


Jun Otsuka



# Thinking About Statistics

The Philosophical Foundations



# THINKING ABOUT STATISTICS

Simply stated, this book bridges the gap between statistics and philosophy. It does this by delineating the conceptual cores of various statistical methodologies (Bayesian/frequentist statistics, model selection, machine learning, causal inference, etc.) and drawing out their philosophical implications. Portraying statistical inference as an epistemic endeavor to justify hypotheses about a probabilistic model of a given empirical problem, the book explains the role of ontological, semantic, and epistemological assumptions that make such inductive inference possible. From this perspective, various statistical methodologies are characterized by their epistemological nature: Bayesian statistics by internalist epistemology, classical statistics by externalist epistemology, model selection by pragmatist epistemology, and deep learning by virtue epistemology.

Another highlight of the book is its analysis of the ontological assumptions that underpin statistical reasoning, such as the uniformity of nature, natural kinds, real patterns, possible worlds, causal structures, etc. Moreover, recent developments in deep learning indicate that machines are carving out their own “ontology” (representations) from data, and better understanding this—a key objective of the book—is crucial for improving these machines’ performance and intelligibility.

## Key Features

- Without assuming any prior knowledge of statistics, discusses philosophical aspects of traditional as well as cutting-edge statistical methodologies.
- Draws parallels between various methods of statistics and philosophical epistemology, revealing previously ignored connections between the two disciplines.
- Written for students, researchers, and professionals in a wide range of fields, including philosophy, biology, medicine, statistics and other social sciences, and business.
- Originally published in Japanese with widespread success, has been translated into English by the author.

**Jun Otsuka** is Associate Professor of Philosophy at Kyoto University and a visiting researcher at the RIKEN Center for Advanced Intelligence Project in Saitama, Japan. He is the author of *The Role of Mathematics in Evolutionary Theory* (Cambridge UP, 2019).



**Taylor & Francis**

Taylor & Francis Group

<http://taylorandfrancis.com>

# THINKING ABOUT STATISTICS

The Philosophical Foundations

*Jun Otsuka*

Designed cover image: © Jorg Greuel/Getty Images

First published in English 2023

by Routledge

605 Third Avenue, New York, NY 10158

and by Routledge

4 Park Square, Milton Park, Abingdon, Oxon, OX14 4RN

*Routledge is an imprint of the Taylor & Francis Group, an informa business*

© 2023 Jun Otsuka

The right of Jun Otsuka to be identified as author of this work has been asserted in accordance with sections 77 and 78 of the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this book may be reprinted or reproduced or utilised in any form or by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying and recording, or in any information storage or retrieval system, without permission in writing from the publishers.

*Trademark notice:* Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Originally published in Japanese by the University of Nagoya Press, Japan.

ISBN: 978-1-032-33310-6 (hbk)

ISBN: 978-1-032-32610-8 (pbk)

ISBN: 978-1-003-31906-1 (ebk)

DOI: 10.4324/9781003319061

Typeset in Bembo

by Apex CoVantage, LLC

*For my parents, Yuzuru and Kuniko Otsuka*



**Taylor & Francis**

Taylor & Francis Group

<http://taylorandfrancis.com>

# CONTENTS

<i>Preface to the English Edition</i>	x
Introduction	1
What Is This Book About?	1
The Structure of the Book	4
1 The Paradigm of Modern Statistics	10
1.1 Descriptive Statistics	10
1.1.1 Sample Statistics	11
1.1.2 Descriptive Statistics as “Economy of Thought”	13
1.1.3 Empiricism, Positivism, and the Problem of Induction	16
1.2 Inferential Statistics	17
1.2.1 Probability Models	18
1.2.2 Random Variables and Probability Distributions	21
1.2.3 Statistical Models	27
1.2.4 The Worldview of Inferential Statistics and “Probabilistic Kinds”	34
Further Reading	38
2 Bayesian Statistics	41
2.1 The Semantics of Bayesian Statistics	41
2.2 Bayesian Inference	46
2.2.1 Confirmation and Disconfirmation of Hypotheses	47
2.2.2 Infinite Hypotheses	48
2.2.3 Predictions	50



2.3	Philosophy of Bayesian Statistics	51
2.3.1	Bayesian Statistics as Inductive Logic	51
2.3.2	Bayesian Statistics as Internalist Epistemology	53
2.3.3	Problems with Internalist Epistemology	57
2.3.4	Summary: Epistemological Implications of Bayesian Statistics	70
	Further Reading	72
3	Classical Statistics	74
3.1	Frequentist Semantics	75
3.2	Theories of Testing	79
3.2.1	Falsification of Stochastic Hypotheses	79
3.2.2	The Logic of Statistical Testing	80
3.2.3	Constructing a Test	81
3.2.4	Sample Size	84
3.3	Philosophy of Classical Statistics	85
3.3.1	Testing as Inductive Behavior	85
3.3.2	Classical Statistics as Externalist Epistemology	87
3.3.3	Epistemic Problems of Frequentism	94
3.3.4	Summary: Beyond the Bayesian vs. Frequentist War	104
	Further Reading	106
4	Model Selection and Machine Learning	109
4.1	The Maximum Likelihood Method and Model Fitting	109
4.2	Model Selection	112
4.2.1	Regression Models and the Motivation for Model Selection	112
4.2.2	A Model's Likelihood and Overfitting	114
4.2.3	Akaike Information Criterion	115
4.2.4	Philosophical Implications of AIC	118
4.3	Deep Learning	124
4.3.1	The Structure of Deep Neural Networks	124
4.3.2	Training Neural Networks	126
4.4	Philosophical Implications of Deep Learning	129
4.4.1	Statistics as Pragmatist Epistemology	129
4.4.2	The Epistemic Virtue of Machines	131
4.4.3	Philosophical Implications of Deep Learning	136
	Further Reading	142
5	Causal Inference	144
5.1	The Regularity Theory and Regression Analysis	145
5.2	The Counterfactual Approach	149
5.2.1	The Semantics of the Counterfactual Theory	149
5.2.2	The Epistemology of Counterfactual Causation	151

