

Clarendon Lectures in Economics

Andrei Shleifer

INEFFICIENT MARKETS

AN INTRODUCTION
TO BEHAVIORAL
FINANCE



OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford OX2 6DP

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide in

Oxford New York

Athens Auckland Bangkok Bogotá Buenos Aires Calcutta
Cape Town Chennai Dar es Salaam Delhi Florence Hong Kong Istanbul
Karachi Kuala Lumpur Madrid Melbourne Mexico City Mumbai
Nairobi Paris São Paulo Singapore Taipei Tokyo Toronto Warsaw

with associated companies in Berlin Ibadan

Oxford is a registered trade mark of Oxford University Press
in the UK and in certain other countries

Published in the United States
by Oxford University Press Inc., New York

© Andrei Shleifer, 2000

The moral rights of the author have been asserted

Database right Oxford University Press (maker)

First published 2000

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,
without the prior permission in writing of Oxford University Press,
or as expressly permitted by law, or under terms agreed with the appropriate
reprographics rights organization. Enquiries concerning reproduction
outside the scope of the above should be sent to the Rights Department,
Oxford University Press, at the address above

You must not circulate this book in any other binding or cover
and you must impose this same condition on any acquirer

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

Library of Congress Cataloging in Publication Data

Shleifer, Andrei.

Inefficient markets : an introduction to behavioral finance / Andrei Shleifer.
p.cm.— (Clarendon lectures in economics)

Includes bibliographical references and index.

1. Finance. 2. Investments. 3. Stocks—Prices. 4. Efficient market theory. I. Title.
II. Series.

HG4515 .S54 1999 332.6—dc21 99-057647

ISBN 0-19-829228-7 (Hbk)

ISBN 0-19-829227-9 (Pbk)

Typeset by J&L Composition Ltd, Filey, North Yorkshire
Printed in Great Britain
on acid-free paper by Biddles Ltd
Guildford & King's Lynn

Acknowledgments

This book grows out of the Clarendon Lectures given at Oxford in the Spring of 1996. I appreciate the hospitality of Oxford University Press and its economics editor, Andrew Schuller.

I started working on the efficiency of financial markets as a graduate student in the mid-1980s. At the time, there were only a handful of academic papers in behavioral finance, written by people like Robert Shiller, Larry Summers, and Richard Thaler. The assessment by financial economists of this research was not especially generous. Nonetheless, the area seemed to me to be incredibly exciting. During graduate school, I wrote a paper on stock inclusions into the S&P 500 Index and started on a series of theoretical projects with Brad De Long, Larry Summers, and Robert Waldmann. Behavioral finance experienced an upsurge of research in the late 1980s. I had intended to write a book about it while visiting the Russell Sage Foundation in 1992, but that plan did not materialize. In some ways, this is fortunate, since much work has been done in this area since then. I am still very grateful to the Russell Sage Foundation for getting me started.

During the 1990s, I have worked on a number of projects with Josef Lakonishok and Robert Vishny, both of whom have greatly contributed to the material presented in this book. I also learned a lot about arbitrage from Gabe Sunshine and Nancy Zimmerman, who introduced me to some fascinating analytical issues which arise from a more practical understanding of markets. The invitation to give the Clarendon Lectures came at a perfect moment and allowed me to write this book. I am finally finishing it during a sabbatical year at the Massachusetts Institute of Technology, which both kindly hosted me and encouraged me to present some chapters as Independent Activities Period lectures.

Many of the chapters draw heavily on my joint research with a number of colleagues. Chapter 2 draws on 'Noise Trader Risk in Financial Markets,' written with Brad De Long, Lawrence

Summers, and Robert Waldmann and published in the *Journal of Political Economy* in 1990. Chapter 3 is based on 'Investor Sentiment and the Closed End Fund Puzzle,' written with Charles Lee and Richard Thaler and published in the *Journal of Finance* in 1991. Chapter 4 borrows heavily from 'The Limits of Arbitrage,' written with Robert Vishny and published in the *Journal of Finance* in 1997. Chapter 5 follows 'A Model of Investor Sentiment,' written with Nicholas Barberis and Robert Vishny and published in the *Journal of Financial Economics* in 1998. Chapter 6 uses material from 'Positive Feedback Investment Strategies and Destabilizing Rational Speculation,' written with Brad De Long, Lawrence Summers, and Robert Waldmann and published in the *Journal of Finance* in 1990. While I have borrowed liberally from these articles, I have equally liberally added and subtracted material. I am grateful to all my co-authors for years of collaboration.

