Drew Fudenberg

David K. Levine

The Theory of Learning in Games

The Theory of Learning in Games

Drew Fudenberg
and
David K. Levine

The MIT Press
Cambridge, Massachusetts
London, England

This book was set in Palatino on the Monotype "Prism Plus" PostScript Imagesetter by Asco Trade Typesetting Ltd., Hong Kong.

Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Fudenberg, Drew.
The theory of learning in games / Drew Fudenberg and David K. Levine.
p. cm.—(Economic learning and social evolution)
Includes bibliographical references and index.
ISBN 0-262-06194-5 (hc : alk. paper)
1. Game theory. I. Levine, David K. II. Title. III. Series:
MIT Press series on economic learning and social evolution.
QA269.F83 1998
519.3—dc21 97-39957
CIP

The Theory of Learning in Games

MIT Press Series on Economic Learning and Social Evolution

General Editor

Ken Binmore, Director of the Economic Learning and Social Evolution Centre, University College London.

1. *Evolutionary Games and Equilibrium Selection,* Larry Samuelson, 1997

2. *The Theory of Learning in Games,* Drew Fudenberg and David K. Levine, 1998

To our wives, Geneen O'Brien and Joyce Davidson,
who have taught us much

Series Foreword

The MIT Press series on Economic Learning and Social Evolution reflects the widespread renewal of interest in the dynamics of human interaction. This issue has provided a broad community of economists, psychologists, philosophers, biologists, anthropologists, and others with a sense of common purpose so strong that traditional interdisciplinary boundaries have begun to melt away.

Some of the books in the series will be works of theory. Others will be philosophical or conceptual in scope. Some will have an experimental or empirical focus. Some will be collections of papers with a common theme and a linking commentary. Others will have an expository character. Yet others will be monographs in which new ideas meet the light of day for the first time. But all will have two unifying features. The first will be a rejection of the outmoded notion that what happens away from equilibrium can safely be ignored. The second will be a recognition that it is no longer enough to speak in vague terms of bounded rationality and spontaneous order. As in all movements, the time comes to put the beef on the table—and the time for us is now.

Authors who share this ethos and would like to be part of the series are cordially invited to submit outlines of their proposed books for consideration. Within our frame of reference, we hope that a thousand flowers will bloom.

Ken Binmore
Director
Economic Learning and Social Evolution Centre
University College London
Gower Street
London WC1E 6BT, England

Acknowledgments

As befits a book on learning, we have learned a great deal from our colleagues in the course of this project, so we have quite a number of people to thank. At the preliminary stages, Eddie Dekel and Glenn Ellison made useful suggestions about how to understand and organize the vast and growing literature that the book attempts to survey. Somewhat later, as we realized how little we knew about evolutionary game theory, Ken Binmore and Josef Hofbauer were very helpful in answering our questions and suggesting papers that we should include. On this topic, we also learned a great deal from Josef Hofbauer and Karl Sigmund's classic *The Theory of Evolution and Dynamical Systems*[1] and Jörgen Weibull's *Evolutionary Game Theory*.[2] Obviously the work discussed in chapters 4, 6, and 7 owes a great deal to David Kreps; we would like to give him a special thanks for many years of fascinating discussions and fruitful collaborations in the study of learning in games, and game theory more generally. Dean Foster and Rakesh Vohra introduced us to the computer science literature, much of which is referenced in chapter 8.

Once the first draft of the book was completed, we were fortunate to benefit from the comments and suggestions of many readers. Ken Binmore, Daniel Benjamin, Glenn Ellison, Dan Friedman, Sendhil Mullinaithan, and several anonymous reviewers suggested numerous improvements in the exposition of chapters 1 through 5. Dov Monderer identified some errors in an earlier draft of chapter 2. Josef Hofbauer, Klaus Nitzberger, Larry Samuelson, and Karl Schlag found mistakes in chapter 3; Larry also sent us detailed comments on that chapter and on the book as a whole. Michel Benaim, Glenn Ellison, George Mailath, and Peyton Young

1. English translation published by Cambridge University Press, 1988.
2. MIT Press, 1995.

helped us correct errors in chapter 5. Mathias Erlei caught some mistakes and suggested expositional improvements, especially in chapter 7. Leat Yarriv read chapter 8 very carefully; her suggestions led to many improvements and corrections. The discussion of stimulus-response models and experimental evidence benefited from a number of conversations and correspondence with Ken Binmore, Tilman Börgers, Al Roth, Larry Samuelson, and John Van Huyck, and correspondence with Gary Bolton and Ido Er'ev. We would also like to thank Stefano Dellarigna for careful proofreading.

We would like to thank MIT Press for producing and distributing this book; we especially thank Dana Andrus for her careful copyediting and Terry Vaughn for his encouragement. We are grateful to the UCLA Academic Senate and National Science Foundation grants SBR-93-20695 and SBR-94-24013 for financial support. Most of all, we thank our wives for their forbearance and support those nights when revising a chapter temporarily took priority over spending time with our families.