5.3	Structural Causal Models	158
5.3.1	Causal Graphs	159
5.3.2	Interventions and Back-Door Criteria	162
5.3.3	Causal Discovery	165
5.4	Philosophical Implications of Statistical Causal Inference	167
	Further Reading	171
6	The Ontology, Semantics, and Epistemology of Statistics	172
6.1	The Ontology of Statistics	172
6.2	The Semantics of Statistics	176
6.3	The Epistemology of Statistics	179
6.4	In Lieu of a Conclusion	181
	<i>Bibliography</i>	183
	<i>Index</i>	189

# PREFACE TO THE ENGLISH EDITION

This book is the English edition of a book originally published in Japanese under the title *Tokeigaku wo Tetsugaku Suru (Philosophizing About Statistics)*, by the University of Nagoya Press. Instead of composing a word-to-word translation, I took this occasion to revise the whole book by incorporating feedback to the original edition and adding some new paragraphs, so it might be more appropriate to call it a rewrite rather than a translation. I also replaced some references and book guides at the end of the chapters with those more accessible to English readers.

Translating your own writing into a language of which you don't have a perfect command is a painful experience. It was made possible only by close collaboration with Jimmy Aames, who went through every sentence and checked not only my English but also the content. Needless to say, however, I am responsible for any errors that may remain. I also owe a great debt to Yukito Iba, Yoichi Matsuzaka, Yusaku Ohkubo, Yusuke Ono, Kentaro Shimatani, Shohei Shimizu, and Takeshi Tejima for their comments on the original Japanese manuscript and book, and Donald Gillies, Clark Glymour, Samuel Mortimer, and an anonymous reviewer for Routledge for their feedback on the English manuscript, all of which led to many improvements.

This book, like all of my other works, was made possible by the support of my mentors, teachers, and friends, including but not limited to: Steen Andersson, Yasuo Deguchi, Naoya Fujikawa, Takehiko Hayashi, Chunfeng Huang, Kunitake Ito, Manabu Kuroki, Lisa Lloyd, Guilherme Rocha, Robert Rose, and Tomohiro Shinozaki. I am also grateful to Kenji Kodate and his colleagues at the University of Nagoya Press, and Andrew Beck and his colleagues at Routledge, for turning the manuscript into books.

Finally and most of all, I am grateful to my family, Akiko and Midori Otsuka, for their moral support during my writing effectively two books in a row in the midst of the global upheaval caused by the COVID-19 pandemic.

# INTRODUCTION

## What Is This Book About?

This book explores the intersection between statistics and philosophy, with the aim of *introducing philosophy to data scientists and data science to philosophers*. By “data science” I am not referring to specific disciplines such as statistics or machine learning research; rather, I am using the term to encompass all scientific as well as practical activities that rely on quantitative data to make inferences and judgments. But why would such a practical science have anything to do with philosophy, often caricatured as empty armchair speculation? Statistics is usually regarded as a rigid system of inferences based on rigorous mathematics, with no room for vague and imprecise philosophical ideologies. A philosophically minded person, on the other hand, might dismiss statistics as merely a practical tool that is utterly useless in tackling deep and ineffable philosophical mysteries.

The primary aim of this book is to dispel these kinds of misconceptions. Statistics today enjoys a privileged role as *the* method of deriving scientific conclusions from observed data. For better or worse, in most popular and scientific articles, “scientifically proven” is taken to be synonymous with “approved by a proper statistical procedure.” But on what theoretical ground is statistics able to play, or at least expected to play, such a privileged role? The justification of course draws its force from sophisticated mathematical machinery, but how is such a mathematical framework able to justify scientific—that is, empirical—knowledge in the first place? This is a philosophical question *par excellence*, and various statistical methods, implicitly or explicitly, have some philosophical intuitions at their root. These philosophical intuitions are seldom featured in common statistics textbooks, partly because they do not provide any extra tools

## 2 Introduction

that readers could use to analyze data they collect for their theses or research projects. However, understanding the philosophical intuitions that lie behind the various statistical methods, such as Bayesian statistics and hypothesis testing, will help one get a grip on their inferential characteristics and make sense of the conclusions obtained from these methods, and thereby become more conscious and responsible about what one is really doing with statistics. Moreover, statistics is by no means a monolith: it comprises a variety of methods and theories, from classical frequentist and Bayesian statistics to the rapidly developing fields of machine learning research, information theory, and causal inference. It goes without saying that the proper application of these techniques demands a firm grasp of their mathematical foundations. At the same time, however, they also involve philosophical intuitions that cannot be reduced to mathematical proofs. These intuitions prescribe, often implicitly, how the world under investigation is structured and how one can make inferences about this world. Or, to use the language of the present book, each statistical method embodies a distinct approach to inductive inference, based on its own characteristic ontology and epistemology. Understanding these ideological backgrounds proves essential in the choice of an appropriate method *vis-à-vis* a given problem and for the correct interpretation of its results, i.e., in making sound inferences rather than falling back on the routine application of ready-made statistical packages. This is why I believe philosophical thinking, despite its apparent irrelevance, can be useful for data analysis.