I was fortunate to receive many extremely perceptive comments on this book from Nicholas Barberis, Olivier Blanchard, John Campbell, Edward Glaeser, Paul Gompers, Oliver Hart, Simon Johnson, Daniel Kahneman, David Laibson, Rafael La Porta, Florencio Lopez-de-Silanes, Andrew Metrick, Sendhil Mullainathan, Michael Rashes, Lawrence Summers, Richard Thaler, Daniel Wolfenzon, and Jeff Wurgler. Their comments had a significant impact on the final draft. Nicholas Barberis and David Laibson were particularly helpful. I am also very grateful to Clare MacLean, who worked tirelessly to get this book out the door, and to Malcolm Baker for preparing the index.

My two greatest debts, however, are to Larry Summers and Nancy Zimmerman. Larry got me started as an economist and introduced me to behavioral finance. He has been a shining example in many ways. Nancy's insight into financial markets is always inspiring. This book is dedicated to her.

Contents

<i>Acknowledgments</i>	<i>vii</i>
1. Are Financial Markets Efficient?	1
2. Noise Trader Risk in Financial Markets	28
3. The Closed End Fund Puzzle	53
4. Professional Arbitrage	89
5. A Model of Investor Sentiment	112
6. Positive Feedback Investment Strategies	154
7. Open Problems	175
<i>Bibliography</i>	198
<i>Index</i>	211

Are Financial Markets Efficient?

THE efficient markets hypothesis (EMH) has been the central proposition of finance for nearly thirty years. In his classic statement of this hypothesis, Fama (1970) defined an efficient financial market as one in which security prices always fully reflect the available information. The efficient markets hypothesis then states that real-world financial markets, such as the U.S. bond or stock market, are actually efficient according to this definition. The power of this statement is dazzling. Perhaps most radically, the EMH 'rules out the possibility of trading systems based only on currently available information that have expected profits or returns in excess of equilibrium expected profit or return' (Fama 1970). In plain English, an average investor—whether an individual, a pension fund, or a mutual fund—cannot hope to consistently beat the market, and the vast resources that such investors dedicate to analyzing, picking, and trading securities are wasted. Better to passively hold the market portfolio, and to forget active money management altogether. If the EMH holds, the market truly knows best.

In the first decade after its conception in the 1960s, the EMH turned into an enormous theoretical and empirical success. Academics developed powerful theoretical reasons why the hypothesis should hold. More impressively, a vast array of empirical findings quickly emerged—nearly all of them supporting the hypothesis. Indeed, the field of academic finance in general, and security analysis in particular, was created on the basis of the EMH and its applications. The University of Chicago, where the EMH was invented, justly became the world's center of academic finance. In 1978, Michael Jensen—a Chicago graduate and one of the creators of the EMH—declared that 'there is no other proposition in economics which has more solid empirical evidence supporting it than the Efficient Markets Hypothesis' (Jensen 1978, p. 95).

Such strong statements portend reversals, and the EMH is no exception. In the last twenty years, both the theoretical foundations

of the EMH and the empirical evidence purporting to support it have been challenged. The key forces by which markets are supposed to attain efficiency, such as arbitrage, are likely to be much weaker and more limited than the efficient markets theorists have supposed. Moreover, new studies of security prices have reversed some of the earlier evidence favoring the EMH. With the new theory and evidence, behavioral finance has emerged as an alternative view of financial markets. In this view, economic theory does not lead us to expect financial markets to be efficient. Rather, systematic and significant deviations from efficiency are expected to persist for long periods of time. Empirically, behavioral finance both explains the evidence that appears anomalous from the efficient markets perspective, and generates new predictions that have been confirmed in the data.

This book introduces the research in behavioral finance. In this opening chapter, we describe both the theoretical and the empirical foundations of the EMH, as well as some of the cracks that have emerged in these foundations.

The theoretical foundations of the EMH

The basic theoretical case for the EMH rests on three arguments which rely on progressively weaker assumptions. First, investors are assumed to be rational and hence to value securities rationally. Second, to the extent that some investors are not rational, their trades are random and therefore cancel each other out without affecting prices. Third, to the extent that investors are irrational in similar ways, they are met in the market by rational arbitrageurs who eliminate their influence on prices.

