

BOUNDED RATIONALITY and INDUSTRIAL ORGANIZATION



RAN SPIEGLER

Bounded Rationality and Industrial Organization

Ran Spiegler



OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Oxford University Press, Inc., publishes works that further
Oxford University's objective of excellence
in research, scholarship, and education.

Oxford New York

Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in

Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Copyright © 2011 by Oxford University Press

Published by Oxford University Press, Inc.
198 Madison Avenue, New York, New York 10016
www.oup.com

Oxford is a registered trademark of Oxford University Press

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
electronic, mechanical, photocopying, recording, or otherwise,
without the prior permission of Oxford University Press.

Library of Congress Cataloging-in-Publication Data
Spiegler, Ran.

Bounded rationality and industrial organization / Ran Spiegler.
p. cm.

Includes bibliographical references and index.

ISBN 978-0-19-539871-7 (cloth : alk. paper)

1. Consumer behavior. 2. Decision making. 3. Consumption
(Economics)—Psychological aspects. 4. Industrial organization
(Economic theory) I. Title.

HF5415.32.S67 2011

658.8'342—dc22 2010015046

3 5 7 9 8 6 4 2

Printed in the United States of America
on acid-free paper

■ PREFACE

This book summarizes and synthesizes recent developments in the theory of Industrial Organization that incorporate departures from the standard model of consumer behavior in the direction of a “richer psychology.” It cannot be denied that the challenge of enriching the psychology of decision makers in economic models has been at the very frontier of theoretical research in the last decade, given an enormous boost by the behavioral economics movement.

However, while the subject has given rise to a large and varied literature that includes numerous research articles as well as a number of surveys and anthologies, one can think of very few proper *textbooks* that could serve a graduate-level course in economic theory. Indeed, the most recent theory-oriented textbook of any relevance that I am aware of is Ariel Rubinstein’s 1998 *Modeling Bounded Rationality*, which preceded many of the developments that caused such a stir in our profession. This book aims to narrow this gap, albeit in the very specific domain of industrial organization.

The book is meant to serve as a textbook for graduate courses in microeconomic theory, as well as theory-oriented courses in industrial organization or behavioral economics. It is partly based on lecture notes from courses given at Tel Aviv University, University College London and the Helsinki Center of Economic Research. In the course of writing this book, I have greatly benefitted from financial support by the European Research Council and the ESRC (UK). Parts of the research it is based on were supported by BSF and ISF grants.

I wish to thank several colleagues who gave comments on earlier drafts of book chapters: Eddie Dekel, Kfir Eliaz, Susana Esteban, Ignacio Esponda, Erik Eyster, Yves Guéron, Michael Grubb, Paul Heidhues, Philippe Jehiel, Barton Lipman, Marco Mariotti, Erik Mohlin, Michele Piccione, Jidong Zhou, and especially Ayala Arad for her consistently superb feedback. I am hugely indebted to Ariel Rubinstein, who not only contributed concrete suggestions that helped me improve the exposition and proofs of several results, but also provided valuable encouragement throughout this project. Ariel got me “hooked” on the subject of bounded rationality fifteen years ago, and I have continued to benefit from his work and his thoughts ever since. Ayala Arad and Yves Guéron have provided truly spectacular research assistance. In particular, Yves helped typesetting the manuscript and wrote the solution manual. I am forever grateful to both. My coauthors, Kfir Eliaz and Michele Piccione, deserve some of the credit for the material in the book, as our joint work has formed the basis of several chapters. The Oxford University Press editors, Terry Vaughn and Joe Jackson, provided warm and reliable support. Finally, I wish to thank the composer Steve Reich, whose music provided an indispensable soundtrack for the writing process. If the book feels monotonous or repetitive at times, blame it on him.

