



MORALITY, COMPETITION, AND THE FIRM

**THE MARKET FAILURES APPROACH
TO BUSINESS ETHICS**

JOSEPH HEATH

Morality, Competition, and the Firm

Morality, Competition, and the Firm

THE MARKET FAILURES APPROACH TO
BUSINESS ETHICS

Joseph Heath

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

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.

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

Oxford is a registered trademark of Oxford University Press
in the UK and certain other countries.

Published in the United States of America by

Oxford University Press

198 Madison Avenue, New York, NY 10016

© Oxford University Press 2014

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,
by license, or under terms agreed with the appropriate reproduction rights organization.
Inquiries 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 work in any other form
and you must impose this same condition on any acquirer.

Library of Congress Cataloging-in-Publication Data

Heath, Joseph, 1967–

Morality, competition, and the firm : the market failures approach to business ethics / Joseph Heath.
pages cm

ISBN 978-0-19-999048-1 (hardback)

1. Business ethics. 2. Profit—Moral and ethical aspects. 3. Competition.

4. Corporations—Moral and ethical aspects. I. Title.

HF5387.H435 2014

174'.4—dc23

2013050434

9 8 7 6 5 4 3 2 1

Printed in the United States of America

on acid-free paper

CONTENTS

Acknowledgments vii

Introduction 1

PART I } **The Corporation and Society**

1. A Market Failures Approach to Business Ethics 25
2. Stakeholder Theory, Corporate Governance, and Public Management
(with Wayne Norman) 42
3. Business Ethics without Stakeholders 68
4. An Adversarial Ethic for Business: or, When Sun-Tzu Met the
Stakeholder 93
5. Business Ethics and the “End of History” in Corporate Law 116

PART II } **Cooperation and the Market**

6. Contractualism: Micro and Macro 145
7. Efficiency as the Implicit Morality of the Market 173
8. The History of the Invisible Hand 205
9. The Benefits of Cooperation 230

PART III } **Extending the Framework**

10. The Uses and Abuses of Agency Theory 263
11. Business Ethics and Moral Motivation: A Criminological
Perspective 294
12. Business Ethics after Virtue 322
13. Reasonable Restrictions on Underwriting 345

Bibliography 373

Index 395

ACKNOWLEDGMENTS

I owe many thanks to the various people who have played an important role in the development of this project. First of all, the book would never have seen the light of day without the interest, encouragement, support, and oversight of Peter Ohlin at Oxford University Press in New York. Academically, I owe the greatest debt to Wayne Norman, first for convincing me that this was an exciting and worthwhile field of study, and second for serving as a collaborator, interlocutor, and source of unflagging enthusiasm over the years. Philosophers are a somewhat truculent lot, and so I have never before had the experience of seeing exactly eye-to-eye with anyone on some complex set of philosophical issues. My exchanges with Wayne over the years have provided welcome relief from the usual solitude of philosophical inquiry.

I have learned an enormous amount from John Boatright, who always manages to live up to his reputation as “the smartest guy in the room” (in the non-pejorative sense of the term). My work in this area has also benefited a great deal from conversations over the years with Sareh Pouryousefi, Dominic Martin, Idil Boran, Pierre-Yves Néron, Chris MacDonald, and Thorsten Busch.

I have benefited from the financial support of the Social Sciences and Humanities Research Council of Canada as well as the Pierre Elliot Trudeau Foundation.

Some of these papers have been published previously. I have not made any substantive changes, other than correcting the occasional error, switching all spelling to the American, as well as changing the format of the references in some papers to create a consistent style throughout the book. Special thanks to Jessica Brown for assistance with this task. I was not able to eliminate the occasional repetitions that occur, due to the initially independent publication of the papers, and so the most I can do is apologize and ask for a measure of forbearance on the part of the reader. I would like to thank both the editors and referees of the journals for the work that they put into the improvement of these papers, as well as to the publishers for the permission to reproduce them here.

