Revealed Preference Theory

Pioneered by American economist Paul Samuelson, revealed preference theory is based on the idea that the preferences of consumers are revealed in their purchasing behavior. Researchers in this field have developed complex and sophisticated mathematical models to capture the preferences that are "revealed" through consumer choice behavior. This study of consumer demand and behavior is closely tied up with econometrics (especially nonparametric econometrics), where testing the validity of different theoretical models is an important aspect of research. The theory of revealed preference has a very long and distinguished tradition in economics, but until now there has been no systematic presentation of the theory. This book deals with basic questions in economic theory, such as the relation between theory and data, and studies the situations in which empirical observations are consistent or inconsistent with some of the best-known theories in economics.

Christopher P. Chambers is an associate professor of economics at the Department of Economics, University of California, San Diego. He has published articles in the *American Economic Review*, *Econometrica*, *Games and Economic Behavior*, *Journal of Economic Theory*, and *Theoretical Economics*. He is an associate editor of the *Journal of Economic Theory*, *Journal of Mathematical Economics*, and *Social Choice and Welfare*. He is also co-editor of *Mathematical Social Sciences*.

Federico Echenique is a professor of economics at the Division of the Humanities and Social Science, California Institute of Technology. He has published articles in the *American Economic Review*, *Econometrica*, *Journal of Political Economy*, *Games and Economic Behavior*, *Journal of Economic Theory*, and *Theoretical Economics*. He is an associate editor of the *Journal of Economic Theory*, *Economic Journal*, *Economic Theory*, and *ACM Transactions on Economics and Computation*.

Econometric Society Monographs

Editors:

Jeffrey Ely, Northwestern University
Donald W. K. Andrews, Yale University

The Econometric Society is an international society for the advancement of economic theory in relation to statistics and mathematics. The Econometric Society Monograph series is designed to promote the publication of original research contributions of high quality in mathematical economics and theoretical and applied econometrics.

Other titles in the series:

G. S. Maddala, *Limited dependent and qualitative variables in econometrics*, 9780521241434, 9780521338257
Gerard Debreu, *Mathematical economics: Twenty papers of Gerard Debreu*, 9780521237369, 9780521335614
Jean-Michel Grandmont, *Money and value: A reconsideration of classical and neoclassical monetary economics*, 9780521251419, 9780521313643
Franklin M. Fisher, *Disequilibrium foundations of equilibrium economics*, 9780521378567
Andreu Mas-Colell, *The theory of general equilibrium: A differentiable approach*, 9780521265140, 9780521388702
Truman F. Bewley, Editor, *Advances in econometrics – Fifth World Congress (Volume I)*, 9780521467261
Truman F. Bewley, Editor, *Advances in econometrics – Fifth World Congress (Volume II)*, 9780521467254
Hervé Moulin, *Axioms of cooperative decision making*, 9780521360555, 9780521424585
L. G. Godfrey, *Misspecification tests in econometrics: The Lagrange multiplier principle and other approaches*, 9780521424592
Tony Lancaster, *The econometric analysis of transition data*, 9780521437899
Alvin E. Roth and Marilda A. Oliviera Sotomayor, Editors, *Two-sided matching: A study in game-theoretic modeling and analysis*, 9780521437882
Wolfgang Hardle, *Applied nonparametric regression*, 9780521429504
Jean-Jacques Laffont, Editor, *Advances in economic theory – Sixth World Congress (Volume I)*, 9780521484596
Jean-Jacques Laffont, Editor, *Advances in economic theory – Sixth World Congress (Volume II)*, 9780521484602
Halbert White, *Estimation, inference and specification*, 9780521252805, 9780521574464
Christopher Sims, Editor, *Advances in econometrics – Sixth World Congress (Volume I)*, 9780521444590, 9780521566100
Christopher Sims, Editor, *Advances in econometrics – Sixth World Congress (Volume II)*, 9780521444606, 9780521566094
Roger Guesnerie, *A contribution to the pure theory of taxation*, 9780521629560
David M. Kreps and Kenneth F. Wallis, Editors, *Advances in economics and econometrics – Seventh World Congress (Volume I)*, 9780521589833
David M. Kreps and Kenneth F. Wallis, Editors, *Advances in economics and econometrics – Seventh World Congress (Volume II)*, 9780521589826
David M. Kreps and Kenneth F. Wallis, Editors, *Advances in economics and econometrics – Seventh World Congress (Volume III)*, 9780521580137, 9780521589819
Donald P. Jacobs, Ehud Kalai, and Morton I. Kamien, Editors, *Frontiers of research in economic theory: The Nancy L. Schwartz Memorial Lectures*, 1983–1997, 9780521632225, 9780521635387