But, then, what is the point for a *philosopher* to learn statistics? The standard philosophy curriculum in Japanese and American universities is mostly logic-oriented and does not include much training in statistics, with the possible exception of some basic probability calculus under the name of “inductive logic.” Partly because of this, statistics is not in most philosophers’ basic toolbox. I find this very unfortunate, because statistics is like an ore vein that is rich in fascinating conceptual problems of all kinds. One of the central problems of philosophy from the era of Socrates is: how can we acquire *episteme*, or true knowledge? This question has shaped the long tradition of epistemology that runs through the modern philosophers Descartes, Hume, and Kant, leading up to today’s analytic philosophy. In the course of its history, this question has become entwined with various ontological and/or metaphysical issues such as the assumption of the uniformity of nature, the problem of causality, natural kinds, and possible worlds, to name just a few. As the present book aims to show, statistics is the modern scientific variant of philosophical epistemology that comprises all these themes. That is, statistics is a scientific epistemology that rests upon certain ontological assumptions. Therefore, no one working on epistemological problems today can afford to ignore the impressive development and success of statistics in the past century. Indeed, as we will see, statistics and contemporary epistemology share not only common objectives and interests; there is also

a remarkable parallelism in their methodologies. Attending to this parallelism will provide a fruitful perspective for tackling various issues in epistemology and philosophy of science.

Given what has been said thus far, a reader might expect that this book is intended as an introduction to the philosophy of statistics in general. It is not, for two reasons. First, this book does not pretend to introduce the reader to the field of the *philosophy of statistics*, a well-established branch of contemporary philosophy with a wealth of discussions concerning the theoretical ground of inductive inference, interpretations of probability, the everlasting battle between Bayesian and frequentist statistics, and so forth (Bandyopadhyay and Forster 2010). While these are all important and interesting topics, going through them would make a huge volume, and in any case far exceeds the author's capability. Moreover, as these discussions often tend to be highly technical and assume familiarity with both philosophy and statistics, non-specialists may find it difficult to follow or keep motivated. Some of these topics are of course covered in this book, and in the case of others I will point to the relevant literature. But instead of trying to cover all these traditional topics, this book cuts into philosophical issues in statistics with my own approach, which I will explain in a moment. Thus readers should keep in mind that this book is not intended as a textbook-style exposition of the standard views in the philosophy of statistics.

The second reason why this book is not entitled *An Introduction to the Philosophy of Statistics* is that it does not aim to be an "introduction" in the usual sense of the term. The Japanese word for introduction literally means "to enter the gate," with the implication that a reader visits a particular topic and stays there as a guest for a while (imagine visiting a temple) in order to appreciate, experience, and learn its internal atmosphere and architectural art. This book, however, is not a well-mannered tour guide who quietly stays at one topic, either statistics or philosophy. It is indeed a restless traveler, entering the gate of statistics, quickly leaving and entering philosophy from a different gate, only to be found in the living room of statistics at the next moment. At any rate, the goal of this book is not to make the reader proficient in particular statistical tools or philosophical ideas. This does not mean that it presupposes prior familiarity with statistics or philosophy: on the contrary, this book is designed to be as self-contained as possible, providing plain explanations for every statistical technique and philosophical concept at their first appearance (so experts may well want to skip these introductory parts). The aim of these explanations, however, is not to make the reader a master of the techniques and ideas themselves; rather, they are meant to elucidate the conceptual relationships among these techniques and ideas. Throughout this book we will ask questions like: how is a particular statistical issue discussed in the context of philosophy? How does a particular philosophical concept contribute to our understanding of statistical thinking? Through such questions, this book aims to bridge statistics and philosophy and reveal the conceptual parallelism between them. Because

## 4 Introduction

of this interdisciplinary character, this book is not entitled “Introduction” and is not intended to be read as such. That is, this book does not pretend to train the reader to become a data scientist or philosopher. Rather, this is a book for border-crossers: it tempts the data analyst to become a little bit of a philosopher, and the philosophy lover to become a little bit of a data scientist.

### The Structure of the Book

What kind of topics, then, are covered in this book? This book may be likened to a fabric, woven with philosophy as its warp and statistics as its weft. The philosophy warp consists of three threads: ontology, semantics, and epistemology. *Ontology* is the branch of philosophy that studies the real nature of things existing in the world. Notable historical examples include the Aristotelian theory of the four elements, according to which all subcelestial substances are composed from the basic elements of fire, air, water, and earth; and the mechanical philosophy of the 17th century, which aimed to reduce all physical matter to microscopic particles. But ontology is not monopolized by philosophers. Indeed, every scientific theory makes its own ontological assumptions as to what kinds of things constitute the world that it aims to investigate. The world of classical mechanics, for example, is populated by massive bodies, while a chemist or biologist would claim that atoms and molecules, or genes and cells, also exist according to their worldview. We will not be concerned here with issues such as the adequacy of these ontological claims, or which entities are more “fundamental” and which are “derivative.” What I am pointing out is simply the truism that every scientific investigation, insofar as it is an empirical undertaking, must make clear what the study is *about*.

Unlike physics or biology, which have a concrete domain of study, statistics *per se* is not an empirical science and thus may not seem to rely on any explicit assumption about what exists in the world. Nevertheless, it still makes ontological assumptions about the structure of the world in a more abstract way. What are the entities posited by statistics? The first and foremost thing that must exist in statistics is obvious: data. But this is not enough—the true value of statistics, especially its primary component known as inferential statistics, lies in its art of inferring the unobserved from the observed. Such an inference that goes beyond the data at hand is called *induction*. As the 18th-century Scottish philosopher David Hume pointed out, inductive inference relies on what he called the *uniformity of nature* behind the data. Inferential statistics performs predictions and inferences by mathematically modeling this latent uniformity behind the data (Chapter 1). These mathematical models come in various forms, with differing shades of ontological assumptions. Some models assume more “existence” in the world than others, in order to make broader kinds of inferences possible. Although such philosophical assumptions often go unnoticed in statistical practice, they also sometimes rear their head. For instance, questions

such as “In what sense are models selected by AIC considered good?” or “Why do we need to think about ‘possible outcomes’ in causal inference?” are ontological questions *par excellence*. In each section of this book, we will try to reveal the ontological assumptions that underpin a given statistical method, and consider the implications that the method has on our ontological perspective of the world.