When investors are rational, they value each security for its fundamental value: the net present value of its future cash flows, discounted using their risk characteristics. When investors learn something about fundamental values of securities, they quickly respond to the new information by bidding up prices when the news is good and bidding them down when the news is bad. As a consequence, security prices incorporate all the available information almost immediately and prices adjust to new levels corresponding to the new net present values of cash flows. Samuelson (1965) and Mandelbrot (1966) proved some of the first theorems showing how, in competitive markets with rational risk-neutral

investors, returns are unpredictable—security values and prices follow random walks. Since then, economists have characterized efficient securities prices for risk-averse investors, with both varying levels of risk over time and varying tolerances toward risk. In these more complicated models security prices are no longer predicted to follow random walks. Still, investor rationality implies the impossibility of earning superior risk-adjusted returns, just as Fama wrote in 1970. The EMH is thus first and foremost a consequence of equilibrium in competitive markets with fully rational investors.

But remarkably, the EMH does not live or die by investor rationality. In many scenarios where some investors are not fully rational, markets are still predicted to be efficient. In one commonly discussed case, the irrational investors in the market trade randomly. When there are large numbers of such investors, and when their trading strategies are uncorrelated, their trades are likely to cancel each other out. In such a market, there will be substantial trading volume as the irrational investors exchange securities with each other, but the prices are nonetheless close to fundamental values. This argument relies crucially on the lack of correlation in the strategies of the irrational investors, and, for that reason, is quite limited. The case for the EMH, however, can be made even in situations where the trading strategies of investors are correlated.

This case, as made by Milton Friedman (1953) and Fama (1965), is based on arbitrage. It is one of the most intuitively appealing and plausible arguments in all of economics. A textbook definition (Sharpe and Alexander 1990) defines arbitrage as 'the simultaneous purchase and sale of the same, or essentially similar, security in two different markets at advantageously different prices.' Suppose that some security, say a stock, becomes overpriced in a market relative to its fundamental value as a result of correlated purchases by unsophisticated, or irrational, investors. This security now represents a bad buy, since its price exceeds the properly risk adjusted net present value of its cash flows or dividends. Noting this overpricing, smart investors, or arbitrageurs, would sell or even ~~sell short~~ this expensive security and simultaneously purchase other, 'essentially similar,' securities to hedge their risks. If such substitute securities are available and arbitrageurs are able to trade them, they can earn a profit, since

they are short expensive securities and long the same, or very similar, but cheaper securities. The effect of this selling by arbitrageurs is to bring the price of the overpriced security down to its fundamental value. In fact, if arbitrage is quick and effective enough because substitute securities are readily available and the arbitrageurs are competing with each other to earn profits, the price of a security can never get far away from its fundamental value, and indeed arbitrageurs themselves are unable to earn much of an abnormal return. A similar argument applies to an undervalued security. To earn a profit, arbitrageurs would buy underpriced securities and sell short essentially similar securities to hedge their risk, thereby preventing the underpricing from being either substantial or very long-lasting. The process of arbitrage brings security prices in line with their fundamental values even when some investors are not fully rational and their demands are correlated, as long as securities have close substitutes.

Arbitrage has a further implication. To the extent that the securities that the irrational investors are buying are overpriced and the securities they are getting rid of are underpriced, such investors earn lower returns than either passive investors or arbitrageurs. Relative to their peers, irrational investors lose money. As Friedman (1953) points out, they cannot lose money forever: they must become much less wealthy and eventually disappear from the market. If arbitrage does not eliminate their influence on asset prices instantaneously, market forces eliminate their wealth. In the long run, market efficiency prevails because of competitive selection and arbitrage.

It is difficult not to be impressed with the full range and power of the theoretical arguments for efficient markets. When people are rational, markets are efficient by definition. When some people are irrational, much or all of their trading is with each other, and hence has only a limited influence on prices even without countervailing trading by the rational investors. But such countervailing trading does exist and works to bring prices closer to fundamental values. Competition between arbitrageurs for superior returns ensures that the adjustment of prices to fundamental values is very quick. Finally, to the extent that the irrational investors do manage to transact at prices that are different from fundamental values, they only hurt themselves and bring

about their own demise. Not only investor rationality, but market forces themselves bring about the efficiency of financial markets.