■ CONTENTS

<i>Preface</i>	<i>ix</i>
----------------	-----------

1 Introduction	1
1.1 Bibliographic Notes	7

PART ONE ■ Anticipating Future Preferences

2 Dynamically Inconsistent Preferences I: Unconstrained Contracting	11
2.1 The Multi-Selves Model	12
2.1.1 Naivete	13
2.2 Monopoly Pricing	14
2.2.1 Optimal Price Schemes for Sophisticated Consumers	15
2.2.2 Optimal Price Schemes for Naive Consumers	15
2.2.3 Screening the Consumer's Type	18
2.3 Competitive Pricing	18
2.4 Welfare Analysis	20
2.5 Educating Naive Consumers	21
2.6 The Interpretation of Naivete	22
2.7 Two Applications	23
2.8 Other Topics	25
2.8.1 The (β, δ) Model	25
2.8.2 Preference Heterogeneity	26
2.9 Summary	28
2.10 Bibliographic Notes	28
3 Dynamically Inconsistent Preferences II: Constrained Contracting	30
3.1 Two-Part Tariffs	30
3.1.1 Departure from Marginal-Cost Pricing	31
3.1.2 Welfare Analysis	33
3.2 Destabilization of Commitment Devices: Renegotiation and Spot Market Competition	34
3.3 Self-Control	36
3.3.1 Implications for Monopoly Pricing	39
3.3.2 Do Self-Control Costs Hamper Competition?	40
3.4 Summary	41
3.5 Bibliographic Notes	42

4	Dynamically Inconsistent Preferences III: Partial Naivete	43
4.1	Magnitude Naivete	43
4.1.1	Monopoly Pricing	43
4.1.2	Are More Sophisticated Consumers Always Better Off?	45
4.2	Frequency Naivete	46
4.2.1	First-Best Monopoly Pricing	47
4.2.2	Second-Best Monopoly Pricing	48
4.2.3	Does Competition Curb Exploitation?	51
4.3	Summary	52
4.4	Bibliographic Notes	52
5	Biased Beliefs without Dynamic Inconsistency	53
5.1	Monopoly Pricing with Over-Optimistic Consumers	54
5.1.1	Comparison with Related Models	57
5.2	Overconfidence: Three-Part Tariffs	60
5.3	Unforeseen Contingencies: Add-On Pricing	62
5.4	A Summary Exercise: Insurance Markets with Biased Consumers	67
5.4.1	Equilibrium Analysis When Subjective Beliefs Are Observable	69
5.4.2	Equilibrium Analysis When Subjective Beliefs Are Private Information	70
5.5	Summary	73
5.6	Bibliographic Notes	73
PART TWO ■ Responding to Market Complexity		
6	Sampling-Based Reasoning I: Price Competition and Product Differentiation	77
6.1	A Sampling-Based Choice Procedure	77
6.2	Price Competition and Technology Adoption	78
6.2.1	Nash Equilibrium	80
6.2.2	Welfare Analysis	82
6.3	Spurious Product Differentiation	84
6.3.1	Nash Equilibrium	85
6.3.2	Product Complexity as a Differentiation Device	86
6.4	Can the Market Educate Consumers?	90
6.5	Summary	92
6.6	Bibliographic Notes	92
7	Sampling-Based Reasoning II: Obfuscation	94
7.1	A Model of Competitive Obfuscation	94
7.1.1	Nash Equilibrium	96
7.1.2	Welfare Analysis	100
7.2	Production Inefficiencies	100
7.3	Multi-Dimensional Prices	103

7.4	A Market Intervention: Introducing “Simple” Options	104
7.5	Summary	107
7.6	Bibliographic Notes	108
8	Coarse Reasoning	110
8.1	A Modeling Framework	110
8.2	Complex Price Patterns as a Discrimination Device	112
8.2.1	“DeBruijn” Price Sequences	114
8.2.2	Conditions for Profitability of Complex Price Patterns	115
8.3	Limited Understanding of Adverse Selection	118
8.3.1	A Buyer-Seller Example	119
8.3.2	A Benchmark: A Bayesian-Rational Buyer	120
8.3.3	A “Coarse” Buyer	120
8.3.4	Action-Dependent Feedback	122
8.4	Summary	123
8.5	Bibliographic Notes	124
PART THREE ■ Reference Dependence		
9	Loss Aversion	127
9.1	Expected Price as a Reference Point: Monopoly Pricing	128
9.1.1	Reduced Price Variability	130
9.1.2	Impact on Average Prices	134
9.2	Price Uniformity in a Duopoly Setting: “Kinked” Demand	135
9.3	Expected Consumption as a Reference Point: An “Attachment Effect”	138
9.3.1	Personal Equilibrium	139
9.3.2	Price Randomization	141
9.4	Discussion	143
9.4.1	Actual Prices as Reference Points	143
9.4.2	Pleasant Surprises	144
9.5	Summary	145
9.6	Bibliographic Notes	146
10	Inertia I: Price Competition	147
10.1	Price Competition under Consumer Inertia	148
10.2	Price-Frame Competition	151
10.2.1	Nash Equilibrium	153
10.2.2	Equilibrium Properties	157
10.2.3	Two Market Interventions	158
10.3	Consumer Switching	159
10.4	Asymmetric Default Assignment	160
10.5	A Few General Remarks	161
10.5.1	More Than Two Frames	161
10.5.2	Revealed Preferences	163

