayes' theorem:

$$\begin{aligned} & \Pr[\text{chance } H = \frac{2}{3} / \text{outcome}] \\ &= \frac{\Pr[\text{outcome} / \text{chance } H = \frac{2}{3}] \Pr[\text{chance } H = \frac{2}{3}]}{\text{Sum}_i \Pr[\text{outcome} / \text{chance } H = i] \Pr[\text{chance } H = i]} \end{aligned}$$

The outcome evidence figures in Bayes' theory solely through the conditional probabilities of the outcome given the chance hypotheses. But for a given chance hypothesis, *every outcome sequence with seven heads out of ten tosses has the same probability conditional on the chance hypothesis*. (This is because the trials are independent, conditional on the chance hypotheses.) Any other outcome sequence of ten trials with the same relative frequency of heads would lead to the same calculation (order is not important). With regard to our inferences about chances here, all the relevant information in the outcome sequence is captured by a report of relative frequency. Relative frequency is said here to be a *sufficient statistic*. (It is a sufficient statistic when we are considering outcome sequences of fixed length. The pair consisting of the length of the outcome sequence together with the relative frequency constitutes a sufficient statistic in general for all outcome sequences.) Whenever the trials are independent, conditional on the chance hypotheses, relative frequency is a sufficient statistic. In this typical situation, reports of relative frequency capture all the salient information in the experimental results.

Reports of relative frequency become more compelling as the number of observed trials becomes large. Seven out of ten heads might incline us to believe that the coin is biased toward heads, but our inclination would turn into something stronger on the evidence of 70 out of 100 heads, let alone on 700 out of 1000. These gambler's intuitions are mathematically well founded. Suppose that you are in the coin tossing case, with a finite number of chance hypotheses, each of which has positive prior (degree of belief) probability and that the trials are independent, conditional on the chance hypotheses. Then as the number of trials becomes larger and larger, your probability that the relative frequency approximates the true chances to any tiny preassigned error tolerance approaches 1. This is a consequence of a theorem of the probability calculus known as the *weak law of large numbers*. Here, you *must* believe strongly that the relative frequencies will approximate the chances in the long run if your degrees of belief are coherent.

Here, observed relative frequencies can serve as evidence—evidence that becomes more and more conclusive as the number of trials approaches infinity—about the chances *via* Bayes' theorem.

Let us look at the use of Bayes' theorem a little more closely. The quantities that are crucial to the way in which the observed frequencies cut for and against the chance hypotheses are the conditional probabilities of outcomes conditional on the chance hypotheses. Now I am going to ask a question that may seem trivial. *What is the probability of getting heads conditional on the chance of heads being $\frac{2}{3}$?* The natural answer is $\frac{2}{3}$, and it may sound tautological. But remember, the conditional probability in question is a degree of belief. What we have is a principle connecting degree of belief and chance that is so obvious that it is constantly used but rarely stated. If our evidence for an outcome comes only from knowledge of the chances involved, then:

$$\text{Degree of belief}[\text{O given chance}(\text{O}) = x] = x$$

It is this principle that allows us to use Bayes' theorem in inductive reasoning about chances.

It also allows us to determine in a natural way degrees of belief about outcomes from degrees of belief about chances. Suppose that you are certain that the chance of heads is either $\frac{2}{3}$ or $\frac{1}{3}$ and that you regard those chance hypotheses as equally likely (that is, you give them each degree of belief $\frac{1}{2}$). Then what should your degree of belief be that the next coin flip comes up heads? A natural answer is to average the possible chances 50/50, getting $(\frac{1}{2})(\frac{2}{3}) + (\frac{1}{2})(\frac{1}{3}) = \frac{1}{2}$. The answer is a consequence of the foregoing principle and the probability calculus:

$$\begin{aligned} \text{Pr}(\text{H}) &= \text{Pr}[\text{Ch}(\text{H}) = \frac{2}{3} \ \& \ \text{H}] + \text{Pr}[\text{Ch}(\text{H}) = \frac{1}{3} \ \& \ \text{H}] \\ &= \text{Pr}[\text{Ch}(\text{H}) = \frac{2}{3}] \text{Pr}[\text{H}/\text{Ch}(\text{H}) = \frac{2}{3}] + \text{Pr}[\text{Ch}(\text{H}) = \frac{1}{3}] \text{Pr}[\text{H}/\text{Ch}(\text{H}) = \frac{1}{3}] \\ &= \text{Pr}[\text{Ch}(\text{H}) = \frac{2}{3}] \quad \left(\frac{2}{3}\right) \quad + \text{Pr}[\text{Ch}(\text{H}) = \frac{1}{3}] \quad \left(\frac{1}{3}\right) \end{aligned}$$

The proper degree of belief that the next flip comes up heads is an average of the possible chances, with the weights of the average being the degrees of belief in the proper chance hypotheses. We have met such weighted averages before when we discussed rational decision. They are called *expectations*. In connection with rational decision, we were interested in expected utility. The principle we have here is that *rational degree of belief is the expectation of chance*.

The preceding two paragraphs round out our quick survey of the connections between chance and the other conceptions of probability (although they may leave lingering questions about the credentials of the “natural” principles invoked). But, what is *chance* in and of itself? This question is highly controversial. I will attempt a rough and ready sketch of the main positions,

but the reader should be warned that I will not be able to do justice here to the subtle and sophisticated variations on these positions to be found in the literature. Basically, metaphysical views of chance fall into three main categories: (I) chance as primitive and irreducible, (II) chance as reducible in some way to relative frequency, (III) chance as reducible in some way to degree of belief.

Those who think of chance as an irreducible notion of physical tendency or propensity usually think of it as a theoretical concept on a par with the concept of force in physics. Physical theories stated in terms of chance permit predictions about the chances. But all that we can observe are the relative frequencies in sequences of outcomes. These are inductively rather than deductively linked to statements about chance. That is, our rational degrees of belief about the chances are influenced by observed relative frequencies. Conversely, our beliefs about the chances influence our anticipations about relative frequencies. In short, all that we know about chance in general is its connections with rational degree of belief and relative frequency. The main shortcoming of this view is that it has so little to say; its main strength is the shortcomings of competing views.

Frequentists think of chance as relative frequency *in the long run*—more precisely, as the limiting relative frequency in an infinite sequence of trials. Some versions of the theory (such as that of von Mises) add the proviso that the sequence of trials be random; some versions (such as that of Reichenbach) do not. Frequency views derive a certain motivation from limit theorems of the probability calculus, such as the law of large numbers. If we consider infinite sequences of independent trials,¹ the law of large numbers can be stated in a strong way: The probability that the chance equals the relative frequency is 1. Why not just say that chance *is* limiting relative frequency.

This view has its difficulties. The most obvious is that the appropriate infinite sequences may not exist. Suppose, with some plausibility, that my biased coin does not last for an infinite number of flips. Still, we want to say that on each flip there was a definite chance of that flip coming up heads. There is another problem as well in finding the appropriate infinite sequences. The physical factors determining chance may change. The coin may wear in such a way as to change the bias after a few million flips. Or, to vary the example, the chance of a newborn baby living to age 50 may change from decade to decade and country to country, and so on. For these reasons, sophisticated relative frequency views from the time of John Venn (1866) onward have had to talk about *hypothetical relative frequencies: what the limiting relative frequency*

¹And strengthen the additivity rule of the probability calculus to allow infinite additivity.

would have been if the experiment had been repeated (independently?) an infinite number of times such that the factors determining the chances did not change. Notice that this definition of chance, as it stands, is circular. It uses both the notion of the factors determining the chances, and, apparently, the notion of independent trials. A noncircular hypothetical relative frequency definition of chance would require some way of eliminating these references to chance.