Chapter 1. “A Market Failures Approach to Business Ethics,” in Bernard Hodgson, ed., *Studies in Economic Ethics and Philosophy*, Vol. 9 (Berlin: Springer, 2004).

Chapter 2. “Stakeholder Theory, Corporate Governance and Public Management,” with Wayne Norman, *Journal of Business Ethics*, 53 (2004): 247–265.

Chapter 3. “Business Ethics without Stakeholders,” *Business Ethics Quarterly*, 16 (2006): 533–557. I would like to thank Wayne Norman and Alexei Marcoux for their input and advice with the writing of this paper.

Chapter 4. “An Adversarial Ethic for Business: or, When Sun-Tzu met the Stakeholder,” *Journal of Business Ethics*, 69 (2006): 359–374. Thanks to the Social Sciences and Humanities Research Council of Canada for financial support in the development of this project.

Chapter 5. “Business Ethics and the ‘End of History’ in Corporate Law,” *Journal of Business Ethics*, 102 (2011): 5–20. The idea for this paper arose during a discussion with John Boatright. It benefited from subsequent discussion with Margaret Blair, Wayne Norman, Alexei Marcoux, Eric Orts, Alan Strudler, members of the Legal Studies and Business Ethics program at the Wharton School of Business, as well as participants in the Trans-Atlantic Business Ethics Conference held at York University.

Chapter 6. “Contractualism: Micro and Macro.” This paper was written for the New York University Colloquium in Legal, Social and Political Philosophy, convened by Thomas Nagel and Ronald Dworkin. I benefited from many of the comments made by participants, but especially the careful attention given to it by Professors Nagel and Dworkin. Thanks to Shlomi Segall for comments on the manuscript as well.

Chapter 7. “Efficiency as the Implicit Morality of the Market.” This paper revises some material previously published as “Ideal Theory in an n -th Best World: The Case of Pauper Labor,” *Journal of Global Ethics*, 9 (2013): 159–172. My thanks to Kathryn Walker, Monique Deveaux, and participants in the “Thinking Beyond Distribution” workshop at the University of Toronto Centre for Ethics for valuable comments. This paper was also presented at the *Iustitia Amplificata* conference on “Relational Injustice” at Goethe University; my thanks to James Gledhill in particular for the invitation to participate. Thanks to Dominic Martin for comments on the manuscript as well.

Chapter 8. “The History of the Invisible Hand.” This paper was presented at the summer school of the Netherlands School for Research in Practical Philosophy, organized by Ingrid Robeyns and Rutger Claassen on the theme of “Ethics and Economics.” My thanks to the participants for helpful feedback, as well as to Manuel Wörsdörfer for comments that led to a number of specific improvements.

Chapter 9. “The Benefits of Cooperation,” *Philosophy and Public Affairs*, 34 (2006): 313–351. I would like to thank Christopher Morris, Thomas Hurka, and Edward McClennen for encouragement at key points in the development of this article. Thanks also to audiences at the Canada School of Public Service, the University of Guelph, and the Université de Montréal, along with the Editors of *Philosophy & Public Affairs*, for helpful comments.

Chapter 10. “The Uses and Abuses of Agency Theory,” *Business Ethics Quarterly*, 19 (2009): 497–528.

Chapter 11. “Business Ethics and Moral Motivation: A Criminological Perspective,” *Journal of Business Ethics*, 83 (2008): 595–614.

Chapter 12. “Business Ethics after Virtue.” My thanks to Christine Tappolet for comments and advice on the improvement of this paper. I have also drawn some material from a paper co-authored with Wayne Norman and Jeffrey Moriarty, “Business Ethics and (or as) Political Philosophy,” *Business Ethics Quarterly*, 10 (2010): 427–452.

Chapter 13. “Reasonable Restrictions on Underwriting,” in Patrick Flanagan, Patrick Primeaux, and William Ferguson, eds., *Research in Ethical Issues in Organizations*, Vol. 7 (Amsterdam: Elsevier, 2006).