10.6	Summary	163
10.7	Bibliographic Notes	164
11	Inertia II: Costly Marketing	166
11.1	A Model of Competitive Marketing	167
11.2	Nash Equilibrium	170
11.3	The Effective Marketing Property	176
11.4	Discussion	178
11.5	Summary	180
11.6	Bibliographic Notes	180
PART FOUR ■ Discussion		
12	Recurring Themes	183
12.1	Complex Pricing Strategies	183
12.2	Spurious Variety	184
12.3	Market Transactions as a Form of Speculative Trade	185
12.4	How Effective Are Competition and Consumer Protection Policies?	186
12.5	Externalities between Rational and Boundedly Rational Consumers	187
12.6	Conclusion	187
13	But Can't We Get the Same Thing with a Standard Model?	188
13.1	Rationalization via Modified Information	190
13.2	Rationalization via Modified Preferences	194
13.3	Rationalization via Endogenization	196
13.4	Discussion	199
13.5	Epilogue	200
13.6	Bibliographic Notes	201
A	Appendix to Part I: A Decision-Theoretic Perspective	202
A.1	The Multi-Selves Model	202
A.2	Self-Control Preferences	204
A.3	The Relation between Self-Control Preferences and the Multi-Selves Model	208
A.4	Other Classes of Temptation-Driven Preferences	210
A.5	Bibliographic Notes	211
	<i>Bibliography</i>	213
	<i>Index</i>	219

1

Introduction

Consider the following market situations you may have encountered as a consumer:

- You receive an invitation to accept a new credit card. You are reluctant to accept it because you are concerned that, armed with this new credit card, you will be tempted to spend more than you wish.
- You consider acquiring a new service at what seems like an attractive price. You suspect, however, that you may be unaware of certain future contingencies that will make it considerably more expensive than you can estimate at present.
- You try to choose a mutual fund from a vast array of market options, and you find the task of evaluating and comparing the numerous alternatives daunting. It is not that you lack information. On the contrary, information is superabundant, and it overwhelms you.
- You have spent the last hour reading the latest telephone bill, and you are still trying to figure out how the bottom line was reached.
- You listen to a sales pitch / watch a TV commercial / stare at the package of a product, and although these have no informational content, you feel more willing to consider buying the product.

These situations should be familiar to most members of an advanced economy. However, they present the modern economic theorist with a dilemma. Does conventional economic theory offer adequate tools for making sense of these situations and their welfare implications, or do we need new modeling tools? This book springs from the belief that the latter is true. In each of the above situations, there is an aspect of consumer psychology that takes the situation outside the scope of standard economic theory. My objective in this book is to present and synthesize recent theoretical research that has tried to formalize some of these aspects of decision making and examine their implications for market interactions—specifically, for the subfield of microeconomic theory that studies individual markets, namely Industrial Organization (henceforth I.O.).

The theory of I.O. has been developed in several major waves. The first wave involved the formulation of basic models of market structure (particularly perfect competition and monopoly). The second and third waves came more or less simultaneously in the 1970s. One wave saw the introduction of Game Theory into I.O., enabling thorough exploration of static and dynamic models of oligopoly. The other incorporated asymmetric information into I.O. models. Economists have come to realize that informational asymmetries constitute a major source of market friction, and that certain market and non-market institutions may be explained as responses to this market failure.

Throughout these developments, a maintained assumption has been that all market agents are rational. They are assumed to hold well-defined preferences. Although the rational choice paradigm allows preferences to be defined over very general domains, in virtually all I.O. applications rationality is narrowly practiced: preferences are defined over “simple” consequences, fully specified by the amount of money the consumer pays and the quantity or quality of the product he consumes. In addition, agents in standard models have full understanding of the market model. When they are imperfectly informed, they have perfect ability to draw Bayesian inferences in accordance with correct knowledge of the market model and market equilibrium.

Bounded rationality is another potential source of market friction. When some agents have limited understanding of their market environment (including their own behavior in certain circumstances), limited ability to process information, and preferences that are highly unstable, context-dependent, and malleable, market outcomes may differ in interesting and economically significant ways from the rational-consumer benchmark. Moreover, introducing boundedly rational agents into our market models may challenge conventional wisdom regarding the welfare properties of market interactions.