(continued after Index)

Revealed Preference Theory

Christopher P. Chambers
Department of Economics,
University of California San Diego

Federico Echenique
Division of the Humanities and Social Science,
California Institute of Technology

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE
UNIVERSITY PRESS

University Printing House, Cambridge CB2 8BS, United Kingdom

Cambridge University Press is part of the University of Cambridge.

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

www.cambridge.org
Information on this title: www.cambridge.org/9781107087804

First published 2016

Printed in the United States of America by Sheridan Books, Inc.

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

Library of Congress Cataloguing in Publication data
Names: Chambers, Christopher P., 1975– author. | Echenique, Federico, author.
Title: Revealed preference theory / Christopher P. Chambers, Federico Echenique.
Description: New York : Cambridge University Press, [2016] |
Series: Econometric society monographs | Includes bibliographical references and index.
Identifiers: LCCN 2015040826| ISBN 9781107087804 (hardback) |
ISBN 9781107458116 (pbk.)
Subjects: LCSH: Revealed preference theory. | Consumption (Economics) |
Consumers' preferences–Econometric models.
Classification: LCC HB801.C3945 2016 | DDC 339.4/7–dc23
LC record available at http://lccn.loc.gov/2015040826

ISBN 978-1-107-08780-4 Hardback
ISBN 978-1-107-45811-6 Paperback

To Evan, Helena and Natalia

Contents

Preface

0.1 WHAT IS REVEALED PREFERENCE THEORY?

"Revealed preference" is a term with several interpretations in economic theory, all closely related but possessing subtle philosophical differences. The central theme common to all interpretations is that of understanding what economic models say about the observable world. Most practitioners of revealed preference theory recognize that an economic model is useful for organizing data and making predictions, but they would not go so far as to admit that the model is, or represents, "reality." Economic models usually consist of multiple interacting parts. Some parts are theoretical in nature, and are not meant to be observed. These unobservable parts are then tied to objects which could potentially be observed and measured.

Let us begin with the canonical example of revealed preference theory. Economists often view individuals as making decisions consistent with some objective function, usually interpreted as "utility." So, the theory posits that an individual chooses that option which gives her the highest utility among all feasible options. This has proved to be a useful and tractable model in many branches of economics; in fact, the concept is ubiquitous. But it has been recognized at least since Pareto (1906) that, even if an individual optimizes a utility, the cardinal structure of said utility function (that is, the numerical values assigned to potential choices by the utility) cannot be inferred from choice behavior. Utility is a "theoretical" concept. The only empirically meaningful statements that can be gleaned from data relate to whether one option exhibits a higher utility than another. Once this is recognized, it is natural to ask: what are the predictions of the utility maximization model for observable choice behavior? Obviously, the predictions will have to be made *across* different choice situations. A theory that claims an individual chooses the best available option from those in front of her will have little to say without imposing constraints on the notion of "best." So, the next step was to understand what these predictions are. These questions form the first few chapters of this book.

The revealed preference approach seeks to understand what a given model says about data. Observable economic data are either consistent with a given model or not. If the data are consistent with the model, this does not provide evidence that the model is right, but instead tells us that – at least in terms of economically observable phenomena – we may without loss of generality treat the economic system under consideration as if the model is right.

As an illustration of this phenomenon, the words "as if" are frequently applied in revealed preference analysis: a decision maker behaves "as if" she possesses a utility function, a firm behaves "as if" it is profit-maximizing, a group of agents behave "as if" they make agreements according to a Nash bargaining solution. Importantly, common to all uses of the "as if" qualification is the idea that there is some procedural aspect to behavior which is unobserved, and which cannot be observed using economic data.