Even if this were accomplished, there would be some question as to the grounding of the hypothetical. Improbability, even probability 0, is not impossibility. There is nothing in probability theory that says that it is impossible for a fair coin to come up heads every time in an infinite sequence of independent trials. What, then, does the frequentist take as grounding the truth of hypotheticals about what the limiting relative frequency would have been, in such a way that they capture the concept of chance? The idea of a relative frequency interpretation of chance offers to bring the concept of chance down to earth, but as the relative frequency interpretation becomes more sophisticated, it becomes more remote from the real events that constitute sequences of actual trials.

Personalists want to do without the concept of chance. The primary conception of probability for them is rational degree of belief. But they need to provide some explanation of what appears to be rational degrees of belief *about chance*, as in the case of the biased coin. They do this by showing how ordinary degrees of belief about the experimental setup and the outcome sequence can look like degrees of belief about chance.

Presumably there are certain overt physical facts that you take as determining the chances, for example, the shape and distribution of density of the coin. Let us say that your personal chance ("p-chance," for short) of heads is your degree of belief in heads conditional on the specification of these physical facts. For example, let LUMPH describe a coin with a lump on the heads side such that your degree of belief that the coin will come up heads conditional on that physical asymmetry of the coin is $\frac{2}{3}$. Then we will say that your p-chance of heads is $\frac{2}{3}$ in any situation in which LUMPH is a true description of the coin. Now the personalist can show that p-chances work just the way that chances are supposed to, and that this is a consequence of the probability calculus. Thus:

$$\text{Pr}[\text{heads}/\text{p-chance}(\text{heads}) = \frac{2}{3}] = \frac{2}{3}$$

from the definition of p-chance, and rational degree of belief is the expectation of p-chance.

Now, p-chances are *personal*; they come from one's degrees of belief. Different rational agents can assign different p-chances to the same possible

circumstances. This subjective character of p-chances, however, gives way to an objective determination in certain special circumstances. One sort of overt fact about the world that one might appeal to in constructing p-chances consists of the *relative frequency* of an outcome type (for example, heads) in an outcome sequence. If two rational agents have degrees of belief such that their p-chances constructed in this way give independent trials, then their p-chances must agree in the limiting case of infinite trials as a consequence of the law of large numbers. Here objective agreement is a consequence of a theorem of the probability calculus.

The analysis of the condition under which a degree of belief has the sort of structure such that conditioning on statements of relative frequency gives independence in the limit is due to de Finetti. It is that the degree of belief probability makes the trials *exchangeable*, which is to say that for any finite length subsequence, relative frequency is a sufficient statistic. In this case, personalism comes surprisingly close to frequentism: *Degree of belief is the expectation of limiting relative frequency*. This version of personalism inherits some of the strengths of frequentism, such as the objective character of laws of large numbers; but it also inherits some of the problems of frequentism, such as questions about the idealized character of infinite sequences of trials.

The question of the nature of chance is a difficult and controversial metaphysical topic. Propensity, frequency, and personalist views of chance seem initially to be quite distinct. But sophisticated versions of each of these views make contact with the central concepts and methods of the others, and there is considerable common ground between them.

Exercise

We both believe that a coin is either (A) biased 2 to 1 in favor of heads, or (B) biased 2 to 1 in favor of tails, but we differ as to how our initial degrees of belief are distributed between (A) and (B). You believe (A) to degree .8 and (B) to degree .2. I believe (A) to degree .3 and (B) to degree .7. We flip the coin five times and get one head and four tails. If we both conditionalize on this evidence, what are our final degrees of belief in (A) and (B)? (Use Bayes' theorem.)

VIII

Probability and Scientific Inductive Logic

VIII.1. INTRODUCTION. In Chapter II we discussed in general terms the relationship between probability and scientific inductive logic. Now that we have studied a bit of the mathematical theory of probability in Chapter VI and the interaction of various kinds of probability in Chapter VII, we can do a little more. In this chapter we revisit some of the topics discussed earlier to see what light our knowledge of probability can throw on them.

VIII.2. HYPOTHESIS AND DEDUCTION. Suppose we formulate a scientific hypothesis, deduce an observational prediction from it, test the prediction, and observe that it is false. Then the hypothesis is false. A deductively valid argument cannot lead from a true premise to a false conclusion, so if the argument from the hypothesis to the prediction is valid and the prediction is false, the hypothesis must also be false. We have investigated one kind of this sort of elimination of hypotheses when we discussed Mill's methods of experimental inquiry in Chapter V. In that chapter, the hypotheses were of some special form: " C is a necessary condition for E ," " C is a sufficient condition for E ," and so forth. The prediction obtained from the hypothesis said that a certain kind of event could not occur, for example " E cannot be present if C is absent." An observation of an occurrence that violates this prediction shows that the hypothesis is false.

Some philosophers, notably Sir Karl Popper, have held that this sort of deductive testing of hypotheses is all that there is to the logic of science. A hypothesis that passes such a test is said to be *corroborated*. Passing many observational tests is better than passing just one, and some tests are more telling than others. For instance, successful prediction of something that is unexpected is better than one of something that would be expected even if the hypothesis weren't true.

This account of the continual testing of hypotheses in science carries useful insights into scientific method, but it also leaves unanswered some important questions. Why are more stringent tests better than less stringent ones? Why is it better to pass more tests? How much better? More importantly, what are we to say about scientific hypotheses from which we cannot, strictly speaking, *deduce* predictions that can be conclusively refuted by observation? As an important example, consider probabilistic hypotheses that do not tell you what occurrences are impossible, but only which are more or less probable. Note

that a physical theory as basic as quantum mechanics generates just this kind of hypothesis.

Fortunately, we can show that probabilistic reasoning preserves the insights of deductive testing without such a drastic restriction of the range of inductive reasoning. Suppose that we have a hypothesis, H , that we wish to test. Our degree of belief in H is neither equal to zero (in which case we would consider it a non-starter) nor equal to one (in which case we would consider it conclusively established), but rather somewhere in between. Now we deduce a prediction, E , from H . Because of this deduction the conditional probability, $\Pr(E \text{ given } H)$ is equal to one and $\Pr(\sim E \text{ given } H)$ is equal to zero. We assume that our prior degree of belief in E is also neither zero nor one, but rather something in between. Now we look to see whether E is true. This may involve performing an experiment, or it may involve pure observation.

Suppose our observation shows that E is false. Then our new degree of belief in the hypothesis, H , is our old conditional probability, $\Pr(H \text{ given } \sim E)$. By Bayes' theorem this is:

$$\Pr(H \text{ given } \sim E) = \Pr(\sim E \text{ given } H) \Pr(H) / \Pr(\sim E).$$

But $\Pr(\sim E \text{ given } H)$ is zero because we deduced E from H . So the new probability of the hypothesis is zero. The hypothesis is refuted.