Here are a few of the theoretical questions I address:

- Can we view certain aspects of firms’ pricing, marketing, and product differentiation strategies as responses to consumers’ bounded rationality?
- To what extent are boundedly rational consumers vulnerable to exploitation by profit-maximizing firms? Does market competition protect consumers from being exploited?
- Does interaction between firms and boundedly rational consumers give rise to inefficiencies, and how are these affected by competition?
- What is the impact of various regulatory interventions in markets with boundedly rational consumers?
- Do market forces impel firms to “educate” or “de-bias” boundedly rational consumers?
- Does greater consumer rationality imply more competitive market outcomes?
- What sort of methods do firms use to discriminate between consumers according to the type and magnitude of their deviation from perfect rationality?

As can be gleaned from these questions, the book maintains an important distinction: firms will always be rational profit maximizers with a correct understanding of the market model, as in standard theory. On the other hand, consumers will depart from the standard model. The primary justification for this simplistic dichotomy is that a firm is more likely to conform to the standard model, in the sense that it focuses its attention, intelligence, and internal organization on a small set of markets. In contrast, consumers devote a fraction of their attention and intelligence to any individual market. Firms interact repeatedly with the market and therefore have many opportunities to learn its regularities. In contrast, consumers often have limited opportunities to learn the market model and the

market equilibrium. Firms deliberately apply systematic reasoning, relying on experts and statistical data, whereas consumers often rely on intuition. These are essentially asymmetries in *rationality*.

This is not to say that firms do not have their own types of non-rational behavior. For example, firms are perhaps more vulnerable than individual consumers to the follies of groupthink. Communication failures and frequent staff changes within the firm can lead to weak organizational memory. However, this book completely abstracts from these issues. The distinction between rational, profit-maximizing firms and boundedly rational consumers is maintained throughout.

■ WHY BOUNDED RATIONALITY?

The motivation for market models that involve boundedly rational consumers is threefold. First, there are the casual, everyday observations of the type that traditionally inspired economic thought, with which I opened this book. Second, there is a growing sentiment among economists that certain phenomena (advertising, consumption of addictive goods, complicated financial products) are not captured in a satisfactory manner by standard rational-choice models. Finally, experimental psychologists have made a powerful case for the claim that decision makers systematically deviate from the model of rational choice as it is typically practiced by economists.

The term “bounded rationality” is notoriously vague. Many readers (not to mention authors of some of the works I shall discuss) would justly view some of the decision models as entirely compatible with rationality—even if not in the narrow sense described above and practiced in traditional I.O. models. Rubinstein (1998) classifies decision models in the bounded-rationality category when they explicitly incorporate procedural elements of decision making that are absent from the standard model of rational choice. Gilboa & Schmeidler (2001) use bounded rationality as a way to judge decision-making quality: a decision is not rational if it embarrasses the decision maker once the situation is explained to him. Other economists view decision models with unstable and context-dependent preferences as manifestations of bounded rationality, even if the modeling tools they end up using are more or less standard.

Therefore, the book’s title is inescapably a misnomer. It might just as well have been called “I.O. Applications of Non-standard Models of Consumer Behavior.” My main excuse for using the B-word is that I think it will create the right expectations among its potential readership. My other excuse is that it is no less a misnomer than “rational expectations.”

■ WHY I.O.?

Bounded rationality manifests itself in many environments. Why did I choose to focus entirely on applications to I.O.? First, this is an area of fundamental importance in economics. The consumer is the decision unit that receives the

greatest attention in our undergraduate textbooks, and the analysis of how rational consumers interact with profit-maximizing firms in single markets is the bedrock of undergraduate microeconomics. For this reason, there is something very basic about investigating the implications of introducing boundedly rational consumers into individual markets.

Second, there is a methodological advantage in focusing on I.O. models. Throughout this book, consumers will be *non-strategic* agents, just as in the basic market models familiar from undergraduate microeconomics. The only strategic agents in these models will be the firms. Since firms will be assumed to be rational profit maximizers, it will be legitimate to analyze their behavior with standard tools (constrained maximization, Nash equilibrium). This greatly simplifies analysis and allows us to gain more mileage than if we also tried to construct and analyze models in which boundedly rational agents are also strategic players.