It is true, however, that this "as if" approach often generates parameters which are intended to be applied outside of the model. For example, the theory of subjective expected utility can be used to elicit a probability measure over states of the world. Whether or not this probability measure actually represents a probability measure the agent in question possesses, it is often treated as such. Researchers often assume that the agent possesses a probability measure, and that the behaviorally elicited probability measure is a good enough proxy. Assumptions like these obviously cannot be tested using economic data, and thus fall in the realm of philosophy.

What does the standard revealed preference exercise look like? A model is specified by hypothesizing the existence of some objects, which are not observable as data, even in principle. Thus, a revealed preference theory is existential in nature. We can refer to these unobservable objects as theoretical objects. These objects are hypothesized to interact with observable objects in a certain way. Generally speaking, theoretical objects live in a large space (though this need not always be the case). The standard example of the exercise we have in mind is the utility function over commodity space: the set of utility functions is obviously quite large. We can never see a utility function, but what we might be able to see are demand observations at a finite list of prices. Thus the question arises: does there exist a utility function which could generate these demand observations by a process of maximization? In general, such questions are by no means trivial to answer. For the utility example, because there are an infinite number of potential utility functions, searching through every possible utility function to verify whether or not one of them could lead to observed data is an impossible task. A different approach is needed, some way of simplifying the problem. This approach is revealed preference theory.

We wish to emphasize that the role of revealed preference theory is not to make any claim on the notion of what constitutes an economic theory, or to discuss the "realism" or meaning of theoretical concepts in an economic theory.[1] In fact, revealed preference theory has nothing to say about the

[1] Some essays by Paul Samuelson (1963, 1964, 1965) have sometimes been interpreted as suggesting that a theory is equivalent to its theoretical predictions.

interpretation of these theoretical concepts. Rather, the driving force of revealed preference theory is understanding the implications of our theories and models for the real world.[2] Similarly, while the theory has been formulated for choice data, there is nothing inherent in revealed preference theory that precludes using other sources of datasets.[3]

0.2 DIFFERENT APPROACHES IN REVEALED PREFERENCE

In broad terms, revealed preference theory can be divided into two main subfields. The first assumes that all relevant and potentially observable economic data are in fact observed. For example, the integrability question of demand theory asks when a fully specified demand function can be rationalized by an (unobserved) preference relation which the consumer maximizes. Another example is standard in decision theory, where one observes a complete preference relation, meaning that one observes the outcome of all pairwise comparisons between objects of choice. We study a few such questions in this book, but not many. On the other hand, the second subfield assumes that economic data are only partially observable. We observe consumer demand a finite number of times. We observe production a finite number of times. Or we see how agents interact a finite number of times.

Within this second subfield, there are again two underlying themes, which are two sides of the same coin. One theme deals with refutability, or falsifiability. This approach studies the revealed preference implications of a model by studying what the model rules out. Given data, we want to check whether or not these data falsify the model in question. A canonical example is the equivalence of the generalized axiom of revealed preference (GARP) and rationalizability. GARP is a statement that rules out "revealed preference cycles," any demonstration of which falsifies the model. So, the idea is to characterize those potential observations which would directly falsify the model. GARP characterizes these observations as those which exhibit "preference cycles." The absence of a preference cycle, on the other hand, says nothing about the truth of the model. It only establishes that the model has not been falsified.

An alternative approach attempts to verify that given data are consistent with the model. The driving force in this framework is that of "reducing" the search for unobservables to a smaller space. The canonical example of this approach is the equivalence of the Afriat inequalities and the hypothesis of maximizing a utility function. The Afriat inequalities are a finite list of linear inequalities that are derived directly from observed data. If the inequalities have a common

[2] In a final chapter, we will describe a "theory" as a kind of class of potential relations that might hold in the real-world. This turns out to be a useful construction for understanding the empirical content that a theory has, but it may not conform to the notion of a theory as philosophers of science define it.