On the other hand, suppose that the hypothesis passes the test; E is shown to be true. Then the new probability of the hypothesis should be equal to $\Pr(H \text{ given } E)$. Let's say that the hypothesis is corroborated by the evidence if the new probability of the hypothesis is greater than its old probability, that is to say if the ratio:

$$\Pr(H \text{ given } E) / \Pr(H)$$

is greater than one. But probability theory tells us that:

$$\Pr(H \text{ given } E) / \Pr(H) = \Pr(E \text{ given } H) / \Pr(E)$$

Since we deduced E from H , $\Pr(E \text{ given } H)$ is equal to one, so:

$$\Pr(H \text{ given } E) / \Pr(H) = 1 / \Pr(E)$$

Since $\Pr(E)$ is less than one, the hypothesis is corroborated by passing the test, just as the deductivist believes that it should be. Furthermore, we get, as a little probabilistic bonus, an explanation of additional deductivist intuitions. The more surprising the corroborating phenomenon [that is, the smaller $\Pr(E)$] the greater the corroboration of the hypothesis [the greater $\Pr(H \text{ given } E) / \Pr(H)$].

Popper's deductive story of observational test of scientific theory is preserved as a special case by probabilistic reasoning, but notice that the

same sort of story now has been shown to apply much more generally. Suppose that the prediction, E , cannot strictly be deduced from the hypothesis, but can only be shown to have high probability given the hypothesis. For example, suppose that $\Pr(E \text{ given } H) = .999$. In this case, when we test for corroboration, we get:

$$\Pr(H \text{ given } E)/\Pr(H) = .999/\Pr(E)$$

Now E corroborates H if our prior degree of belief in E was not already .999 or more. The more improbable E was initially thought to be, the more powerfully it corroborates the hypothesis.

VIII.3. QUANTITY AND VARIETY OF EVIDENCE. Suppose that a new drug is being tested for serious side effects. It would be laughable if a pharmaceutical company applied to the Food and Drug Administration for approval to market the drug on the basis of experience with only one patient. Extensive testing is required to provide a reasonable assurance of public safety. Passing more tests is better than passing less. Why is this so? Deductive reasoning alone does not really provide an answer. Probabilistic reasoning does.

Let us simplify and suppose that the hypothesis, H , is *No liver damage when this drug is administered* and that evidence $E1$ is *There was no liver damage when this drug was administered to patient 1* and evidence $E1000$ is *There was no liver damage when this drug was administered to patients 1 through 1000*. We would like to know why $E1000$ is better evidence for H than $E1$. Assuming as background knowledge that the drug was administered, we have, as in the last section:

$$\begin{aligned} \Pr(H \text{ given } E1)/\Pr(H) &= 1/\Pr(E1) \text{ and} \\ \Pr(H \text{ given } E1000)/\Pr(H) &= 1/\Pr(E1000) \end{aligned}$$

Prior to the test, when the safety of the drug is in serious doubt, the probability that one patient escapes liver damage is higher than the probability 1000 do. That is to say, $\Pr(E1)$ is greater than $\Pr(E1000)$. It follows that H is better corroborated by 1000 trials than by just one. $\Pr(H \text{ given } E1000)$ is greater than $\Pr(H \text{ given } E1)$.

The same sort of reasoning explains the importance of variety of evidence. Compare two sorts of evidence: One, EH , reports no liver damage from 1000 trials in a homogeneous population of patients who are all between 20 and 15 years of age and in general good health. Another, EV , reports no liver damage in a sample of 1000 trials in a heterogeneous population of people of all ages and various degrees of health. With the safety of the drug unknown, the probability that liver damage would not show up in the homogeneous population is

higher than the probability that it would not show up in the varied population. That is to say, $\Pr(EH)$ is greater than $\Pr(EV)$. So, as before, $\Pr(H \text{ given } EV)$ is greater than $\Pr(H \text{ given } EH)$.

These examples explain the force of quantity and variety of evidence in a deductive setting where the hypothesis entails the evidence. The same general considerations apply when the connection between the hypothesis and the evidence is probabilistic.

Exercises

1. Suppose that an urn either has 99 red balls and 1 white ball (Hypothesis 1) or has 99 white balls and 1 red ball (Hypothesis 2). Suppose both hypotheses have initial probability of $\frac{1}{2}$. Consider two samples drawn with replacement. Sample 1 consists of 1 red ball. Sample 2 consists of 10 red balls. Find the $\Pr(\text{Hypothesis 1 given Sample 1})$ and $\Pr(\text{Hypothesis 1 given Sample 2})$.
2. Suppose that there are two urns, each of which has either 99 red balls and 1 white ball or 99 white balls and 1 red ball. For each urn both hypotheses have initial probability of $\frac{1}{2}$, and composition of the urns are independent of one another. (You could think of it this way: For each urn, a fair coin has been flipped to see whether it has mostly red balls or mostly white balls. The coin flips are independent. Consider the hypothesis, H , that both urns have 99 red balls and 1 white ball. Consider two possible samples. Sample 1 consists of 20 red balls from urn 1. Sample 2 consists of 10 red balls from urn 1 and 10 red balls from urn 2. Compute $\Pr(H)/\Pr(H \text{ given Sample 1})$ and $\Pr(H)/\Pr(H \text{ given Sample 2})$.

VIII.4. TOTAL EVIDENCE. Why should you use your total evidence in updating your degrees of belief? It doesn't cost you anything to use the evidence that you already have, and it is plausible to do so, but what is the reason that this is the rational thing to do? It can be shown that we should always expect better results from making decisions based on more knowledge rather than less.

Suppose you are bitten by a dog that was foaming at the mouth. You suspect that it may have had rabies. You now have to choose between taking the Pasteur treatment or not. The Pasteur treatment is inconvenient and painful; but if the dog did have rabies the treatment can save your life. Given your current estimate about how likely it is that the dog had rabies and your current utilities, you maximize expected utility by playing it safe and opting for the treatment.

But now the dog is captured, and you can have it tested for rabies before you make your decision whether to have the Pasteur treatment or not. The test will be quick and will not cost you anything. Will you decide now, or get the new evidence from the test results and then decide? It is obvious that the




Figure VIII.1 (Source: From *The Dynamics of Rational Deliberation* by Brian Skyrms. Copyright © 1990 by the President and Fellows of Harvard College. Reprinted by permission of Harvard University Press.)

best thing to do is to get the test results before deciding. The reason is equally obvious. If it turns out that the dog is not rabid you will be better off making a different decision than you would on the basis of your current information; you will be better off not taking the treatment. You maximize expected value by choosing to make your decision on probabilities that incorporate more information rather than less.

You can see this reasoning graphically in Figure VIII.1. The expected utilities of treatment and of no treatment are graphed as a function of the probability that the dog is rabid. The expected utility of each act is graphed as a straight line. At the right edge of the graph, the probability that the dog had rabies is equal to one, and the utility of treatment is higher than that of non-treatment. At the left edge of the graph, the probability that the dog has rabies is zero, and the expected utility of treatment is lower than that of non-treatment. The lines are straight because expected utility is an average.

For instance, if the probability that the dog is rabid is one-half, then the expected utility of treatment is midway between its value when the dog is surely rabid and its value when the dog is surely not rabid. The bold parts of the lines show which act is best at that probability. The upper horizontal dotted line

connecting treatment on the right and none on the left gives us the expected utility of getting the information as to whether the dog is rabid and then deciding. The upper horizontal line is higher than the boldfaced line (except at the edges) and this shows that the expected utility of getting information and then deciding is higher than the expected utility of making a decision without getting new information. This is an illustration of a general principle. If new information may change your optimal act, the expected utility of making an informed act is always greater than that of making an uninformed act.

Why then, should you use your total evidence to update your degrees of belief? Because you expect that this is the best way to get degrees of belief that are to be used for making rational decisions.

VIII.5. CONVERGENCE TO THE TRUTH. What about the problem of induction? Probability theory is not going to solve Hume's problem in Hume's terms. Probability theory isn't magic! But it can help us understand some aspects of the problem.