Third, narrowing the scope of investigation to a certain class of I.O. models enables me to achieve, within these restrictions, greater coherence and generality than is often the case with treatments of bounded rationality. It will allow me to trace recurring questions and themes and suggest generalizations.

However, none of these arguments would be relevant if this were not an area that has seen significant progress in the last decade. The literature on I.O. models with boundedly rational consumers is not huge. At this stage of its development, it does not justify a “bible” akin to Jean Tirole’s 1988 *Theory of Industrial Organization*. However, in my opinion, it is ripe for a concise, pedagogical synthesis that distills the literature’s major recurring questions as well as its main theoretical insights. This is my task in this book.

■ PLAN OF THE BOOK

The book presents a few strands in the development of I.O. models with boundedly rational consumers. Each of its first three parts addresses a different major aspect of consumer psychology and examines its implications for pricing, marketing, or product differentiation strategies employed by firms, as well as for welfare analysis.

Part I: Anticipating Future Preferences

This part presents a collection of closely related dynamic models, in which consumers are confronted with price schemes that pertain to future consumption decisions. It mostly deals with *dynamically inconsistent preferences*, covering a large class of situations in which consumer tastes change over time as a result of various psychological forces, such as temptations that appeal to visceral urges, addiction, or changing reference points. A key feature of the models analyzed in this part is that some consumers are “naive,” in the sense that they fail to fully appreciate the likelihood or magnitude of future taste changes. Variants include an extended model that allows consumers to exercise self-control in response to

temptations, and a model with biased beliefs about future preferences without dynamic inconsistency.

Part II: Responding to Market Complexity

This part examines situations in which market alternatives are complicated (either inherently or as a result of firms' deliberate obfuscation), and the ways in which boundedly rational consumers cope with this complexity. First, I introduce a decision model in which consumers evaluate market alternatives by sampling small parts and extrapolating naively. Second, I study a model in which consumers reason in terms of "coarse representations" of market alternatives.

Part III: Reference Dependence

This part examines decision models in which consumer choice is sensitive to reference points. I begin by tackling situations in which preferences exhibit loss aversion, and then turn to models in which consumers' default options exert a pull on their decision process, resulting in consumer inertia.

While each part is based on a different aspect of the psychology of decision making, the treatment is unified by several factors. First, the non-standard decision models are invariably embedded in I.O. models that are among the most basic in economics, such that if consumers were rational, everything would collapse to an undergraduate textbook model.

Second, there is considerable overlap among the I.O. questions analyzed in the three parts. For the most part, I explore the implications of non-standard decision models for the firms' pricing behavior (including price discrimination, the structure of price plans, and marketing effects such as add-on pricing or the use of irrelevant alternatives), product differentiation strategies, and welfare analysis (including the effect of regulatory interventions).

Third, the three topics studied in Parts I–III are interrelated. The formation of biased beliefs about future preferences is closely related to the problem of evaluating a complex contract when the number of possible future contingencies is so large that the consumer may be unaware of some of them. Likewise, the complexity of evaluating market alternatives may strengthen consumer inertia. When preferences are sensitive to a reference point, they may display dynamic inconsistency because the reference point may shift over time. These are only a few of the examples I could give for such links.

Although the selection of topics for Parts I–III covers a significant portion of the theoretical literature on I.O. with boundedly rational consumers, it is not intended to be exhaustive. For example, limited consumer memory is almost entirely absent from the book. I selected those topics I thought I knew how to turn into compelling lecture notes. And of course, there is the natural bias in favor of research I have been involved with myself.

Part IV concludes the book and takes a more global view of the material, by revisiting themes—substantive economic themes as well as methodological ones—that recur throughout the book.

■ HOW TO USE THE BOOK

The book is primarily meant to be used as a set of lecture notes for graduate courses in microeconomic theory, industrial organization, or behavioral economics. I hope that certain features of the presentation—simple pedagogical models, a well-defined domain of economic applications, complete proofs for most of the results, and close attention to choice-theoretic considerations—will make the book appealing to researchers in various fields.

The book is modular. Its three main parts can be studied independently of each other. The themes that recur in Parts I–III are discussed from a bird’s-eye view in Part IV. An instructor who wishes to give graduate students a taste of developments in economic theory that involve elements of bounded rationality can use any of the first three parts, coupled with a glance at the final part.