[3] Indeed, the most common empirical application of revealed preference theory uses data from non-incentivized consumption surveys. These surveys are not, strictly speaking, choice data.

solution, then one can claim that the observed data could have been generated by utility maximization; in other words, the model has not been falsified. On the other hand, if there is no such solution, then we have falsified the model. Directly verifying that data are consistent with utility maximization involves searching an infinite-dimensional space (the space of utility functions). Finding a solution to the Afriat inequalities involves searching a finite dimensional space. Indeed, there are well-known linear programming algorithms which can be used to determine whether the Afriat inequalities have a solution. However, with either approach, the hypothesis refers to "unobservable" or "theoretical" concepts (the utility function in the first place, the solution to the inequalities in the other).

In this book, we focus on the first approach. Our interest in this book is in understanding the data which falsify a given model. It is natural that our interest will therefore be in directly characterizing such datasets. However, it is important to understand that, quite generally, the two approaches are mathematically equivalent, and neither is "better" than the other. At times, one approach, and at other times the other, will be more useful. To illustrate, consider Afriat's Theorem: this beautiful theorem establishes that GARP, satisfaction of the Afriat inequalities, and consistency with the maximization of a utility function are all equivalent. Thus, Afriat's Theorem not only provides a direct method of falsifying the model (GARP), but a simple method of verifying when the model is not refuted (the Afriat inequalities). It is now know that this duality holds relatively generally; this was first described in economics by Brown and Matzkin (1996). In practice, however, the more difficult problem is often in exhibiting classes of data which refute the given model, eliminating all unobservable and theoretical concepts.

0.3 SCOPE OF THE BOOK

This book arose from courses taught by the authors at Caltech and UCSD, and by no means claims to be exhaustive, complete, or definitive: the selection of material covered inevitably reflects our own interests and expertise. There are many areas of revealed preference theory, both old and modern, which we discuss very little. In particular, we include no discussion whatsoever of the theory of household consumption, which is an area of revealed preference theory that has been especially active. In addition, there are already several fine existing surveys of revealed preference theory; Richter (1971), Mas-Colell (1982), Carvajal, Ray, and Snyder (2004), Varian (2006), Cherchye, Crawford, De Rock, and Vermeulen (2009), and Crawford and De Rock (2014) are excellent and highly recommended references which have helped us to shape our own ideas.

Our presentation of the material emphasizes some concepts that we have found particularly important in our own understanding and teaching of revealed preference theory.

In the classical choice or demand theoretic setting, while there are several objects that might rightfully be called a revealed preference relation, the literature often refers to revealed preference as if it is a single concept. This can lead to substantial confusion on the part of students. While everyone readily understands the notion that x is revealed preferred to y if x is chosen when y is available, somewhat more subtle is the notion of revealed strict preference: that x is chosen when y is available and y is not chosen. The subtlety lies in the fact that the revealed strict preference need not be the strict part of the revealed preference. Indeed, the role of the most famous of revealed preference axioms, the weak axiom of revealed preference, is simply to ensure that revealed strict preference does not contradict revealed preference: it is impossible for x to be revealed preferred to y while y is revealed strictly preferred to x. We chose, therefore, to introduce a technology in this book which makes these concepts clear. Instead of referring to a revealed preference, we refer throughout to a revealed preference *pair*: the pair consisting of the revealed preference and the revealed strict preference.