Introduction

This volume brings together a series of papers that I have written over the past ten years on the subject of business ethics, along with a few more general pieces that articulate the normative foundations of the project. Together they provide a basic outline of what I refer to as the “market failures” approach to business ethics. It was not my original intention to make a contribution to this particular literature. My objective in this introduction will therefore be to sketch out a bit of the intellectual history that led to the development of the project and to say something about how the various pieces fit together. Along the way, I hope to make clear some of the more general political motivations that inform the market failures approach, since these have occasionally been the subject of misunderstanding.

I.1. The Intellectual History of the Project

Traditional philosophical business ethics is done within an “applied ethics” paradigm, in which the theorist begins by establishing a commitment to a particular ethical theory, such as Kantianism or virtue theory, then (typically) goes on to discuss some “moral dilemma” that might arise in a business context. This leaves business ethics firmly anchored within a framework of personal ethics, or what Robert Solomon (1992) refers to as the “micro” level of institutional analysis. My own approach, by contrast, arose as something of a byproduct of a larger (or “macro”) project in political economy. The best way to understand it, I believe, is to see it in the context of this more general position.

My initial interest was in the role that the state plays in a modern capitalist economy. This led me (*inter alia*) to write a popular book defending certain features of the Canadian welfare state against the ever-present pull of the American model (*The Efficient Society*, published by Penguin in 2001). One of my central preoccupations in that book was public health care, since at the

time of writing the deficiencies of the American system were not as obvious as they are now, while the Canadian single-payer system was under significant strain, thanks to both cost escalation and government budget constraints. What I found frustrating about the public debate in Canada at the time was that when partisans of the public health care system were called upon to defend it, they inevitably appealed to its egalitarian qualities—the fact that it guaranteed roughly equal care to all citizens. While not having any particular objections to this argument, it also seemed to me that it should not be the first line of defense. Equality is important, but promoting equality is not the only thing that the state does. The most obvious argument in favor of single-payer health systems is that they are more efficient—the Canadian system, for instance, delivers approximately the same volume of health care services and achieves similar health outcomes as the American, but at something approaching half the cost. Furthermore, the basic argument in defense of single-payer appeals to efficiency in the Pareto sense, namely, that by resolving a market failure in the health insurance sector, it corrects a mispricing of insurance that lowers the welfare of everyone in the society.

Of course, there are also egalitarian arguments in favor of public provision (although these turn out to be far more complicated than they may at first appear). But it seemed to me that if there are both efficiency arguments and equality arguments to be made for a particular public program, it is always better to lead with the efficiency arguments, simply because they are inherently less controversial. Equality arguments are essentially about how to resolve distributive conflict, and so always have a win-lose structure. This means that regardless of how compelling they are, there will always be a constituency with an interest in opposing them. Efficiency arguments, on the other hand, appeal to mutual benefit, or win-win transformations, and so do not necessarily create an oppositional constituency (see chapter 6, figure 6.2). Solving collective action problems is not something that anyone should have a stake in preventing.

In developing this intuition, however, I came around to the view that efficiency arguments were not only rhetorically more effective, but actually did a better job of articulating the underlying logic of many welfare-state programs. While the tax system no doubt institutionalizes a set of egalitarian commitments, the social programs that the welfare state provides are primarily driven by efficiency concerns. This is often obscured by the fact that many of the “goods” that the welfare state provides are insurance products and are therefore superficially redistributive. A public pension scheme, for instance, looks like a system of redistribution—since it takes money from one group and transfers it to another—but is in fact a collective retirement insurance scheme. A better way to think of it is as a bundle of collectively purchased life annuities, delivered through the public sector because private insurance fails to price these products at an appropriate level (see chapter 9). Like all

insurance arrangements, it is designed to redistribute from the lucky to the unlucky, not from the rich to the poor. Public health insurance, of course, has the same structure (chapter 13).