Statistics thus mathematically models the world’s structure and expresses it in probabilistic statements. But mathematics and the world are two different things. In order to take such mathematical models as models of empirical phenomena, we must interpret these probabilistic statements in a concrete way. For example, what does it mean to say that the probability of a coin’s landing heads is 0.5? How should we interpret the notorious  $p$ -value? And what kind of state of affair is represented by the statement that a variable  $X$  causes another variable  $Y$ ? *Semantics*, which is the second warp thread of this book, elucidates the meaning of statements and conceptions that we encounter in statistics.

Statistics is distinguished from pure mathematics in that its primary goal is not the investigation of mathematical structure *per se*, but rather the application of its conclusions to the actual world and concrete problems. For this purpose, it is essential to have a firm grasp of what statistical concepts and conclusions stand for, i.e., their semantics. However, just as statistics itself is not a monolith, so the meaning and interpretation of its concepts are not determined uniquely either. In this book we will see the ways various statistical concepts are understood in different schools of statistics, along with the implications that these various interpretations have for actual inferential practices and applications.

The third and last warp thread of this book is *epistemology*, which concerns the art of correctly inferring the entities that are presupposed and interpreted from actual data. As we noted earlier, statistics is regarded as the primary method by which an empirical claim is given scientific approval in today’s society. There is a tacit social understanding that what is “proven” statistically is likely true and can be accepted as a piece of scientific knowledge. What underlies this understanding is our idea that the conclusion of an appropriate statistical method is not a lucky guess or wishful thinking; it is justified in a certain way. But what does it mean for a conclusion to be justified? There has been a long debate over the concept of justification in philosophical epistemology. Similarly, in statistics, justification is understood in different ways depending on the context—what is to be regarded as “(statistically) certain” or counts as statistically confirmed “knowledge” is not the same among, say, Bayesian statistics, classical statistics, and the machine learning literature, and the criteria are not always explicit even within each tradition. This discrepancy stems from their respective philosophical attitudes as to how and why *a priori* mathematical proofs and calculations are able to help us in solving empirical problems like prediction and estimation. This philosophical discordance has led to longstanding conflicts among statistical paradigms, as exemplified by the



notorious battle between Bayesians and frequentists in the 20th century. It is not my intention to fuel this smoldering debate in this book; rather, what I want to emphasize is that this kind of discrepancy between paradigms is rooted in the different ways that they understand the concept of justification. Keeping this in mind is important, not in order to decide on a winner, but in order to fully appreciate their respective frameworks and to reflect on why we are able to acquire empirical knowledge through statistical reasoning in the first place. As will be argued in this book, the underlying epistemology of Bayesian statistics and that of classical testing theory are akin to internalism and externalism in contemporary epistemology, respectively. This parallelism, if it holds, is quite intriguing, given the historical circumstance that statistics and philosophical epistemology developed independently without much interaction, despite having similar aims.

With ontology, semantics, and epistemology as our philosophical warp threads, each chapter of this book will focus on a specific statistical method and analyze its philosophical implications; this will constitute the weft of this book.

Chapter 1 is a preliminary introduction to statistics without tears for those who have no background knowledge of the subject. It reviews the basic distinction between descriptive and inferential statistics and explains the minimal mathematical framework necessary for understanding the remaining chapters, including the notions of sample statistics, probability models, and families of distributions. Furthermore, the chapter introduces the central philosophical ideas that run through this book, namely that this mathematical framework represents an ontology for inductive reasoning, and that each of the major statistical methods provides an epistemological apparatus for inferring the entities thus postulated.

With this basic framework in place, Chapter 2 takes up Bayesian statistics. After a brief review of the standard semantics of Bayesian statistics, namely the subjective interpretation of probability, the chapter introduces Bayes' theorem and some examples of inductive inference based on it. The received view takes Bayesian inference as a process of updating—through probabilistic calculations and in accordance with evidence—an epistemic agent's degree of belief in hypotheses. This idea accords well with internalist epistemology, according to which one's beliefs are to be justified by and only by other beliefs, via appropriate inferential procedures. Based on this observation, it will be pointed out that well-known issues of Bayesian statistics, like the justification of prior probabilities and likelihood, have exact analogues in foundationalist epistemology, and that if such problems are to be avoided, inductive inference cannot be confined to internal calculations of posterior probabilities but must be opened up to holistic, extra-model considerations, through model-checking and the evaluation of predictions.

Chapter 3 turns to so-called classical statistics, and in particular the theory of statistical hypothesis testing. We briefly review the frequentist interpretation of probability, which is the standard semantics of classical statistics, and then we

introduce the basics of testing theory, including its key concepts like significance levels and  $p$ -values, using a simple example. Statistical tests tell us whether or not we should reject a given hypothesis, together with a certain error probability. Contrary to a common misconception, however, they by no means tell us about the truth value or even probability of a hypothesis. How, then, can such test results justify scientific hypotheses? We will seek a clue in externalist epistemology: by appealing to a view known as reliabilism and Nozick's tracking theory, I argue that good tests are reliable epistemic processes, and their conclusions are therefore justified in the externalist sense. The point of this analogy is not simply to draw a connection between statistics and philosophy, but rather to shed light on the well-known issues of testing theory. In particular, through this lens we will see that the misuse of  $p$ -values and the replication crisis, which have been a topic of contention in recent years, can be understood as a problem concerning the reliability of the testing process, and that the related criticism of classical statistics in general stems from a suspicion about its externalist epistemological character.