Suppose you are flipping a coin with unknown bias. You think that the coin is either biased 2 to 1 in favor of heads [$\text{Chance}(\text{heads}) = \frac{2}{3}$] or 2 to 1 in favor of tails [$\text{Chance}(\text{tails}) = \frac{2}{3}$]. You think that the tosses are objectively independent. That is to say that they are independent in the chance probabilities. You are uncertain as to whether the coin is biased toward heads or tails. For a definite example we can suppose that your degree of belief that the coin is biased toward heads is $\frac{1}{2}$ as is your degree of belief that it is biased toward tails, but what I am about to say would hold good for other non-extreme values of degree of belief.

You are going to flip the coin over and over, and you will update your degrees of belief in the chance hypotheses using Bayes' theorem. Now, given all this, how skeptical can you be, concerning your learning from experience? There are limits. Given your degrees of belief you must believe that if the coin is biased 2 to 1 in favor of heads, then in the long run the observed relative frequency of heads will get closer and closer to $\frac{2}{3}$, as discussed in Chapter VII. Then if you update your degrees of belief using Bayes' theorem, your degree of belief that $\text{Chance}(\text{heads}) = \frac{2}{3}$ will get closer and closer to 1. Likewise, you must believe that if the coin is biased 2 to 1 in favor of tails, then in the long run the relative frequency of heads will get closer and closer to $\frac{1}{3}$, and your degree of belief that $\text{Chance}(\text{heads}) = \frac{1}{3}$ will get closer and closer to 1. You must believe that you will learn the true chances! Knowing the true chances does not tell you for sure whether that coin will come up on the next toss or not. It only gives you the chances. But we can't expect any more. That's as good as it gets.

You might legitimately worry about use of "the long run" in this discussion. Must you be prepared to wait forever? Not really. I won't work out the details

here, but I can say this: Given that your degrees of belief are as stated in the example, you must be very sure that you will assign very high probability to the true chances after about 100 flips of the coin.

So far, so good. But what if you are wrong about the coin either being biased so that $\text{Chance}(\text{heads}) = \frac{2}{3}$ or being biased so that $\text{Chance}(\text{heads}) = \frac{1}{3}$? Suppose that the coin is fair and $\text{Chance}(\text{heads}) = \frac{1}{2}$. Then no matter how many coin flips you observe, you will never learn the true chances! How can you be certain that you will learn the true chances? You are certain that you will learn the true chances because you are certain that either $\text{Chance}(\text{heads}) = \frac{2}{3}$ or $\text{chance}(\text{heads}) = \frac{1}{3}$. But I have a strong belief that the coin is fair, and I will have a strong belief that you will not learn the true chances. *You* cannot be skeptical about whether you will learn the true chances, but *someone else* can be skeptical about whether you will learn the true chances.

Can you do better than this? You can. The reason that you cannot learn that the coin is fair if it is, is that you have closed your mind to the possibility that the true chances might be anything other than $\frac{1}{3}$ or $\frac{2}{3}$. How can you open your mind to other possibilities? You need to have some prior degrees of belief such that they give probability greater than zero to the true chances being in every little interval around every point from zero to one. This is not so restrictive. You could do this and still be quite sure (for example, with probability .999999) that the chances are $\frac{1}{3}$ or $\frac{2}{3}$, with just a little uncertainty spread over the other possibilities. Then both you and I would be sure that you would eventually learn the true bias of the coin, no matter what it is. All that was required was a little open-mindedness.

Have we solved the problem of induction? Not quite. We assumed some structure in order to get our nice results. We assumed that the true chances made the coin flips independent, and that the chance of heads did not change from flip to flip. If this is the true chance structure and what is unknown is just the bias of the coin then we are fine. But we had to assume something to get some positive results.

We could be more open-minded, starting with some degree of belief in the foregoing coin flipping structure and some degree of belief in alternative structures; for example, a structure where the outcome of a flip influences the chances on the next flip. We could have open-minded degrees of belief that assure eventually learning the true chances if any of these structures is the true chance structure, but with no assumptions at all, we get no results. Probability theory is not magic, and in its strongest pure form, the skepticism of David Hume is unanswerable. Still, I hope that the small sample served up in this chapter convinces the reader that probabilistic analysis has a lot to offer in the way of insight into the problems of scientific induction.

Answers to Selected Exercises

CHAPTER I

Section I.2

2. "Dennis" is the referring expression. "Is tall" is the characterizing expression. It is a property expression.

4. "Arizona," "New Mexico," and "California" are all referring expressions. ". . . is between . . . and . . ." is the characterizing expression. It is a relational expression used to relate three things.

Section I.3

1.b. Mutually Exclusive

	p	q	$\sim p$	$\sim q$	$\sim p \vee \sim q$	$p \& q$
Case 1:	T	T	F	F	F	T
Case 2:	T	F	F	T	T	F
Case 3:	F	T	T	F	T	F
Case 4:	F	F	T	T	T	F

1.d. Mutually Exclusive

	p	q	$\sim p$	$\sim q$	$\sim p \vee q$	$p \& \sim q$
Case 1:	T	T	F	F	T	F
Case 2:	T	F	F	T	F	T
Case 3:	F	T	T	F	T	F
Case 4:	F	F	T	T	T	F

1.f. Logically Equivalent

	p	q	$\sim p$	$\sim q$	$\sim p \vee q$	$\sim(\sim p \vee q)$	$p \& \sim q$
Case 1:	T	T	F	F	T	F	F
Case 2:	T	F	F	T	F	T	T
Case 3:	F	T	T	F	T	F	F
Case 4:	F	F	T	T	T	F	F

2.b. Contingent

	p	q	p	$pvqvr$
Case 1:	T	T	T	T
Case 2:	T	T	F	T
Case 3:	T	F	T	T
Case 4:	T	F	F	T
Case 5:	F	T	T	T
Case 6:	F	T	F	T
Case 7:	F	F	T	T
Case 8:	F	F	F	F

2.d. Tautology

	p	q	$\sim p$	$\sim q$	$pv\sim q$	$\sim(pv\sim q)$	$(pv\sim q)v\sim(pv\sim q)$
Case 1:	T	T	F	F	T	F	T
Case 2:	T	F	F	T	T	F	T
Case 3:	F	T	T	F	F	T	T
Case 4:	F	F	T	T	T	F	T

2.f. Tautology

	p	$\sim p$	$pv\sim p$	$\sim(pv\sim p)$	$\sim\sim(pv\sim p)$
Case 1:	T	F	T	F	T
Case 2:	F	T	T	F	T

Section I.4

1.b. Logically Equivalent

	F	G	$\sim F$	$\sim G$	$F\&G$	$\sim(F\&G)$	$\sim Fv\sim G$
Case 1:	P	P	A	A	P	A	A
Case 2:	P	A	A	P	A	P	P
Case 3:	A	P	P	A	A	P	P
Case 4:	A	A	P	P	A	P	P

1.d. Neither

	F	G	$F\&G$	$\sim(F\&G)$	$Fv\sim(F\&G)$	$\sim(F\&G)\&F$
Case 1:	P	P	P	A	P	A
Case 2:	P	A	A	P	P	P
Case 3:	A	P	A	P	P	A
Case 4:	A	A	A	P	P	A