The empirical foundations of the EMH

Strong as the theoretical case for the EMH may seem, the empirical evidence that appeared in the 1960s and the 1970s was even more overwhelming. At the most general level, the empirical predictions of the EMH can be divided into two broad categories. First, when news about the value of a security hits the market, its price should react and incorporate this news both quickly and correctly. The 'quickly' part means that those who receive the news late—for instance by reading it in the newspapers or in company reports—should not be able to profit from this information. The 'correctly' part means that the price adjustment in response to the news should be accurate on average: the prices should neither underreact nor overreact to particular news announcements. There should be neither price trends nor price reversals following the initial impact of the news. Second, since a security's price must be equal to its value, prices should not move without any news about the value of the security. That is, prices should not react to changes in demand or supply of a security that are not accompanied by news about its fundamental value. The quick and accurate reaction of security prices to information, as well as the non-reaction to non-information, are the two broad predictions of the efficient markets hypothesis.

The principal hypothesis following from quick and accurate reaction of prices to new information is that stale information is of no value in making money, as Fama (1970) points out. To evaluate this hypothesis empirically, researchers needed to define 'stale information' and 'making money.' The first definition turns out to be relatively straightforward. Defining 'making money,' in contrast, is enormously controversial. The reason is that 'making money' in finance means making a superior return after an adjustment for risk. Showing that a particular strategy based on exploiting stale information on average earns a positive cash flow over some period of time is not, therefore, by itself evidence of market inefficiency. To earn this profit, an investor may have to bear risk and his profit may just be a fair market compensation for risk-bearing. The trouble is that measuring the risk of a particular

investment strategy is both difficult and controversial, and requires a model of the fair relationship between risk and return. One widely-accepted model is the Capital Asset Pricing Model (Sharpe 1964), but it is not the only possibility. The dependence of most tests of market efficiency on a model of risk and expected return is Fama's (1970) deepest insight, which has pervaded the debates in empirical finance ever since. Whenever researchers have found a money-making opportunity resulting from trading on stale information, critics have been quick to suggest a model of risk—convincing or otherwise—that would reduce these profits to a fair compensation for risk-taking.

The definition of stale information is far less controversial. Fama distinguishes between three types of stale information, giving rise to three forms of the EMH. For the so-called *weak form* efficiency, the relevant stale information is past prices and returns. The weak form EMH posits that it is impossible to earn superior risk-adjusted profits based on the knowledge of past prices and returns. Under the assumption of risk neutrality, this version of the EMH reduces to the random walk hypothesis, the statement that stock returns are entirely unpredictable based on past returns (Fama 1965).

Past returns are not the only stale information that investors have. The *semi-strong form* of the EMH states that investors cannot earn superior risk-adjusted returns using *any* publicly available information. Put differently, as soon as information becomes public, it is immediately incorporated into prices, and hence an investor cannot gain by using this information to predict returns. A semi-strong form efficient market is obviously weak form efficient as well, since past prices and returns are a proper subset of the publicly available information about a security.

It is still possible that while an investor cannot profit from trading on publicly available information, he can still earn abnormal risk-adjusted profits by trading on information that is not yet known to market participants, sometimes described as inside information. The *strong form* of the EMH states that even these profits are impossible because the insiders' information quickly leaks out and is incorporated into prices. To be fair, most evaluations of the EMH have focused on weak and semi-strong form efficiency, and have not taken the extreme position that there is no such thing as profitable insider trading, as would be required if

the strong form EMH were to hold. Indeed, the insider traders occupying minimum security prisons for making illegal profits themselves represent some evidence against the strong form EMH. But there is more systematic evidence as well that insiders earn some abnormal returns even when they trade completely legally (Seyhun 1998, Jeng *et al.* 1999).

When economists have set out to test these predictions, their evidence was broadly supportive of the EMH. With respect to weak form efficiency, Fama (1965) finds that stock prices indeed approximately follow random walks. He finds no systematic evidence of profitability of 'technical' trading strategies, such as buying stocks when their prices just went up or selling them when their prices just went down. On a given day, the price of a stock is as likely to rise after a previous day's increase as after a previous day's decline. Early tests of more complicated trading rules have yielded similar failures to earn profits on average by predicting returns based on past returns, consistent with the weak form EMH.