While the aforementioned chapters deal with classical themes in statistics, the fourth and fifth chapters will focus on more recent topics. The main theme of Chapter 4 is prediction, with an emphasis on the recently developed techniques of model selection and deep learning. Model selection theory provides criteria for choosing the best among multiple models for the purpose of prediction. One of its representative criteria, the Akaike Information Criterion (AIC), shows us that a model that is too complex, even if it allows for a more detailed and accurate description of the world, may fare worse in terms of its predictive ability than a simpler or more coarse-grained model. This result prompts us to reconsider the role of models in scientific inferences, suggesting the pragmatist idea that modeling practices should reflect and depend on the modeler's practical purposes (such as the desired accuracy of predictions) as well as limitations (the size of available data). On the other hand, deep learning techniques allow us to build highly complex models, which are able to solve predictive tasks with big data and massive computational power. The astonishing success of this approach in the past decade has revolutionized scientific practice and our everyday life in many aspects. Despite its success, however, deep learning models differ from traditional statistical models in that much of their theoretical foundations and limitations remain unknown—in this respect they are more like accumulations of engineering recipes developed through trial and error. But in the absence of theoretical proofs, how can we trust the outcomes or justify the conclusions of deep learning models? We will seek a clue to this question in virtue epistemology, and argue that the reliability of a deep learning model can be evaluated in terms of its model-specific epistemological capability, or epistemic virtue. This perspective opens up the possibility of employing philosophical discussions about understanding the epistemic abilities of other people and species for thinking about what “understanding a deep learning model” amounts to.

Chapter 5 changes gears and deals with causal inference. Every student of statistics knows that causality is not probability—but how are they different? In the language of the present book, they correspond to distinct kinds of entities; in other words, probabilistic inference and causal inference are rooted in different ontologies. While predictions are inferences about this actual world, causal inferences are inferences about possible worlds that would or could have been. With this contrast in mind, the chapter introduces two approaches to causal inference: counterfactual models and structural causal models. The former encodes situations in possible worlds using special variables called potential outcomes, and estimates a causal effect as the difference between the actual and possible worlds. The latter represents a causal relationship as a directed graph over variables and studies how the topological relationships among the graph’s nodes determine probability distributions and vice versa. Crucial in both approaches is some assumption or other concerning the relationship between, on the one hand, the data observed in the actual world and, on the other, the possible worlds or causal structures which, by their very nature, can never be observed. The well-known “strongly ignorable treatment assignment” assumption and the “causal Markov condition” are examples of bridges between these distinct ontological levels, without which causal relationships cannot be identified from data. In causal inference, therefore, it is essential to keep in mind the ontological level to which the estimand (the quantity to be estimated) belongs, and what assumptions are at work in the estimation process.

On the basis of these considerations, the sixth and final chapter takes stock of the ontological, semantic, and epistemological aspects of statistics, with a view toward the fruitful and mutually inspiring relationship between statistics and philosophy.

Figure 0.1 depicts the logical dependencies among the chapters. Since philosophical issues tend to relate to one another, the parts of this book are written




FIGURE 0.1 Flowchart of the book

in such a way that they reflect as many of these organic connections as possible. Readers who are interested in only certain portions of the book will find the diagram useful for identifying relevant contexts and subsequent material. At the end of each chapter I have included a short book guide for the interested reader. I stress, however, that the selection is by no means exhaustive or even standard: rather, it is a biased sample taken from a severely limited pool. There are many good textbooks on both statistics and philosophy, so the reader is encouraged to consult works that suit their own needs and tastes.

## References

- Adadi, Amina , and Mohammed Berrada . 2018. "Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI)." *IEEE Access* 6: 52138–52160.
- Akaike, Hirotugu . 1974. "A New Look at the Statistical Model Identification." *IEEE Transactions on Automatic Control* 19 (6): 716–723.
- Anderson, David R. 2008. *Model Based Inference in the Life Sciences: A Primer on Evidence*. Springer.
- Bandyopadhyay, Prasanta S. , and Malcolm Forster , eds. 2010. *Philosophy of Statistics. Handbook of the Philosophy of Science*. Elsevier.
- Battaly, Heather . 2008. "Virtue Epistemology." *Philosophy Compass* 3 (4): 639–663.
- Berger, James O. , and Robert L. Wolpert . 1988. *The Likelihood Principle*. 2nd ed. Institute of Mathematical Statistics.
- Bianchini, Stefano , Moritz Müller , and Pierre Pelletier . 2020. "Deep Learning in Science." arXiv:2009.01575 [cs.CY].
- Bickel, P.J. , E.A. Hammel , and J.W. O'Connell . 1975. "Sex Bias in Graduate Admissions: Data from Berkeley." *Science* 187 (4175): 398–404.
- Bird, Alexander , and Emma Tobin . 2018. "Natural Kinds." In *The Stanford Encyclopedia of Philosophy*, Spring 2018, edited by Edward N. Zalta . Meta-physics Research Lab, Stanford University Press.
- Birnbaum, Allan . 1962. "On the Foundations of Statistical Inference." *Journal of the American Statistical Association* 57 (298): 269.
- Bonjour, Laurence . 1980. "Externalist Theories of Empirical Knowledge." *Midwest Studies in Philosophy* 5 (1): 54–74.
- Box, George E.P. , Alberto Luceño , and Maria del Carmen Paniagua-Quinones . 2009. *Statistical Control by Monitoring and Adjustment*. Wiley.
- Boyd, Richard . 1999. "Homeostasis, Species, and Higher Taxa." In *Species: New Interdisciplinary Essays*, edited by Robert A. Wilson , 141–186. MIT Press.
- Bradley, Darren . 2015. *A Critical Introduction to Formal Epistemology*. Bloomsbury.
- Cartwright, Nancy . 1983. *How the Laws of Physics Lie*. Oxford University Press.
- Cartwright, Nancy . 1999. *The Dappled World*. Cambridge University Press.
- Childers, Timothy . 2013. *Philosophy and Probability*. Oxford University Press.
- Conee, Earl , and Richard, Feldman . 1998. "The Generality Problem for Reliabilism." *Philosophical Studies* 89 (1): 1–29.
- Dennett, Daniel C. 1991. "Real Patterns." *The Journal of Philosophy* 88 (1): 27–51.
- Earman, John . 1992. *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. The MIT Press.
- Edwards, Ward , Harold Lindman , and Leonard J. Savage . 1963. "Bayesian Statistical Inference for Psychological Research." *Psychological Review* 70 (3): 193–242.
- Efron, Bradley , and Trevor Hastie . 2016. *Computer Age Statistical Inference: Algorithms, Evidence, and Data Science*. Cambridge University Press.
- Forster, Malcom , and Elliott Sober . 1994. "How to Tell when Simpler, More Unified, or Less Ad Hoc Theories Will Provide More Accurate Predictions." *The British Journal for the Philosophy of Science* 45 (1): 1–35.
- Fumerton, Richard . 2002. "Theories of Justification." In *The Oxford Handbook of Epistemology*, edited by Paul K. Moser , 204–233. Oxford University Press.
- Galton, Francis . 1886. "Regression Towards Mediocrity in Hereditary Stature." *The Journal of the Anthropological Institute of Great Britain and Ireland* 15: 246–263.
- Gelman, Andrew , John Carlin , Hal Stern , and Donald Rubin . 2004. *Bayesian Data Analysis*. Chapman & Hall/CRC.
- Gelman, Andrew , and Cosma Rohilla Shalizi . 2012. "Philosophy and the Practice of Bayesian Statistics." *British Journal of Mathematical and Statistical Psychology* 66 (1): 8–38.
- Gettier, Edmund L. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23 (6): 121–123.
- Gillard, Jonathan . 2020. *A First Course in Statistical Inference*. Springer.
- Gillies, Donald . 2000. *Philosophical Theories of Probability*. 1st ed. Routledge.