1.f. Mutually Exclusive

	F	G	H	$F \vee G \vee H$	$\sim(F \vee G \vee H)$
Case 1:	P	P	P	P	A
Case 2:	P	P	A	P	A
Case 3:	P	A	P	P	A
Case 4:	P	A	A	P	A
Case 5:	A	P	P	P	A
Case 6:	A	P	A	P	A
Case 7:	A	A	P	P	A
Case 8:	A	A	A	A	P

2.b. Null

	F	$F \vee F$	$\sim(F \vee F)$	$(F \vee F) \& \sim(F \vee F)$
Case 1:	P	P	A	A
Case 2:	A	A	P	A

2.d. Contingent

	F	G	$\sim F$	$\sim G$	$F \vee \sim G$	$G \vee \sim F$	$(F \vee \sim G) \& (G \vee \sim F)$
Case 1:	P	P	A	A	P	P	P
Case 2:	P	A	A	P	P	A	A
Case 3:	A	P	P	A	A	P	A
Case 4:	A	A	P	P	P	P	P

2.f. Universal

	F	G	$\sim F$	$\sim G$	$F \& \sim G$	$\sim(F \& \sim G)$	$G \vee \sim F$	$\sim(G \vee \sim F)$	$\sim(F \& \sim G) \vee \sim(G \vee \sim F)$
Case 1:	P	P	A	A	A	P	P	A	P
Case 2:	P	A	A	P	P	A	A	P	P
Case 3:	A	P	P	A	A	P	P	A	P
Case 4:	A	A	P	P	A	P	P	A	P

Section I.5

2. Valid

	<u>p</u>	<u>$p \& p$</u>
Case 1:	T	T
Case 2:	F	F

4. Invalid, as shown by case 2 in which the premise is true and the conclusion false

	<u>p</u>	<u>q</u>	<u>$\sim p$</u>	<u>$p \& q$</u>	<u>$\sim(p \& q)$</u>
Case 1:	T	T	F	T	F
Case 2:	T	F	F	F	T
Case 3:	F	T	T	F	T
Case 4:	F	F	T	F	T

6. Valid. The premises are all true only in case 7, where the conclusion is also true.

	<u>p</u>	<u>q</u>	<u>r</u>	<u>$\sim p$</u>	<u>$\sim q$</u>	<u>$p \vee q \vee r$</u>
Case 1:	T	T	T	F	F	T
Case 2:	T	T	F	F	F	T
Case 3:	T	F	T	F	T	T
Case 4:	T	F	F	F	T	T
Case 5:	F	T	T	T	F	T
Case 6:	F	T	F	T	F	T
Case 7:	F	F	T	T	T	T
Case 8:	F	F	F	T	T	F

CHAPTER II**Section II.2**

- 1.b. statement
- d. statement
- f. not a statement
- h. not a statement
- j. statement

- 2.b. not an argument
d. argument

The specific gravity of water is less than that of iron.

Iron will not float when put in water.

Section II.4

2. Inductively strong
4. Neither

CHAPTER IV

Section IV.2

2. *Projectible*: Sugar dissolves in hot coffee. Conservation of momentum.

Unprojectible: Smoking hasn't killed me. After 1980 the stock market goes up every year.

Section IV.3

1. X is grue at t if X is blue or green but not bleen at t . [(Blue \vee Green) $\& \sim$ Bleen]

Section IV.4

X is an insect and X is green

or

X is a ball of wax and X is yellow

or

X is a feather and X is purple

or

X is some other type of thing and X is blue

CHAPTER V

Section V.2

In solving these problems use the definition of necessary and sufficient conditions and the logic of simple and complex properties developed in Chapter I. You can use presence tables to verify the logic steps.

2. If $\sim D$ is a necessary condition for $\sim C$, then
 whenever $\sim C$ is present $\sim D$ is present (definition)
 whenever D is present C is present (logic)
 C is a necessary condition for D (definition)

Presence Table Supporting the Logic:

	C	D	$\sim C$	$\sim D$
Case 1:	P	P	A	A
Case 2:	P	A	A	P
Case 3:	A	P	P	A
Case 4:	A	A	P	P

To say that “whenever $\sim C$ is present $\sim D$ is present” is to say that case 3 doesn’t happen. But if case 3 is ruled out, whenever D is present (case 1) C is present.

Example:

If *motor not running* is a necessary condition for *no fuel*, then *fuel* is a necessary condition for *motor running*.

4. If $\sim C$ is a sufficient condition for $\sim D$, then
 whenever $\sim C$ is present $\sim D$ is present (definition)
 whenever D is present C is present (logic)
 C is a necessary condition for D (definition)
6. If $A \& B$ is a necessary condition for E , then
 whenever E is present $A \& B$ is present (definition)
 whenever E is present A is present (logic)
 A is a necessary condition for E (definition)
 whenever E is present B is present (logic)
 B is a necessary condition for E (definition)
8. If $A \vee B$ is a sufficient condition for E , then
 whenever $A \vee B$ is present E is present (definition)
 whenever A is present E is present (logic)
 whenever B is present E is present (logic)
 A is a sufficient condition for E (definition)
 B is a sufficient condition for E (definition)

10. If A is a sufficient condition for E , then
 whenever A is present E is present (definition)
 whenever $A \& F$ is present E is present (logic)
 $A \& F$ is a sufficient condition for E (definition)

Section V.4

2.

	A	B	C	$\sim A$	$\sim B$	$\sim C$	$A \vee C$	$\sim B \vee C$
Occ.1:	P	P	P	A	A	A	P	P
Occ.2:	P	P	A	A	A	P	P	A
Occ.3:	P	A	P	A	P	A	P	P
Occ.4:	P	P	A	A	A	P	P	A
Occ.5:	A	A	P	P	P	A	P	P
Occ.6:	A	A	P	P	P	A	P	P

All but $A \vee C$ are eliminated because they are absent in some occurrence where E is present.

Section V.5

2. $A, B, C, A \& C, B \& C, A \& B$ are eliminated. The others are not.
 4. $B \& C$ is not eliminated.

Section V.6

2.

A	B	C	D	$\sim A$	$\sim B$	$\sim C$	$\sim D$	E
A	A	A	P	P	P	P	A	A

would eliminate all the candidates except A .

4.

A	B	C	D	$\sim A$	$\sim B$	$\sim C$	$\sim D$	E	
P	A	A	A	A	P	P	P	A	(eliminates A)
A	A	P	P	P	P	A	A	A	(eliminates $\sim B$)
A	P	A	A	P	A	P	P	A	(eliminates $\sim C$)

6. The list of possible conditioning properties did not contain a sufficient condition for E present in occurrence *.

8. Either your answers to 2 and 4 taken together eliminate all of the candidates or they eliminate all but one. If they eliminate all, you could conclude that the list of possible conditioning properties did not contain a sufficient condition for E that was present in occurrence *. If they eliminate all but one, then that one is a sufficient condition for E provided that the initial list contained a sufficient condition for E that was present in occurrence *.

Section V.7

2. It is B . The other properties present in occurrence 1, A , $\sim C$, $\sim D$, are all present in occurrence 2 in which the conditioned property, E , is absent, so they cannot be sufficient conditions for E . This is the joint method of agreement and difference (taking occurrence 1 as occurrence *).

4. B is both necessary and sufficient. Occurrences 1 and 3 eliminate all but B as a necessary condition for E . This is the method of agreement. Occurrences 2 and 4 eliminate all but B as a sufficient condition for E . This is the inverse method of agreement. Altogether, it is the double method of agreement.

CHAPTER 6

Section VI.2

2. $q \& q$ is logically equivalent to q , as can be shown by a truth table, so by rule 3, $\Pr(q \& q) = \Pr(q) = \frac{1}{4}$.