This insight led me to regard efficiency as one of the underappreciated virtues of the welfare state—and therefore to take more seriously the traditional “public economics” view that the welfare state is primarily in the business of correcting market failure, not achieving distributive justice, even in many cases where it does not look like it is promoting efficiency (see Heath 2011). This in turn suggested that one could make the case for many of the core features of the welfare state (specifically, public education, health care, and pensions) without appealing to controversial egalitarian commitments (Moss 2004). The only concept required is that of market failure, followed by some account of the resources that the state is able to deploy in order to resolve particular forms of such failure.

Once I had the position formulated in this way, it became clear to me that the conceptual resources required to mount a defense of the welfare state were in essence no different from the ones that underlay a standard transaction cost account of the firm. Transaction cost theory takes as its point of departure the observation that we have a toolkit of different institutional forms that we can use to organize economic cooperation, with the two most important being markets and administrative hierarchies. Depending on the nature of the “transaction” in question, each different institutional form will have a different profile of costs. These costs must be understood broadly, to include not only direct monetary costs (e.g., hiring a lawyer to draw up a contract), indirect costs (e.g., losses due to inadequate enforcement of contracts), but also purely invisible costs (e.g., deadweight losses, due to potentially advantageous exchanges that do not occur because of fear of fraud).

The crux of the theory is that neither markets nor hierarchies dominate the other as an organizational form, and so which one will impose greater costs depends upon the nature of the transaction in question. As a result, what tends to arise in a market economy—when the appropriate legal devices are made available—is a mixture of organizational forms, with some production being organized in a decentralized fashion among various individuals or firms using the market to coordinate their relations, and other production occurring “in house,” within the administrative hierarchy of the firm and subject to the authority of management. Ronald Coase’s (1937) great insight was that—granted certain idealizing assumptions—the boundary of the firm will be determined by the relative cost of organizing production using these different governance structures. The boundary will also be quite dynamic. Mergers and acquisitions are processes through which transactions that were once mediated by the market are brought within the scope of managerial authority, while outsourcing is a process through which an administered transaction is dissolved and replaced by market contracting.

Now a “market failure” is defined simply as a circumstance in which markets fail to achieve a Pareto optimum (which is to say, where they leave room for improving at least one person’s position without worsening anyone else’s [Bator 1958]). So what the transaction cost theory of the firm claims is that the organization of economic activity through administrative hierarchies is everywhere explained by the existence of market failure. Less intuitively, it also implies that the organization of economic activity through markets is everywhere explained by “administrative failure.” All of this is to say that neither the concept of market failure nor that of administrative failure does much work, taken alone, what matters is simply the relative cost profile of these different modes of economic organization.

This is an extremely powerful theory, and a far more subtle one than most people realize when they first encounter it. It also has both explanatory and ideological virtues. Most importantly, it is very useful when it comes to helping people to unlearn some of the more extreme or Panglossian views of the market that are often inculcated through early exposure to the introductory economics curriculum. It is important to point out, for instance, that the invisible hand of the market, despite its many virtues, is not magical, and it solves no more than a fraction of our economic problems. If the market actually produced perfectly efficient outcomes, *then there would be no need for corporations*. And yet corporations exist. Therefore, there must be non-trivial limitations on the efficiency properties of markets.

Once all of this has been established, it is fairly easy to make the further point that the corporation possesses certain inherent limitations when it comes to its ability to correct market failure, whereas the state has two qualities that give it a different cost profile when organizing economic cooperation, namely, that within its territory membership is universal and it exercises a monopoly over the powers of compulsion (Stiglitz et al. 1989). For certain transactions, this can result in lower costs. For instance, the state has the power to control adverse selection in a way that no private insurer does, which results in the state being the lowest-cost provider of a variety of different forms of insurance (including, typically, health insurance). This means that a straightforward transaction cost analysis is going to suggest that certain goods and services should be funded by taxation and provided through the public sector.