- Gillies, Donald . 2009. "Problem-Solving and the Problem of Induction." In *Rethinking Popper*, edited by Zuzana Parusniková and Robert S. Cohen , 103–115. Springer.
- Glymour, Clark . 1981. *Theory and Evidence*. Chicago University Press.
- Glymour, Clark , and Madelyn R. Glymour . 2014. "Commentary: Race and Sex Are Causes." *Epidemiology* 25 (4): 488–490.
- Goldman, Alvin I. 1975. "Innate Knowledge." In *Innate Ideas*, edited by Stephen P. Stich , 111–120. University of California Press.
- Goldman, Alvin I. 2009. "Internalism, Externalism, and the Architecture of Justification." *The Journal of Philosophy* 106 (6): 309–338.
- Goodfellow, Ian , Yoshua Bengio , and Aaron Courville . 2016. *Deep Learning*. The MIT Press.
- Hacking, Ian . 1980. "The Theory of Probable Inference: Neyman Peirce and Braithwaite." In *Science, Belief and Behavior: Essays in Honour of R. B. Braithwaite*, edited by D. H. Mellor , 141–160.
- Hacking, Ian . 1990. *The Taming of Chance*. Cambridge University Press.
- Hacking, Ian . 2006. *The Emergence of Probability: A Philosophical Study of Early Ideas about Probability, Induction and Statistical Inference*. Cambridge University Press.
- Hacking, Ian . 2016. *Logic of Statistical Inference*. Cambridge University Press.
- Hájek, Alan . 2007. "The Reference Class Problem is Your Problem Too." *Synthese* 156 (3): 563–585.
- Hand, David J. 2008. *Statistics: A Very Short Introduction*. Oxford University Press.
- Hasan, Ali , and Richard Fumerton . 2018. "Foundationalist Theories of Epistemic Justification." In *The Stanford Encyclopedia of Philosophy*, Fall 2018, edited by Edward N. Zalta . Metaphysics Research Lab, Stanford University.
- Hempel, Carl G. , and Paul Oppenheim . 1948. "Studies in the Logic of Explanation." *Philosophy of Science* 15 (2): 135–175.
- Hendricks, Lisa Anne , Zeynep Akata , Marcus Rohrbach , Jeff Donahue , Bernt Schiele , and Trevor Darrell . 2016. "Generating Visual Explanations." *Computer Vision—ECCV 2016*. ECCV 2016. Lecture Notes in Computer Science 9908.
- Hoff, Peter D. 2009. *A First Course in Bayesian Statistical Methods*. Springer.
- Holland, Paul W. 1986. "Statistics and Causal Inference." *Journal of the American Statistical Association* 81 (396): 945–960.
- Howson, Colin , and Peter Urbach . 2006. *Scientific Reasoning: The Bayesian Approach*. 3rd ed. Open Court.
- Hume, David . 1748. *An Enquiry Concerning Human Understanding*.
- Imaizumi, Masaaki , and Kenji Fukumizu . 2019. "Deep Neural Networks Learn Non-Smooth Functions Effectively." In *Proceedings of Machine Learning Research*, edited by Kamalika Chaudhuri and Masashi Sugiyama , 89: 869–878.
- Iten, Raban , Tony Metger , Henrik Wilming , Lidia del Rio , and Renato Renner . 2020. "Discovering Physical Concepts with Neural Networks." *Physical Review Letters* 124 (1): 010508.
- James, William . 1907. *Pragmatism: A New Name for Some Old Ways of Thinking*. Harvard University Press.
- Jaynes, E. T. 1957. "Information Theory and Statistical Mechanics." *Physical Review* 106 (4): 620–630.
- Jeffrey, Richard . 2004. *Subjective Probability: The Real Thing*. Cambridge University Press.
- Kasuya, Eiiti . 2015. "Misuse of Akaike Information Criterion in Ecology—AIC Does Not Choose Correct Models Because AIC Is Not Aimed at the Identification of Correct Models (in Japanese)." *Japanese Journal of Ecology* 65 (2): 179–185.
- Kelleher, John D. 2019. *Deep Learning*. The MIT Press.
- Klein, Peter D. 1999. "Human Knowledge and the Infinite Regress of Reasons." *Philosophical Perspectives* 13: 297–325.
- Konishi, Sadanori , and Genshiro Kitagawa . 2008. *Information Criteria and Statistical Modeling*. Springer Science & Business Media.
- Krantz, David H. , Patrick Suppes , R. Duncan Luce , and A. Tversky . 1971. *Foundations of Measurement (Additive and Polynomial Representations)*. Vol. 1. Academic Press, Inc.