Much of this book focuses on the idea of "eliminating" existentially quantified theoretical concepts, as described in our discussion of utility maximization above. The elimination of quantifiers can often be facilitated by some simple mathematical results which prove to be extremely powerful. One of the important mathematical results that we use over and over is called the Theorem of the Alternative. Geometrically, it is very simple. At the same time, it is very powerful, and often lets us convert existential quantification in one set of variables to universal quantification in another. It relates pairs of linear inequalities. The theorem tells us that one of the systems (the primal) has a solution if and only if the other (the dual) does not. The key aspect is that these two systems are systems of inequalities in different variables. The variables of the primal correspond to the constraints of the dual, and conversely. Often, we are lucky enough to be able to reduce the search for a theoretical object to a list of linear inequalities. For example, it may not be obvious, but the existence of a utility function representing a preference on a finite set is equivalent to the existence of a solution to a collection of linear inequalities, where the utility function is viewed as a vector living in a finite-dimensional space. By the Theorem of the Alternative, there is no solution to this collection of linear inequalities if and only if the dual set of linear inequalities has a solution. Thus, falsifying the hypothesis of the existence of a utility is simplified: instead of demonstrating that all potential solutions cannot be solutions, we simply have to demonstrate one solution to the dual set of linear inequalities. And the dual set of linear inequalities has a variable for each constraint of the primal set. It turns out that in this context, one can show that a solution to the dual set of linear inequalities corresponds to a revealed preference cycle, so that a utility function exists if and only if there is no revealed preference cycle. Importantly, what the Theorem of the Alternative has allowed us to do is to remove any existential mention of the theoretical object.

This technique proves useful because tests for many economic models can be reduced to collections of linear inequalities. But this is not always the case. Tests for some economic models turn out to be best studied by systems of polynomial inequalities. And in fact, for polynomial inequalities, there are similar notions to the Theorem of the Alternative. One, which has not yet found widespread application in economics, is called the Positivstellensatz. The Positivstellensatz claims that a system of polynomial inequalities has a solution if and only if some dual set has no solution. Though in general it may be difficult to test whether this dual set of polynomial inequalities has a solution, there are algorithms that seem to work well in practice. Another related idea is the Tarski–Seidenberg result. Tarski–Seidenberg asks whether a given collection of polynomial inequalities has a solution. It turns out that there is an algorithm for checking whether or not this is so (this algorithm, unfortunately, is generally computationally very costly).

The Tarski–Seidenberg algorithm is a beautiful application of a model-theoretic construct called quantifier elimination. Quantifier elimination is a rare property that certain theories can have; and it just so happens that the theory of real closed fields (of which the real numbers are an example) has this property. Quantifier elimination is said to hold if existentially quantified first-order sentences can always be shown to be equivalent to sentences with no existential quantification. As most revealed preference exercises are existential in nature (that is, one wants to search for the existence of a utility function, or preference, or choice correspondence), quantifier elimination proves to be an extremely powerful tool in revealed preference theory.

The book is laid out into chapters whose titles are, for the most part, self-explanatory. Chapter 1 begins with some mathematical preliminaries which will be useful throughout the text. Chapters 2–6 investigate the revealed preference question in classical environments, ranging from abstract choice, to demand, to production. Chapter 7 is an investigation into the stochastic choice model, which generalizes the abstract choice-theoretic model to allow choice over options to be random. Chapter 8 describes some of the basic economic models of uncertainty, and how they can be investigated with revealed preference techniques. Chapters 9–11 are devoted to studying revealed preference-style questions in models of group choice; general equilibrium theory, game theory, and political science topics are addressed here.

The final two chapters are more methodological and abstract. Chapter 12 formalizes an underlying theme of the book: that finding rationalizations in many revealed preference models turns out to be equivalent to finding a satisfaction to a collection of polynomial inequalities. This observation is not new, of course, but we describe some tools deriving from the mathematical discipline of real algebraic geometry which we believe could be useful for practitioners. Chapter 13 is devoted to studying notions of revealed preference, empirical content, and falsifiability from an abstract perspective. This chapter is mostly based on the authors' joint work with Eran Shmaya.

Finally, we are very grateful to Kim C. Border, Jun Chen, SangMok Lee, Ce Liu, Efe Ok, Juan Sebastián Pereyra, John N. Rehbeck, Eran Shmaya, and Gerelt Tserenjigmid for corrections, references, and very useful comments on early drafts of the book. Kevin Li, in particular, proofread most of the material and helped us improve the writing and the exposition in the book significantly. We also received very useful feedback from five anonymous reviewers from Cambridge University Press. Special thanks are due to our coauthors Yaron Azrieli, Takashi Hayashi, PJ Healy, SangMok Lee, Kota Saito, Matt Shum, Eran Shmaya, and Bumin Yenmez for work that has shaped our understanding of revealed preference theory, and for letting us use our joint research for some of the material included in the book.