This analysis provided the basic rhetorical strategy of *The Efficient Society*. The objective was to show that once one accepts both the need for and the legitimacy of corporations, then one cannot but accept both the need for and the legitimacy of the modern welfare state. The rhetorical aspect of the argument is important because at the time that I presented it I did not actually know very much about the theory of the firm or transaction cost theory. Like many philosophers, I had spent a lot of time reading “macro” arguments about capitalism, communism, and the state, but had spent very little time studying the “meso” level—the medium-sized institutions that do most of the organizational

work in the operations of a capitalist economy, such as the corporation in its myriad forms, stock exchanges, financial institutions, insurance companies, government agencies, and so on. For example, one can read both John Rawls's *A Theory of Justice* and Robert Nozick's *Anarchy, State and Utopia* and learn absolutely nothing about any of these institutions. Nozick, despite providing the most sophisticated libertarian defense of the market economy, has essentially nothing to say about the corporation (and in fact mentions it only twice over the course of his book). Rawls writes in considerable detail about various institutional features of the state, but again has nothing to say about the corporation. (This is because he does not consider it part of the basic structure, and so treats it as outside the scope of his principles of justice [Rawls 1999: 126].)

Perhaps because of this shortage of literature, the fact that *The Efficient Society* contained a chapter on the subject of the corporation (and was written in a popular, accessible style), led to its widespread adoption in undergraduate business ethics courses. In particular, Richard Wellen at York University began teaching the book in his very large “business and society” course, and subsequently invited me to meet with a group of enthusiastic students to discuss my views—such as they were—on the corporation. At around the same time, the late Bernard Hodgson invited me to a conference at Trent University to discuss “the invisible hand and the common good.” Together, these led me to develop a stand-alone version of my argument about the firm and its normative implications—the paper that became “A Market Failures Approach to Business Ethics” and forms the first chapter of this book.

One of the things about this paper that will undoubtedly strike many readers is that, apart from being rather polemical in tone, it contains practically no references to the business ethics literature. Indeed, the only two articles I cite are Milton Friedman's famous piece (1970), as well as Andrew Stark's “What's the Matter with Business Ethics?” (1993). The reason is that I had, at the time, essentially no knowledge of business ethics (I had read Stark's paper only because he is a colleague of mine at the University of Toronto). The position that I presented was what I took to be a fairly straightforward implication of the “political economy” perspective that I had been developing. I called it a “market failures” approach to business ethics, although strictly speaking it would be more accurate to have called it a “Paretian” approach, since my major claim is that the market is essentially a staged competition, designed to promote Pareto efficiency, and in cases where the explicit rules governing the competition are insufficient to secure the class of favored outcomes, economic actors should respect the spirit of these rules and refrain from pursuing strategies that run contrary to the point of the competition.

This is a very natural position to take if one approaches the basic question of corporate social responsibility from the perspective of modern economics (understood broadly, to include both the transaction cost theory of the firm and the standard “public economics” understanding of the welfare state; e.g.,

Barr 1998). In particular, in the wake of the “socialist calculation” debate of the early twentieth century, as well as Friedrich Hayek’s information-theoretic reformulation of the classic “invisible hand” argument for the market (discussed in chapter 8), it has become common to regard marketplace competition as essentially a system designed to generate a set of prices, which can in turn be used to achieve a more efficient allocation of productive resources and consumer goods. The major thrust of regulatory interventions in the market—most obviously in the case of environmental and consumer protection—is to correct imperfections that are distorting price signals and thereby leading to the misallocation of resources (such as overproduction of environmental “bads”). It is quite natural—it seems to me—to think that the basic thrust of “business ethics” is the same as that of these regulatory interventions, namely, to discourage firms from taking advantage of market imperfections, even in cases where legal regulation is not feasible (see chapter 1, figure 1.1).