- Krohn, Jon , Grant Beyleveld , and Aglaé Bassens . 2019. *Deep Learning Illustrated: A Visual, Interactive Guide to Artificial Intelligence*. Addison-Wesley Professional.
- Kutach, Douglas . 2014. *Causation*. Polity Press Ltd.
- Leemis, Lawrence M. , and Jacquelyn T. McQueston . 2008. "Univariate Distribution Relationships." *American Statistician* 62 (1): 45–53.
- de Leeuw, Jan . 1992. "Introduction to Akaike (1973) Information Theory and an Extension of the Maximum Likelihood Principle." In *Breakthroughs in Statistics*, edited by Kotz S. and N. L. Johnson , 599–609. Springer Series in Statistics. Springer.
- Lewis, David . 1973. "Causation." *The Journal of Philosophy* 70 (17): 556–567.
- Lewis, David . 1980. "A Subjectivist's Guide to Objective Chance." In *Studies in Inductive Logic and Probability*, edited by Richard C. Jeffrey , II:263–293. University of California Press.
- Lindley, Dennis V. , and Lawrence D. Phillips . 1976. "Inference for a Bernoulli Process (a Bayesian View)." *American Statistical* 30 (3): 112–119.
- Marcellesi, Alexandre . 2013. "Is Race a Cause?" *Philosophy of Science* 80 (5): 650–659.
- Mayo, Deborah G. 1996. *Error and the Growth of Experimental Knowledge*, 493. University of Chicago Press.
- Mayo, Deborah G. 2018. *Statistical Inference as Severe Testing: How to Get Beyond the Statistics Wars*. Cambridge University Press.
- McGrayne, Sharon Bertsch . 2011. *The Theory that Would Not Die: How Bayes' Rule Cracked the Enigma Code, Hunted Down Russian Submarines, & Emerged Triumphant from Two Centuries of Controversy*. Yale University Press.
- Millikan, Ruth Garrett . 1984. *Language, Thought, and Other Biological Categories*. The MIT Press.
- Morgan, Stephen L. , and Christopher Winship . 2007. *Counterfactuals and Causal Inference*. Cambridge University Press.
- Mumford, Stephen , and Rani Lill Anjum . 2014. *Causation: A Very Short Introduction*. Oxford University Press.
- Nagel, Jennifer . 2014. *Knowledge: A Very Short Introduction*. Oxford University Press.
- Narens, Louis . 2007. *Introduction to the Theories of Measurement and Meaningfulness and the Use of Symmetry in Science*. Psychology Press.
- Neyman, Jerzy . 1957. "'Inductive Behavior' as a Basic Concept of Philosophy of Science." *Revue de l'Institut International de Statistique/Review of the International Statistical Institute* 25 (1/3): 7.
- Neyman, Jerzy , and Egon S. Pearson . 1933. "On the Problem of the Most Efficient Tests of Statistical Hypotheses." *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 231 (694–706): 289–337.
- Nozick, Robert . 1981. *Philosophical Explanations*. Harvard University Press.
- Otsuka, Jun . 2019. "Design Problems in Life and AI (in Japanese)." *Journal of the Japan Association for Philosophy of Science* 46 (2): 71–77.
- Pappas, George . 2017. "Internalist vs. Externalist Conceptions of Epistemic Justification." In *The Stanford Encyclopedia of Philosophy*, Fall 2017, edited by Edward N. Zalta . Metaphysics Research Lab, Stanford University.
- Pearl, Judea . 2000. *Causality: Models, Reasoning, and Inference*. Cambridge University Press.
- Pearl, Judea . 2009. *Myth, Confusion, and Science in Causal Analysis*. Technical Report.
- Pearl, Judea , Madelyn Glymour , and Nicholas P. Jewell . 2016. *Causal Inference in Statistics: A Primer*. John Wiley & Sons.
- Pearson, Karl . 1892. *The Grammar of Science*. Adam/Charles Black.
- Peters, Jonas , Dominik Janzing , and Bernhard Schölkopf . 2017. *Elements of Causal Inference: Foundations and Learning Algorithms*. The MIT Press.
- Ponciano, José Miguel , and Mark L. Taper . 2019. "Model Projections in Model Space: A Geometric Interpretation of the AIC Allows Estimating the Distance Between Truth and Approximating Models." *Frontiers in Ecology and Evolution* 7.
- Popper, Karl Raimund . 1959. *The Logic of Scientific Discovery*, 513. Hutchinson & Co.

Porter, Theodore M. 1996. *Trust in Numbers*, 324. Princeton University Press.

Porter, Theodore M. 2001. "Statistical Tales." *American Scientist* 89 (5): 469–470.

Pritchard, Duncan . 2014. *What Is This Thing Called Knowledge?* 3rd ed. Routledge.

Quine, Willard Van Orman . 1951. "Two Dogmas of Empiricism." *The Philosophical Review* 60(1): 20–43.