In terms of necessary background, the book presupposes that the reader knows basic microeconomic theory, preferably in the form of a core graduate course. It does not require any additional mathematical knowledge. Part I is perhaps the one most suited for advanced undergraduates.

■ STYLE OF THE BOOK

The pedagogical orientation of this book dictates a style of presentation that seeks the simplest framework that could express the ideas I wish to convey. Thus, instead of presenting a catalogue of the literature, in each chapter I study thoroughly a single model (or a family of closely related models) and use it to synthesize and adapt ideas from a number of papers, which themselves often follow diverse modeling approaches. This has the unavoidable consequence that my mode of presentation and emphases often depart from the original works. In this sense, this book occasionally fails to do full justice to the works it is based on. This aspect of the book’s style also forces me to relegate virtually all references to the bibliographic notes that appear at the end of almost every chapter.

Two decades ago, the quest for psychologically richer models of decision making was pursued by a small group of economic theorists. Since then, the success and wide exposure of behavioral economics have changed the field. As a result, many intellectual traditions have been brought into contact with issues that concern the psychology of decision making. This has often led to intense disagreements about the “right” way of doing research in this area.

Against this background, this book is eclectic, in the sense that its themes have been explored by economists from diverse traditions. Nevertheless, I maintain a uniform, distinctive style. Although the study of I.O. models with boundedly rational consumers may provide an intellectual background for consumer protection policies (analogous to the role that conventional I.O. theory plays in competition policy), this book has a clear theoretical orientation, and discussions of applications or policy implications are highly stylized. Models are constructed for their pedagogical and “story-telling” value. Although the models are deliberately simple, they are pitched at a level of abstraction that

immediately suggests generalizations and encourages links with the decision-theoretic literature. The I.O. models in which I embed the non-standard decision models are themselves extremely basic, and I do not complicate them in any way other than in the bounded rationality dimension. Although I invoke relevant psychological research to motivate certain models, there is almost no discussion of experimental psychology, and readers who are interested in this material are referred to excellent available sources.

■ PREVIEW OF THE MAIN THEMES

The book's repeated use of the most basic I.O. models as a template and the recurrence of certain general questions allow us to draw a few economic "lessons." It may be useful to list them at this stage:

- The presence of boundedly rational consumers may impel firms to adopt price schemes that are more complex than if consumers were rational (e.g., three-part tariffs, add-on pricing, excessive price variation over time and across circumstances).
- Boundedly rational consumers are often vulnerable to exploitative contracts. Competitive forces do not necessarily mitigate the exploitation, and may sometimes exacerbate it. Ordinary competition and consumer protection policies (increasing the number of competitors, introducing simple alternatives or simplifying existing ones) may be ineffective, and even counterproductive when consumers are boundedly rational.
- Consumers' bounded rationality is often a force that generates greater product differentiation, which is "spurious" in the sense that it does not enhance consumer welfare.
- It is not necessarily the case that the more rational the consumers, the more competitive the market outcome. A related point is that rational consumers do not always exert a positive externality on boundedly rational consumers.

The economic message that emerges from these "lessons" is critical of conventional accounts of the market mechanism. In particular, aspects of firms' behavior that are conventionally viewed as "healthy" responses to consumers' preferences and constraints are here viewed as irrelevant for consumer welfare at best and harmful at worst. At the same time, I exhibit traditional caution and skepticism regarding the power of regulatory intervention to curb the frictions that result from consumers' bounded rationality. To put it in coarse terms, consumers' bounded rationality can generate market failure, but it is far from clear whether it is a failure that can be "fixed."

■ 1.1 BIBLIOGRAPHIC NOTES

Thanks to the flourishing of behavioral economics, there is a multitude of excellent sources on the psychology of decision making and its relevance for economics. Kahneman, Slovic & Tversky (1982), Kahneman & Tversky (2000) and Thaler

(1994) are authoritative. Recent popular books include Gilbert (2006), Iyengar (2010), and Thaler & Sunstein (2008).

There are fewer texts on economic models of bounded rationality. Simon (1982) is a collection of Herbert Simon's early writings on the subject. Rubinstein (1998) offers a collection of bounded rationality models in various economic contexts. Ellison (2006) and Armstrong (2008) are surveys of I.O. models that incorporate bounded rationality, the latter being more policy-oriented than the former. Rabin (1998) is an early survey of behavioral economics. DellaVigna (2009) is a survey of empirical work in behavioral economics.



PART ONE

Anticipating Future Preferences