I was therefore unsurprised to discover, after having articulated the “market failures” perspective, that essentially the same thought had occurred to other people. Kenneth Arrow (1973), for example, had expressed very similar views. And I have since discovered a strong current of similar thinking about business ethics in the *Wirtschaftsethik* tradition in Germany, most obviously in the work of Peter Koslowski (2001: 26–30) and Karl Homann (1993). Again, I find all of this unsurprising, simply because the “market failures” perspective is a very natural consequence of taking a broadly liberal theory of justice (such as the “minimally controversial contractualism” that I articulate in chapter 6) and adopting a fairly standard economic perspective on the market.

What was surprising, at least to me, was the resistance that I encountered when I first presented my ideas to business ethicists. During several of my early presentations I discovered that a whole series of steps that I took to be self-evident were in fact highly controversial. For example, I was extremely surprised to find business ethicists resisting the suggestion that markets are competitive, or that competition is somehow important to their function. Of course, they were not exactly denying it, they were mainly resisting it, based largely on the intuition that if we want businesspeople to behave themselves, putting too much emphasis on the competitiveness of market interactions is counterproductive. So, for example, while it seemed natural to me to draw comparisons between business ethics and the ethics of sport and games—given that they are all competitively organized domains of interaction—it turns out there is a long tradition in business ethics of denouncing this exact comparison (the usual object of opprobrium is Carr [1968]).

Thus I resolved to write a series of papers, each one focused on a different stumbling block that I had encountered. I also set out to read more carefully the business ethics literature, in order to understand better the assumptions and motivations of those who were critical of my project. The first part of this book is essentially a record of these efforts. The second part then tries to

articulate some of the broader “political economy” considerations that inform the project—in particular, my understanding of contractualism, the relationship between cooperation and social institutions, and the general justification for the market economy. The third part contains what I like to think of as “extensions” of the theory, primarily focused on how the framework can be used to address issues that arise in a management context, or within the firm.

1.1.1. PROFIT

The first stumbling block was the fact that I failed to take issue with the profit-orientation of the firm, and by extension, had no particular objection to shareholder primacy. This obviously generated a headlong conflict with what has arguably been the major trend in North American business ethics, the view that managers should be held accountable to a variety of different “stakeholder” groups, with investors enjoying no special privilege. Coming at things from a background in political economy, the stakeholder perspective seemed to me a strange view, simply because making peace with capitalism essentially involves acknowledging the value of the profit motive (since it is the quest for profit that generates the competitive dynamic that allows the price system to exhibit the desirable properties that it does). Arguing against the profit orientation of firms seemed to me equivalent to arguing for some kind of market socialism, which may be a perfectly respectable position to take, but probably should not be a position *in business ethics*. It seemed to me rather a case of changing the topic (away from the question of how economic actors should behave in a capitalist economy to whether we should have a capitalist economy at all).

Furthermore, there is a deep and sophisticated literature in socialist economics dealing with the problems that were encountered by managers in organizations that are not subject to the discipline of profit-maximization, both in the former communist countries under central planning and in state-owned enterprises under democratic welfare states (Kornai 1992; Nove 1983; Roemer 1994; Stiglitz 1994). For example, there is the well-known fact that investor ownership imposes a “hard budget constraint” that is very difficult to replicate under public ownership. It seemed to me likely that the reorganization of firms along the lines proposed by stakeholder theorists would encounter many of the same difficulties. Thus my first foray into the literature was the collaborative piece with Wayne Norman (chapter 2), the general purpose of which was to encourage business ethicists to confront the literature in public administration that discusses these challenges.

