

Contents

	Preface and Acknowledgments	ix
1	Introduction	1
2	Basics of Impact Evaluation	7
	2.1 The Fundamental Problem of Impact Evaluation	7
	2.2 Characterizing the Impact	8
	2.3 The Problem of Comparing Apples to Oranges	12
3	Experiments (A/B Testing)	19
	3.1 Comparing Apples to Apples	19
	3.2 Behavioral Assumptions and Methods for Analyzing Experiments	22
	3.3 Multiple Interventions	26
	3.4 Use Cases in R	29
	3.5 Use Cases in Python	33
4	Selection on Observables: Aim to Compare Apples with Apples	37
	4.1 Making Groups Comparable in Observed Characteristics	37
	4.2 Behavioral Assumptions	41
	4.3 Methods for Impact Evaluation	42
	4.4 Use Cases in R	49
	4.5 Use Cases in Python	54
5	Causal Machine Learning	59
	5.1 Motivating Causal Machine Learning	59
	5.2 Elements of Causal Machine Learning	63
	5.3 A Brief Introduction to Several Machine Learning Algorithms	64
	5.4 Effect Heterogeneity and Optimal Policy Learning	69
	5.5 Use Cases in R	74
	5.6 Use Cases in Python	80
6	Instrumental Variables	87
	6.1 Instruments and Complier Effects	87
	6.2 Behavioral Assumptions and Methods	89

6.3	Use Cases in R	93
6.4	Use Cases in Python	96
7	Regression Discontinuity Designs	99
7.1	Sharp and Fuzzy Regression Discontinuity Designs	99
7.2	Behavioral Assumptions and Methods	102
7.3	Use Cases in R	107
7.4	Use Cases in Python	109
8	Difference-in-Differences	111
8.1	Difference-in-Differences and the Impact in the Treatment Group	111
8.2	Behavioral Assumptions and Extensions	114
8.3	Use Cases in R	117
8.4	Use Cases in Python	119
9	Synthetic Controls	123
9.1	Impact Evaluation When a Single Unit Receives the Intervention	123
9.2	Behavioral Assumptions and Variants	124
9.3	Use Cases in R	128
9.4	Use Cases in Python	131
10	Conclusion	133
	References	135
	Index	139