The initial tests corroborated semi-strong form efficiency as well. One testing strategy is to look at particular news events pertaining to individual companies and to ask whether prices adjusted to this news immediately or over a period of a few days. These so-called *event studies*, pioneered by Fama *et al.* (1969), became the principal methodology of empirical finance, as multitudes of important corporate news events, such as earnings and dividend announcements, takeovers and divestitures, share issues and repurchases, changes in management compensation, and so on, came to be evaluated empirically through the effects of these news events on share prices. As an illustration, consider the study by Keown and Pinkerton (1981) of returns to targets of takeover bids around the announcement of the bid. The results for returns of an average target, adjusted for market movements, are reproduced in Figure 1.1. They show that share prices of targets begin to rise prior to the announcement of the bid as the news of a possible bid is incorporated into prices, and then jump on the date of the public announcement to reflect the takeover premium offered to target firm shareholders. But the jump in share prices on the announcement is not followed by a continued trend up or a reversal down, indicating that prices of takeover targets adjust




FIG. 1.1 Cumulative abnormal returns to shareholders of targets of takeover attempts around the announcement date.

Source: Keown and Pinkerton (1981).

to the public news of the bid instantaneously, consistent with the semi-strong form EMH.

Taken together, the early evidence on weak and semi-strong form market efficiency was almost entirely supportive. The same was also the case with the other implication of the EMH, namely that prices do not react to non-information. In the principal early empirical study of this proposition, Scholes (1972) uses the event study methodology to evaluate share price reactions to sales of large blocks of shares in individual companies by substantial stockholders. Scholes's work is particularly important because it deals directly with the issue central to the arbitrage arguments in the efficient markets theory: the availability of close substitutes for individual securities. An exact substitute for a given security is another security (or portfolio of securities) with identical cash flows in all states of the world. A close substitute is a security (or portfolio) with very similar cash flows in all states of the world, and therefore with similar risk characteristics to those of a given

security. The availability of close substitutes for a given security, closely related to the assumption of complete markets, is essential for the working of arbitrage because it allows for more than one way to get a given pattern of cash flows in different states of the world.

Scholes reasons that when arbitrage is needed to make markets efficient, individual stocks must have close substitutes for such arbitrage to work well. When close substitutes are available, arbitrageurs can sell expensive securities and buy cheap close substitutes, thereby equalizing their relative prices and bringing markets to efficiency. When close substitutes are not available, arbitrageurs cannot make such trades. When stocks do have close substitutes, investors are indifferent as to which stocks with a given set of risk characteristics to hold. Sales of large blocks of shares—particularly if made by uninformed sellers—should not have a material impact on the stock price because that price is determined by the stock's value relative to that of its close substitutes rather than by supply. Scholes's argument, which he calls the substitution hypothesis, is virtually identical to the arbitrage argument for market efficiency. When a seller unloads a block of shares on the market, other investors would gladly increase their holdings of that stock a bit in exchange for only a trivial, if any, price concession, and perhaps reduce their holdings of close substitutes to keep the risk of their portfolio constant. Competition between these potential buyers assures that the price concession that an uninformed seller must give them is small.

In his study, Scholes finds relatively small share price reactions to block sales. He accounts for these reactions by the possible, but small, adverse news revealed by the decision of large blockholders to sell their shares. On Scholes's interpretation, this result is consistent with the substitution hypothesis, and therefore with the second crucial prediction of the EMH, namely that stock prices do not react to non-information. More importantly, the non-reaction of prices to non-information points to arbitrage in action: the willingness of investors to adjust their portfolios to absorb more shares without a large influence on the price.

As matters stood at the end of the 1970s, the EMH was indeed one of the great triumphs of twentieth-century economics. Standard economic theory—particularly the theory of arbitrage—predicted that financial markets were efficient. Mountains of

empirical evidence based on some of the most extensive data available in economics, that on security prices, almost universally confirmed the predictions of the theory. Whenever researchers found small money-making opportunities, they could be easily explained away by a variety of arguments, the most pervasive of which was the failure to adjust properly for risk. At the time, Jensen's claim about the best established fact in economics was not all that outrageous.

Theoretical challenges to the EMH

Shortly after Jensen's pronouncement, the EMH was challenged on both theoretical and empirical grounds. Although the initial challenges were primarily empirical, it is easier to begin by reviewing some potential difficulties with the theoretical case for the EMH and then turn to the evidence. This chapter only highlights the principal challenges. They are developed in the rest of the book.

To begin, it is difficult to sustain the case that people in general, and investors in particular, are fully rational. At the superficial level, many investors react to irrelevant information in forming their demand for securities; as Fischer Black (1986) put it, they trade on noise rather than information. Investors follow the advice of financial gurus, fail to diversify, actively trade stocks and churn their portfolios, sell winning stocks and hold on to losing stocks thereby increasing their tax liabilities, buy and sell actively and expensively managed mutual funds, follow stock price patterns and other popular models. In short, investors hardly pursue the passive strategies expected of uninformed market participants by the efficient markets theory.