This paper, however, dealt essentially with implementation problems that stakeholder theorists were likely to encounter; it did not take issue with the normative core of the theory. Thus the following paper (chapter 3) represents my first attempt to explain the normative inadequacies of the theory. The central argument is that stakeholder theory, in its standard formulation, expands

managerial obligations to encompass the interests of groups other than just shareholders, but does so in a way that is either too broad or else arbitrary from the moral point of view. Arguments of this sort have become increasingly familiar in recent years and are achieving much broader acceptance (Boatright 2006; Orts and Strudler 2009; Marcoux 2000). There is, however, an important complication, which I became aware of after having acquired greater familiarity with the literature. It is common to talk about “shareholder primacy” as the legal norm in much of the world (in contrast to, say, the “co-determination” arrangement in Germany that gives workers representation in the governance of certain firms). As a matter of fact, what the law actually provides is a menu of options, which permits all sorts of different organizational forms. In particular, as reading the work of Henry Hansmann (1992) impressed upon me, there is nothing to stop any constituency group from serving as owners of the firm, by forming a cooperative rather than a standard business corporation. Furthermore, legal “partnerships” are often *de facto* cooperatives, or multi-constituency ownership structures (that include, say, workers and investors). Furthermore, it is not always the case that in open marketplace competition, standard business corporations do better than cooperatives. The insurance industry, for example, was for decades dominated by consumer cooperatives (or “mutuals”).

Thus it turns out to be incorrect to describe modern market economies as being governed by a norm of shareholder primacy within firms. What they are actually governed by is a norm of *owner primacy*. This, combined with the empirical regularity that most firms are owned by the providers of capital, generates the illusion that shareholder primacy is the legal norm. This raises a whole host of questions, which business ethicists have been slow to confront (with notable exceptions, e.g., Boatright 2002). For example, is the “market for control” unfairly biased against non-shareholder groups, or are there good reasons for the prevalence of investor ownership? (Miller 1989: 83–93; Dow 2003). If one accepts Hansmann’s redescription of the shareholder-owned firm as essentially an “investor’s cooperative,” how does this affect our thinking about managerial responsibility? If we have no problem with “member primacy” in the case of a cooperative, why should we take issue with it when the “members” happen to be the lenders? These are some of the questions that I address in chapter 5. In general, what I find transformative in Hansmann’s analysis is the suggestion that shareholder primacy and profit-maximization be understood as just special instances of owner primacy and maximization of the residual claim. I take issue, however, with some of the normative conclusions that he draws from this analysis.

1.1.2. COMPETITION

The second stumbling block stemmed from the fact that, from the very beginning, I regarded the market failures perspective as implying a system

of *adversarial ethics* for all transactions mediated by the price mechanism. Adversarialism, in my view, involves *deontic weakening* with respect to everyday morality, where certain actions that would ordinarily be obligatory may become optional (and, equivalently, actions that are forbidden may become permissible). Use of the price mechanism implies an adversarial orientation because prices are competitively determined. A competition, in my view (articulated at length in *The Efficient Society*) is essentially an institutionalized collective action problem, where we are released from the everyday-morality obligation to act cooperatively and are actively encouraged to play free-rider strategies. This is a slightly counterintuitive view of competition, although the central idea is well captured by the old saying that “every free market is a failed cartel.” The major objective of chapter 4 is to articulate and defend the close connection between the price system, marketplace competition, and adversarialism in business ethics.

This is why, despite the various disanalogies, I think that it is still illuminating to draw comparisons between business ethics and the ethics of sport (and why I find myself agreeing with Alfred Carr’s unpopular view that business requires individuals to “discard the golden rule” [1968: 145]). Obviously sport is both voluntary and unserious, in a way that having to earn a living in a capitalist economy is not. This would be an issue if I thought that the voluntariness of transactions was important to the normative justification of the market (which I do not). What I find illuminating about the comparison to the ethics of competitive sport is that it focuses on the question of “how far do you go to win?” It seems to me that the structure of reasoning one must employ (in particular, the way that one must think about one’s intentions and goals to resolve this question) can be usefully applied in the business context as well (as I argue in chapter 4).

1.1.3. EFFICIENCY

The third major stumbling block is the almost singular emphasis that I put on the principle of Pareto efficiency in providing normative foundations for the view. This has the potential to give rise to a number of misunderstandings, particularly among those who see the word “efficiency” and assume that it is merely an instrumental principle, or think that there is some sort of conceptual connection between efficiency and the pursuit of individual self-interest. One need only contemplate the structure of a prisoner’s dilemma (as in figure 6.1) to see that efficient outcomes are not an automatic consequence of individual utility-maximization, and that in such interactions the Pareto principle serves as a genuine constraint on the pursuit of self-interest (Gauthier 1986).