4. $\sim(q \& \sim q)$ is a tautology, as can be shown by a truth table, so by rule 1, $\Pr(\sim(q \& \sim q)) = 1$.

q	$\sim q$	$q \& \sim q$	$\sim(q \& \sim q)$
T	F	F	T
F	T	F	T

6. $p \vee \sim p$ is a tautology, $\sim(p \vee \sim p)$ is a contradiction, $\sim\sim(p \vee \sim p)$ is a tautology, so by rule 1, $\Pr(\sim\sim(p \vee \sim p)) = 1$.

8. $q \& (p \vee \sim p)$ is logically equivalent to q , as can be shown by a truth table:

p	q	$\sim p$	$p \vee \sim p$	$q \& (p \vee \sim p)$
T	T	F	T	T
T	F	F	T	F
F	T	T	T	T
F	F	T	T	F

Thus, by rule 3, $\Pr(q \& (p \vee \sim p)) = \Pr(q) = \frac{1}{4}$.

Section VI.3

1.j. Queen and non-spade are not mutually exclusive, so we must use rule 6:

$$\Pr(Q \vee \sim S) = \Pr(Q) + \Pr(\sim S) - \Pr(Q \& \sim S)$$

$$\text{By rule 5, } \Pr(\sim S) = 1 - \Pr(S) = 1 - \frac{13}{52} = \frac{39}{52}.$$

There are three queens that are not the queen of spades, so $\Pr(Q \& \sim S) = \frac{3}{52}$.

$$\Pr(Q \vee \sim S) = \frac{4}{52} + \frac{39}{52} - \frac{3}{52} = \frac{40}{52} = \frac{10}{13}.$$

3. We do not know if r, s are mutually exclusive, so we use rule 6:

$$\Pr(r \vee s) = \Pr(r) + \Pr(s) - \Pr(r \& s)$$

$$\frac{3}{4} = \frac{1}{2} + \frac{1}{4} - \Pr(r \& s)$$

$$\text{So } \Pr(r \& s) = 0.$$

Section VI.4

1.b. They are not mutually exclusive, because $\Pr(p \& q) > 0$. If they were mutually exclusive, $p \& q$ would be a contradiction and would have probability = 0.

2.b. Because of independence, $\Pr(6 \text{ on die A} \& 6 \text{ on die B}) = \Pr(6 \text{ on A}) \cdot \Pr(6 \text{ on B}) = \frac{1}{6} \cdot \frac{1}{6} = \frac{1}{36}$ by rule 8.

3.b. Two heads and one tail on three tosses can come about in one of three mutually exclusive ways: HHT , HTH , THH . Each way has probability $\frac{1}{8} = \frac{1}{2} \cdot \frac{1}{2} \cdot \frac{1}{2}$ by rule 8. Then $\Pr(HHT \vee HTH \vee THH) = \frac{1}{8} + \frac{1}{8} + \frac{1}{8} = \frac{3}{8}$ by rule 4.

d. $\Pr(HTH) = \frac{1}{2} \cdot \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{8}$ by rule 8.

f. No tails is the same as all heads. $\Pr(HHH) = \frac{1}{2} \cdot \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{8}$ by rule 8.

4.b. If an ace is drawn on draw 1 and not replaced, there are 51 cards remaining, including 4 tens. So $\Pr(10 \text{ on } 2 \text{ given ace on } 1) = \frac{4}{51}$.

d. $\Pr(10 \text{ on } 1 \text{ and ace on } 2) = \Pr(10 \text{ on } 1) \cdot \Pr(10 \text{ on } 2 \text{ given ace on } 1) = \frac{4}{52} \cdot \frac{4}{51}$.

f. $\Pr(\text{two aces}) = \Pr(\text{ace on } 1) \cdot \Pr(\text{ace on } 2 \text{ given ace on } 1) = \frac{4}{52} \cdot \frac{3}{51}$.

$$\begin{aligned}
 5. \Pr(\text{pass}) &= \Pr(\text{pass \& study}) \vee (\text{pass \& } \sim\text{study}) (/) \\
 &= \Pr(\text{pass \& study}) + \Pr(\text{pass \& } \sim\text{study}) \\
 &= \Pr(\text{study}) \cdot \Pr(\text{pass given study}) \\
 &\quad + \Pr(\sim\text{study}) \cdot \Pr(\text{pass given } \sim\text{study}) \\
 &= \Pr(\text{study}) \cdot \Pr(\text{pass given study}) \\
 &\quad + (1 - \Pr(\text{study})) \cdot \Pr(\text{pass given } \sim\text{study}) \\
 &= \frac{4}{5} \cdot \frac{3}{5} + \frac{1}{5} \cdot \frac{1}{10} \\
 &= \frac{1}{2}
 \end{aligned}$$

Section VI.5

<i>Outcome</i>	<i>Probability</i>	<i>NetGain</i>	<i>Probability · Gain</i>
<i>H</i>	$\frac{1}{2}$	\$1	\$.50
<i>TH</i>	$\frac{1}{4}$	\$1	\$.25
<i>TTH</i>	$\frac{1}{8}$	\$1	\$.125
<i>TTT</i>	$\frac{1}{8}$	\$-7	\$-.875
		Expected Value =	\$0

4. No. Try it!

Section VI.6

2.b. $\Pr(\text{second ball from urn 3}) = \Pr(\text{red from urn 1}) = \frac{6}{10}$.

d. $\Pr(\text{black given urn 3}) = \frac{5}{10}$.

f. $\Pr(\text{second ball from urn 2 given second ball black})$

$$= \frac{\Pr(\text{urn 2}) \Pr(\text{black given urn 2})}{\Pr(\text{black})} = \frac{(\frac{4}{10}) \cdot (\frac{1}{10})}{(\frac{34}{100})} = \frac{2}{17}$$

$$\begin{aligned} & \text{h. Pr(second ball from urn 2 given second ball red)} \\ &= \frac{\text{Pr(urn 2)} \cdot \text{Pr(red given urn 2)}}{\text{Pr(red)}} = \frac{\left(\frac{4}{10}\right) \cdot \left(\frac{9}{10}\right)}{1 - \left(\frac{34}{100}\right)} = \frac{36}{66} = \frac{6}{11} \end{aligned}$$

$$\begin{aligned} & \text{j. Pr(red on draw 1 given second ball black)} \\ &= \text{Pr(second ball from urn 3 given second ball black)} \\ &= \frac{\text{Pr(urn 3)} \cdot \text{Pr(black given urn 3)}}{\text{Pr(black)}} = \frac{\left(\frac{6}{10}\right) \cdot \left(\frac{5}{10}\right)}{\left(\frac{34}{100}\right)} = \frac{15}{17} \end{aligned}$$

4. F = “Fitch will not be executed”

J = “Jones will not be executed”

S = “Smith will not be executed”

G = “Guard tells Fitch that Jones will be executed”

According to the problem, $\text{Pr}(F) = \text{Pr}(J) = \text{Pr}(S) = \frac{1}{3}$. We need to calculate $\text{Pr}(F \text{ given } G)$. Using Bayes’ theorem:

$$\text{Pr}(F \text{ given } G) = \frac{\text{Pr}(F) \cdot \text{Pr}(G \text{ given } F)}{\text{Pr}(F) \cdot \text{Pr}(G \text{ given } F) + \text{Pr}(J) \cdot \text{Pr}(G \text{ given } J) + \text{Pr}(S) \cdot \text{Pr}(G \text{ given } S)}$$