The Theory of Learning in Games

Contents

1 Introduction

1.1 Introduction

This book is about the theory of learning in games. Most of non-cooperative game theory has focused on equilibrium in games, especially Nash equilibrium and its refinements such as perfection. This raises the question of when and why we might expect that observed play in a game will correspond to one of these equilibria. One traditional explanation of equilibrium is that it results from analysis and introspection by the players in a situation where the rules of the game, the rationality of the players, and the players' payoff functions are all common knowledge. Both conceptually and empirically, these theories have many problems.[1]

This book develops the alternative explanation that equilibrium arises as the long-run outcome of a process in which less than fully rational players grope for optimality over time. The models we will discuss serve to provide a foundation for equilibrium theory. This is not to say that learning models provide foundations for all of the equilibrium concepts in the literature, nor does it argue for the use of Nash equilibrium in every situation; indeed, in some cases most learning models do not lead to any equilibrium concept beyond the very weak notion of rationalizability.

1. First, a major conceptual problem occurs when there are multiple equilibria, for in the absence of an explanation of how players come to expect the same equilibrium, their play need not correspond to any equilibrium at all. While it is possible that players coordinate their expectations using a common selection procedure such as Harsanyi and Selten's (1988) tracing procedure, left unexplained is how such a procedure comes to be common knowledge. Second, we doubt that the hypothesis of exact common knowledge of payoffs and rationality apply to many games, and relaxing this to an assumption of almost common knowledge yields much weaker conclusions. (See, for example, Dekel and Fudenberg 1990 and Borgers 1994.) Third, equilibrium theory does a poor job explaining play in early rounds of most experiments, although it does much better in later rounds. This shift from non-equilibrium to equilibrium play is difficult to reconcile with a purely introspective theory.

Nevertheless, learning models can suggest useful ways to evaluate and modify the traditional equilibrium concepts. Learning models lead to refinements of Nash equilibrium; for example, considerations of the long run stochastic properties of the learning process suggest that risk dominant equilibria will be observed in some games. They lead also to descriptions of long-run behavior weaker than Nash equilibrium; for example, considerations of the inability of players in extensive form games to observe how opponents would have responded to events that did not occur suggests that self-confirming equilibria that are not Nash may be observed as the long-run behavior in some games.

We should acknowledge that the learning processes we analyze need not converge, and even when they do converge, the time needed for convergence is in some cases quite long. One branch of the literature uses these facts to argue that it may be difficult to reach equilibrium, especially in the short run. We downplay this antiequilibrium argument for several reasons. First, our impression is that there are some interesting economic situations in which most of the participants seem to have a pretty good idea of what to expect from day to day, perhaps because the social arrangements and social norms that we observe reflect a process of thousands of years of learning from the experiences of past generations. Second, although there are interesting periods in which social norms change so suddenly that they break down, such as during the transition from a controlled economy to a market one, the dynamic learning models that have been developed so far seem unlikely to provide much insight about the medium-term behavior that will occur in these circumstances.[2] Third, learning theories often have little to say in the short run, making predictions that are highly dependent on details of the learning process and prior beliefs; the long-run predictions are generally more robust to the specification of the model. Finally, from an empirical point of view, it is difficult to gather enough data to test predictions about short-term fluctuations along the adjustment path. For this reason we will focus primarily on the long-run properties of the models we study. Learning theory does, however, make some predictions about rates of convergence and behavior in the medium run, and we will discuss these issues as well.

Even given the restriction to long-run analysis, there is a question of the relative weight to be given to cases where behavior converges and

2. However, Boylan and El-Gamal (1993), Crawford (1995), Roth and Er'ev (1995), Er'ev and Roth (1996), Nagel (1993), and Stahl (1994) use theoretical learning models to try to explain data on short-term and medium-term play in game theory experiments.

cases where it does not. We chose to emphasize the convergence results, in part because they are sharper but also because we feel that these are the cases where the behavior that is specified for the agents is most likely to be a good description of how the agents will actually behave. Our argument here is that the learning models that have been studied so far do not do full justice to the ability of people to recognize patterns of behavior by others. Consequently, when learning models fail to converge, the behavior of the model's individuals is typically quite naive; for example, the players may ignore the fact that the model is locked in to a persistent cycle. We suspect that if the cycles persist long enough, the agents will eventually use more sophisticated inference rules that detect them; for this reason we are not convinced that models of cycles in learning are useful descriptions of actual behavior. However, this does not entirely justify our focus on convergence results: As we discuss in chapter 8, more sophisticated behavior may simply lead to more complicated cycles.

We find it useful to distinguish between two related but different kinds of models that are used to model the processes by which players change the strategies they are using to play a game. In our terminology a "learning model" is any model that specifies the learning rules used by individual players and examines their interaction when the game (or games) is played repeatedly. In particular, while Bayesian learning is certainly a form of learning, and one that we will discuss, learning models can be far less sophisticated and include, for example, stimulus-response models of the type first studied by Bush and Mosteller in the 1950s and more recently taken up by economists.[3] As will become clear in the course of this book, our own views about learning models tend to favor those in which the agents, while not necessarily fully rational, are nevertheless somewhat sophisticated; we will frequently criticize learning models for assuming that agents are more naïve than we feel is plausible.

Individual-level models tend to be mathematically complex, especially in models with a large population of players. Consequently there has also been a great deal of work that makes assumptions directly on the behavior of the aggregate population. The basic assumption here is that some unspecified process at the individual level leads the population as a whole to adopt strategies that yield improved payoffs. The standard practice is to call such models "evolutionary," probably because the first examples of such processes came from the field of evolutionary biology.

3. Examples include Cross (1983), and more recently the Borgers and Sarin (1995), Er'ev and Roth (1996), and Roth and Er'ev (1995) papers discussed in chapter 3.