This evidence of what investors actually do is only the tip of the iceberg. Investors' deviations from the maxims of economic rationality turn out to be highly pervasive and systematic. As summarized by Kahneman and Riepe (1998), people deviate from the standard decision making model in a number of fundamental areas. We can group these areas, somewhat simplistically, into three broad categories: attitudes toward risk, non-Bayesian expectation formation, and sensitivity of decision making to the framing of problems.

First, individuals do not assess risky gambles following the

precepts of von Neumann–Morgenstern rationality. Rather, in assessing such gambles, people look not at the levels of final wealth they can attain but at gains and losses relative to some reference point, which may vary from situation to situation, and display loss aversion—a loss function that is steeper than a gain function. Such preferences—first described and modeled by Kahneman and Tversky (1979) in their ‘Prospect Theory’—are helpful for thinking about a number of problems in finance. One of them is the notorious reluctance of investors to sell stocks that lose value, which comes out of loss aversion (Odean 1998). Another is investors’ aversion to holding stocks more generally, known as the equity premium puzzle (Mehra and Prescott 1985, Benartzi and Thaler 1995).

Second, individuals systematically violate Bayes rule and other maxims of probability theory in their predictions of uncertain outcomes (Kahneman and Tversky 1973). For example, people often predict future uncertain events by taking a short history of data and asking what broader picture this history is representative of. In focusing on such representativeness, they often do not pay enough attention to the possibility that the recent history is generated by chance rather than by the ‘model’ they are constructing. Such heuristics are useful in many life situations—they help people to identify patterns in the data as well as to save on computation—but they may lead investors seriously astray. For example, investors may extrapolate short past histories of rapid earnings growth of some companies too far into the future and therefore overprice these glamorous companies without a recognition that, statistically speaking, trees do not grow to the sky. Such overreaction lowers future returns as past growth rates fail to repeat themselves and prices adjust to more plausible valuations.

Perhaps most radically, individuals make different choices depending on how a given problem is presented to them, so that framing influences decisions. In choosing investments, for example, investors allocate more of their wealth to stocks rather than bonds when they see a very impressive history of long-term stock returns relative to those on bonds, than if they only see the volatile short-term stock returns (Benartzi and Thaler 1995).

A number of terms have been used to describe investors whose preferences and beliefs conform to the psychological evidence rather than the normative economic model. Beliefs based on

heuristics rather than Bayesian rationality are sometimes called 'investor sentiment.' Less kindly, the investors whose conduct is not rational according to the normative model are described as 'unsophisticated' or, following Kyle (1985) and Black (1986), as 'noise traders.'

If the theory of efficient markets relied entirely on the rationality of individual investors, then the psychological evidence would by itself present an extremely serious, perhaps fatal, problem for the theory. But of course it does not. Recall that the second line of defense of the efficient markets theory is that the irrational investors, while they may exist, trade randomly, and hence their trades cancel each other out. It is this argument that the Kahneman and Tversky theories dispose of entirely. The psychological evidence shows precisely that people do not deviate from rationality randomly, but rather most deviate in the same way. To the extent that unsophisticated investors form their demands for securities based on their own beliefs, buying and selling would be highly correlated across investors. Investors would not trade randomly with each other, but rather many of them would try to buy the same securities or to sell the same securities at roughly the same time. This problem only becomes more severe when the noise traders behave socially and follow each others' mistakes by listening to rumors or imitating their neighbors (Shiller 1984). Investor sentiment reflects the common judgment errors made by a substantial number of investors, rather than uncorrelated random mistakes.

Individuals are not the only investors whose trading strategies are difficult to reconcile with rationality. Much of the money in financial markets is allocated by professional managers of pension and mutual funds on behalf of individual investors and corporations. Professional money managers are of course themselves people, and as such are subject to the same biases as individual investors. But they are also agents who manage other people's money, and this delegation introduces further distortions into their decisions relative to what fully-informed sponsors might wish (Lakonishok *et al.* 1992). For example, professional managers may choose portfolios that are excessively close to the benchmark that they are evaluated against, such as the S&P 500 Index, so as to minimize the risk of underperforming this benchmark. They may also herd and select stocks that other managers select, again to avoid falling behind and looking bad (Scharfstein