The guard is said to be truthful, so $\text{Pr}(G \text{ given } J) = 0$. What is $\text{Pr}(G \text{ given } S)$? If Smith will not be executed, then Fitch and Jones will both be executed. But the guard cannot tell Fitch that he will be executed, so in this case he must tell him that Jones will be executed. So $\text{Pr}(G \text{ given } S) = 1$. So far:

$$\text{Pr}(F \text{ given } G) = \frac{\left(\frac{1}{3}\right) \cdot \text{Pr}(G \text{ given } F)}{\left(\frac{1}{3}\right) \cdot \text{Pr}(G \text{ given } F) + \left(\frac{1}{3}\right)}$$


Everything turns on $\text{Pr}(G \text{ given } F)$. If Fitch will not be executed, Jones and Smith will be. Will the guard, in this case, say “Jones” or Smith? The problem gives no reason why he should prefer one rather than another, so he might flip a fair coin, in which case $\text{Pr}(G \text{ given } F) = \frac{1}{2}$. Then $\text{Pr}(F \text{ given } G) = \left(\frac{1}{6}\right) / \left(\frac{3}{6}\right) = \frac{1}{3}$, and Fitch has no better prospects than before.

On the other hand, you might imagine that the guard has special reasons to say “Jones” in this case, so that $\text{Pr}(F \text{ given } G) = 1$. If so,

Fitch has good news, for then: $\Pr(F \text{ given } G) = (\frac{1}{3})/(\frac{2}{3}) = \frac{1}{2}$. But you might just as well imagine that the guard has special reasons to say “Smith” in this case, so that $\Pr(F \text{ given } G) = 0$. If so, Fitch has bad news, for then $\Pr(F \text{ given } G) = 0$!

Section VII.3

- 1.a. $U(\$100) = 1$
- b. $U(\text{nothing}) = 0$
- c. $(\frac{1}{2}) \cdot U(\$100) + (\frac{1}{2}) \cdot U(\text{nothing}) = \frac{1}{2}$.
- d. It is greater than $\frac{1}{2}$, since he is risk averse.
- e.



- 2.b. He believes that each horse in the race has an equal chance of winning.

Section VII.4

- 2.a. The hypothetical gamble “Get Pig and Stewball wins if HH or HT or TH and get no pig and Stewball loses if TT ,” has expected value of $(\frac{3}{4}) \cdot 1 + (\frac{1}{4}) \cdot 0 = \frac{3}{4}$. Farmer Smith is indifferent between this gamble and “Get pig and Molly wins” then his utility for “Get pig and Molly

wins” must also be $\frac{3}{4}$. By similar reasoning, his utility for “No pig and Molly loses” is $\frac{1}{4}$. His utility for the gamble “Get pig and Stewball wins if HH or HT , no pig and Stewball loses if TH or TT ” is $(\frac{1}{2}) \cdot 1 + (\frac{1}{2}) \cdot 0 = \frac{1}{2}$. Since he is indifferent between this gamble and the gamble “Pig if Molly wins, no pig if she loses,” that gamble must also have utility $= \frac{1}{2}$.

So we have:

$$U(\text{Get pig and Stewball wins}) = 1$$

$$U(\text{Get Pig and Molly wins}) = \frac{3}{4}$$

$$U(\text{Get Pig if Molly wins, no pig if she loses}) = \frac{1}{2}$$

$$U(\text{No Pig and Molly loses}) = \frac{1}{4}$$

$$U(\text{No Pig and Stewball loses}) = 0$$

$$\begin{aligned} \text{b. } U(\text{Get Pig if Molly wins, no pig if she loses}) &= \\ & \text{Pr(Molly wins)} \cdot U(\text{Get Pig and Molly wins}) + \\ & (1 - \text{Pr(Molly wins)}) U(\text{No Pig and Molly loses}) \end{aligned}$$

Substituting in numerical values from part (a) we get:

$$\begin{aligned} \frac{1}{2} &= \text{Pr(Molly wins)} \cdot \left(\frac{3}{4}\right) + (1 - \text{Pr(Molly wins)}) \cdot \left(\frac{1}{4}\right) \\ \text{Pr(Molly wins)} &= \frac{1}{2} \end{aligned}$$

Section VII.5

- a. Relative frequency of heads is zero if there are all tails. $\text{Pr}(TTTT) = (\frac{1}{2}) \cdot (\frac{1}{2}) \cdot (\frac{1}{2}) \cdot (\frac{1}{2}) = \frac{1}{16}$.
- c. Relative frequency of heads is $\frac{1}{2}$ if there are two heads and two tails. $\text{Pr}(HHTT \text{ or } HTHT \text{ or } HTTH \text{ or } THHT \text{ or } THTH \text{ or } TTHH) = \frac{6}{16} = \frac{3}{8}$.
- e. $\text{Pr}(HHHH) = \frac{1}{16}$.

Index

A

- Accidental conditions, 101–108
- Agreement
 - direct method of, 74–79
 - double method of, 90–91
 - inverse method of, 79–85
 - joint method of difference and, 91–92
- Analytic sufficient condition, 102
- “And” as logical connective, 2–3
- Argument(s), 12–14
 - conclusion of, 12
 - deductively valid, 17, 18
 - degrees of strength of, 16
 - e-, 31
 - inductive strength of, 19–21
 - measuring, 51–68
 - inductively strong, 17, 19
 - levels of, 36–37
 - logician’s definition of, 13
 - premises of, 12
 - valid, 9–10

B

- Bayes’ theorem, 130–132
 - relative frequency and, 145
- Belief, degree of
 - in frequentism, 150
 - rational, 137–138
- Bet, fair, 128–129

C

- Calculus, probability, 109–136
- Causality
 - necessary and sufficient conditions and, 69–73
 - probability and, 134–135
- Certainty model of knowing, 25–26
- Chance
 - frequentist view of, 148–149
 - personalist view of, 149–150
 - probability and, 146–150
- Characterizing expressions, 1
- Complementary goods, 139
- Complex contingent statement, probability and, 111
- Complex properties, 8–9
- Complex statements, structure of, 2–7
- Conclusion of argument, 12
 - degrees of confidence in, 31
- Conditional probability, 117–121

- Conditioned property, 73
- Conditions
 - accidental, 101–108
 - counterfactual, 107–108
 - lawlike, 101–108
 - necessary. *See* Necessary conditions
 - sufficient. *See* Sufficient conditions
- Confidence in conclusions, degrees of, 31
- Conjunct, 3
- Conjunction, 3–4
 - in inverse method of agreement, 82–83
- Conjunction rule
 - general, 121–122, 124–127
 - special, 122–124
- Connectives, logical, 2–5
- Constant, system-dependent, 99
- Contingent properties, 9
- Contingent statements, 6
- Contradiction, probability of, 138
- Convergence to truth, 156–157
- Corroborated hypothesis, 151
- Counterfactual condition, 107–108
- Counterinductive logic, 41–43

D

- Deduction, hypothesis and, 151–153
- Deductive logic
 - definition of, 21
 - inductive logic versus, 17–23
- Degree of belief
 - in frequentism, 150
 - rational, 137–138
- Dependent variables, 99
- Difference
 - joint method of agreement and, 91–92
 - method of, 85–90
- Disjunction, 4
- Disjunction rule
 - general, 115–116
 - special, 111–113
 - gambling rationale for, 138
- Disjunctive necessary conditions, 76–78
- Disjuncts, 4