Quine, Willard Van Orman . 1960. *Word and Object*. The MIT Press.

Quine, Willard Van Orman . 1986. "Reply to Morton White." In *The Philosophy of W. V. Quine*, edited by L. E. Harn and P. A. Schipp . Open Court.

Reichenbach, Hans . 2008. *The Concept of Probability in the Mathematical Representation of Reality*. Open Court.

Ribeiro, Marco Tulio , Sameer Singh , and Carlos Guestrin . 2016. "Why Should I Trust You?: Explaining the Predictions of Any Classifier." *KDD '16: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144.

Rorty, Richard . 1979. *Philosophy and the Mirror of Nature*. Princeton University Press.

Rosenbaum, Paul . 2017. *Observation and Experiment: An Introduction to Causal Inference*. Harvard University Press.

Roush, Sherrilyn . 2005. *Tracking Truth: Knowledge, Evidence, and Science*. Oxford University Press.

Rowbottom, Darrell P. 2015. *Probability*. Polity Press.

Rubin, Donald B. 1974. "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies." *Journal of Educational Psychology* 66 (5): 688–701.

Russell, Bertrand . 1948. *Human Knowledge: Its Scope and Limits*. Allen & Unwin.

Salsburg, David . 2001. *The Lady Tasting Tea: How Statistics Revolutionized Science in the Twentieth Century*. Macmillan.

Sellars, Wilfrid . 1997. *Empiricism and the Philosophy of Mind*. Harvard University Press.

Shibamura, Ryo . 2004. *R.A. Fisher's Statistical Theory: The Formation and Social Context of Inferential Statistics (in Japanese)*. Kyushu University Press.

Simpson, Edward H. 1951. "The Interpretation of Interaction in Contingency Tables." *Journal of the Royal Statistical Society: Series B Statistical Methodology* 13 (2): 238–241.

Sober, Elliott . 1980. "Evolution, Population Thinking, and Essentialism." *Philosophy of Science* 47: 350–383.

Sober, Elliott . 2008. *Evidence and Evolution*. Cambridge University Press.

Sober, Elliott . 2015. *Ockham's Razors*, 322. Cambridge University Press.

Sosa, Ernest . 1985. "Knowledge and Intellectual Virtue." *Monist* 68 (2): 226–245.

Sosa, Ernest . 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Vol. 1. Oxford University Press.

Sosa, Ernest . 2009. *Reflective Knowledge: Apt Belief and Reflective Knowledge*. Vol. 2. Oxford University Press.

Spirtes, Peter , Clark Glymour , and Richard Scheines . 1993. *Causation, Prediction, and Search*. The MIT Press.

Staley, Kent , and Aaron Cobb . 2011. "Internalist and Externalist Aspects of Justification in Scientific Inquiry." *Synthese* 182 (3): 475–492.

Steup, Matthias , and Ram Neta . 2020. "Epistemology." In *The Stanford Encyclopedia of Philosophy*, Summer 2020, edited by Edward N. Zalta . Meta-physics Research Lab, Stanford University.

Stich, Stephen P. 1990. *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. The MIT Press.

Stigler, Stephen M. 1986. *The History of Statistics: The Measurement of Uncertainty Before 1900*. Harvard University Press.

Stone, Mervyn . 1977. "An Asymptotic Equivalence of Choice of Model by Cross-validation and Akaike's Criterion." *Journal of the Royal Statistical Society* 39 (1): 44–47.

Suárez, Mauricio . 2020. "Philosophy of Probability and Statistical Modelling." In *Elements in the Philosophy of Science*. Cambridge University Press.

Szegedy, Christian , Wojciech Zaremba , Ilya Sutskever , Joan Bruna , Dumitru Erhan , Ian Goodfellow , and Rob Fergus . 2014. "Intriguing Properties of Neural Networks." *arXiv*,1312.6199.



Takeuchi, Kei . 1976. "Distribution of Information Statistics and Criteria for Adequacy of Models (in Japanese)." *Suri-Kagaku (Mathematical Science)* 153: 12–18.

Tanaka, Akinori , Akio Tomiya , and Koji Hashimoto . 2021. *Deep Learning and Physics*. Springer.

VanderWeele, Tyler J. , and Whitney R. Robinson . 2014. "On the Causal Interpretation of Race in Regressions Adjusting for Confounding and Mediating Variables." *Epidemiology* 25 (4): 473–484.

van Fraassen, Bas C. 1977. "Relative Frequencies." *Synthese* 34: 133–166.

van Fraassen, Bas C. 1980. *The Scientific Image*. Oxford University Press.

Vaughan, Simon . 2013. *Scientific Inference: Learning from Data*. Cambridge University Press.

von Mises, Richard . 1928. *Probability, Statistics, and Truth*. 2nd revised English ed. George Allen/Unwin Ltd.

Wasserman, Larry . 2004. *All of Statistics: A Concise Course in Statistical Inference*. Springer.

Wasserstein, Ronald L. , and Nicole A. Lazar . 2016. "The ASA's Statement on p-Values: Context, Process, and Purpose." *American Statistical* 70 (2): 129–133.

Williamson, Jon . 2010. *In Defence of Objective Bayesianism*. Oxford University Press.

Wimsatt, William C. 2007. *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Harvard University Press.

Woodward, James . 2003. *Making Things Happen*. Oxford University Press.

Worrall, John . 2007. "Why There's No Cause to Randomize." *The British Journal for the Philosophy of Science* 58 (3): 451–488.

Xiao, Kai , Logan Engstrom , Andrew Ilyas , and Aleksander Madry . 2020. "Noise or Signal: The Role of Image Backgrounds in Object Recognition." arXiv: 2006.09994 [cs.CV].

Zagzebski, Linda Trinkaus . 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge University Press.