E

- E-argument, 31
- Elimination, principles of, 93
 - in direct method of agreement, 76, 77
 - in inverse method of agreement, 83
- Epistemic probability, 23–26, 30

- Evidence
 quantity of, 153–154
 total, 154–156
 variety of, 153–154
- Evidential links, strength of, 15–16
- Exchangeable trials, 150
- Expectation of chance, rational degree of belief and, 147
- Expected value of gamble, 128–130
- Explanation, science and, 104
- Expressions
 characterizing, 1
 property, 1–2, 8–9
 referring, 1
- Extrapolation, 62–63
- F**
- Factual claim, 12
- Fair bet, 128–129
- Feigl, Herbert, on pragmatic justification of induction, 45
- Frequency, relative, 138, 144–146
 chance and, 150
- Frequentists, chance and, 148–149
- Functional relationships, sufficient conditions and, 97–101
- Functions, 99–100
- G**
- Gamble, expected value of, 128–130
- General conjunction rule, 121–122, 124–127
- General disjunction rule, 115–116
- General negation rule, 113
- Goodman, Nelson, “grue-bleen” paradox of, 56–60
 uniformity of nature and, 66–67
- Goods, complementary, 139
- H**
- Hume, David, on traditional problem of induction, 30–35
- Hypothesis, deduction and, 151–153
- I**
- Impossibility
 logical, 18
 physical, 18
- Inclusive sense of “or,” 4
- Inconsistent statements, 5
- Independence, definition of, 120–121
- Independent variables, 99
- Indicator words in argument, 13
- Induction
 inductive justification of, 35–44
- Induction (*Continued*)
 pragmatic justification of, 44–49
 traditional problem of, 30–50
 Hume on, 30–35
- Inductive justification of induction, 35–44
- Inductive logic
 deductive logic versus, 17–23
 definition of, 21–22
 measuring inductive strength of arguments in, 51–68
 Mill’s methods for finding causes and, 94–95
 pragmatic justification of, 44–49
 probability and, 12–29
 problems of, probability and, 26–29
 scientific, 27
 levels of, 35–44
 system of, constructing, 51–68
 obstacles to, 52–56
- Inductive probability, 20
 definition of, 27
- Inductively strong argument, 17, 19
- Interpolation, 62
- Inverse method of agreement, 79–85
- J**
- Joint method of agreement and difference, 91–92
- Justification of inductive logic
 pragmatic, 44–49
 rational, 31–35, 40–41, 46–49
- L**
- Lawlike conditions, 101–108
- Logic, 14–17
 basics of, 1–11
 counterinductive, 41–43
 deductive
 definition of, 21
 versus inductive, 17–23
 definition of, 15
 inductive. *See* Inductive logic
- Logical connectives, 2–5
- Logically equivalent properties, 9
- Logically equivalent statements, 5
- M**
- Many-place predicates, 2
- Mill, John Stuart, methods of, to find causes, 73–97
 application of, 94–97
 direct method of agreement as, 74–79
 inverse method of agreement as, 79–85
 method of difference as, 85–90

- Mutually exclusive properties, 9, 98
 Mutually exclusive statements, 5
- N**
 Nature, uniformity of, Goodman paradox and, 66–67
 Necessary conditions
 causality and, 69–73
 direct method of agreement for finding, 74–79
 double method of agreement for finding, 90–91
 joint method of agreement and difference for finding, 91–92
 Negation rule, general, 113
 “Not” as logical connective, 2
 Null property, 9
- O**
 One-place predicates, 2
 “Or” as logical connective, 3
- P**
 Personalists, chance and, 149–150
 Physical impossibility, 18
 Popper, Sir Karl, on deductive testing of hypotheses, 151–153
 Positively statistically relevant, 134
 Possible conditioning properties, 73–74
 Pragmatic justification of induction, 44–49
 Predicates
 many-place, 2
 one-place, 2
 Prediction, science and, 104
 Premises of argument, 12
 Presence tables, 8–9
 Principle(s)
 of elimination, 93
 in direct method of agreement, 76, 77
 in inverse method of agreement, 83
 of uniformity of nature, 33–35
 Probabilistic hypotheses, 151–152
 Probability(ies)
 causality and, 134–135
 chance and, 138, 146–150
 conditional, 117–121
 epistemic, 23–26, 30
 inductive, 20
 definition of, 27
 inductive logic and, 12–29
 kinds of, 137–150
 problems of inductive logic and, 26–29
 rational degree of belief and, 137–138
 relative frequencies and, 138, 144–146
 scientific inductive logic and, 151–157
- Probability(ies) (*Continued*)
 utility and, 138–141
 Probability calculus, 109–136
 Projectibility, rules of, for scientific inductive logic, 67
 Projection, regularities and, 52–56
 Property(ies)
 conditioned, 73
 mutually exclusive, 98
 possible conditioning, 73–74
 Property expressions, 1–2, 8–9
 contingent, 9
 logically equivalent, 9
 mutually exclusive, 9
 null, 9
 universal, 9
- R**
 Ramsey, F. P., on utility and probability, 141–144
 Rational degree of belief, 137–138
 Rational justification in inductive logic, 31–35, 40–41, 46–49
 Referring expressions, 1
 Regularity(ies)
 Goodman paradox and, 66–67
 projection and, 52–56
 Reichenbach, Hans, on pragmatic justification of induction, 45–46, 47
 Relationships, functional, sufficient conditions and, 97–101
 Relative frequency, 138, 144–146
 chance and, 150
 Relevant, positively statistically, 134
 Relevant variables, isolating, 101
 Rule S, 53–55
 Rules of projectibility for scientific inductive logic, 67
- S**
 Science, universality of, 104
 Scientific inductive logic, 27
 levels of, 35–44
 probability and, 151–157
 system of, constructing, 51–68
 obstacles to, 52–56
 Self-contradiction, 6, 110
 Sentences in argument, 12
 Simple properties, 8
 Special conjunction rule, 122–124
 Special disjunction rule, 111–113
 gambling rationale for, 138
 Statement(s)
 in argument, 12–13

Statement(s) (*Continued*)

- complex
 - contingent, probability and, 111
 - structure of, 2–7
- contingent, 6
- definition of, 1
- epistemic probability and, 23–26
- inconsistent, 5
- logically equivalent, 5
- mutually exclusive, 5
- probability calculus for, 109
- simple, structure of, 1–2
- Statistic, sufficient, relative frequency and, 145
- Statistically relevant, positively, 134
- Subjective utility, 140
- Sufficient conditions
 - analytic, 102
 - causality and, 69–73
 - double method of agreement for finding, 90–91
 - functional relationships and, 97–101
 - grades of, 102–103
 - inverse method of agreement for finding, 79–85
 - joint method of agreement and difference for finding, 91–92
 - method of difference for finding, 85–90
- Sufficient statistic, relative frequency and, 145
- System-dependent constant, 99

T

- Tautology, 6, 110
 - probability of, 138
- Total evidence condition, 31
- Truth, convergence to, 156–157
- Truth tables, 3–6
 - validity and, 10

U

- Uniformity of nature
 - Goodman paradox and, 66–67
 - principle of, 33–35
- Universal property, 9
- Universality of science, 104
- Utility, probability and, 138–141

V

- Validity, 9–11
 - deductive, 17, 18
- Value, expected, of gamble, 128–130
- Variables
 - dependent, 99
 - independent, 99
 - relevant, isolating, 101
- von Neumann-Morgenstern method for measuring utility, 139–140, 141–142

W

- Weak law of large numbers, 145