In a slightly more sophisticated vein, efficiency is often conflated with utilitarian welfare-maximization, or else the wealth-maximization standard proposed by “law and economics” scholars (Posner 1973). I use the term “efficiency,” by contrast, in the strict Pareto sense, to refer to the principle that,

whenever it is possible to improve at least one person's condition without worsening anyone else's, it is better to do so than not. In practice, this simply commits one to promoting cooperation (in the game-theoretic sense), it says nothing about the specific modalities of cooperation, other than that the benefits should be maximized (i.e., that the set of available Pareto improvements should be exhausted). In economic terms, one way of thinking about the Pareto standard is to regard it as a general prohibition on waste, since if it is possible to rearrange the allocation of resources in such a way as to improve one person's welfare without worsening anyone else's, and yet this is not being done, it means that some resources are being wasted under the status quo.

It is worth emphasizing that I do not assign a central place to efficiency in the assessment of markets because I think it is a foundational value, or that it arises endogenously out of state-of-nature interactions. I think it is an irreducibly normative principle, which constitutes one element of a broader contractualist theory of justice (of the sort that I sketch out in chapter 6). It is only one element because it must to be supplemented with some conception of distributive justice in order to provide a persuasive standard for evaluating the overall system of social cooperation. This does not mean, however, that every domain of interaction is subject to assessment under the full theory of justice. There is, in my view, a division of moral labor within our institutions, with markets being essentially special-purpose institutions designed to promote efficiency (a view defended in greater detail in chapter 7). Thus it is only when embedded within the broader context of a welfare state, which engages in both market-complementing and redistributive policies (primarily through the tax system), that capitalism as a whole can claim to be just. At the same time, this does not mean that market actors are accountable to the same moral demands that the system as a whole must satisfy. Individuals are given license to maximize profits (and to associate in various ways to engage in joint action aimed at maximization of profits), for the narrow reason that, in a reasonably competitive market, this is the best way to get prices that reflect social cost. In order to achieve this, individuals must be given a fairly broad exemption from norms of equality or fairness in the organization of their interactions. To the extent that this is justifiable, it is because the compromises made in the equality dimension (due to the competitiveness of market interactions) are outweighed by the benefits that accrue in the efficiency dimension (due to the operations of the price system). Because of this moral compromise at the heart of capitalism, one cannot hold economic actors engaged in market transactions to a higher standard than that of efficiency promotion. But this is all the more reason to be rigorous in holding them to this standard. A competitive market only serves to promote efficiency under certain conditions, and there are various ways of acting that subvert it. Such actions are not just unethical, but egregiously so, because they fail to satisfy even the artificially low standard that is set for the evaluation of marketplace behavior.

Now to say that the market failures approach to business ethics is “guided” by the Pareto principle is not to say that individuals, when deciding what to do, should ensure that the outcomes of their own actions are always Pareto improving. This is obvious in the context of a market economy, because competitively structured interactions are designed to produce win-lose outcomes (not the win-win outcomes required by the Pareto principle). When a company lowers its asking price, in order to get rid of unsold merchandise, its actions harm its competitors, in a way that is collectively self-defeating when the competitors respond in kind. Our reason for allowing this sort of collective action problem to persist is that it generates, as a byproduct effect, a movement of prices in the direction that will clear the market, and therefore that will maximize the number of efficiency-promoting exchanges with purchasers. The market therefore institutionalizes an indirect strategy for promoting Pareto efficiency, in the form of rules that specify the terms of marketplace competition, in particular, which competitive strategies are permissible and which are not.

Many of the rules of the market are legally enforced, but it is impossible to imagine a circumstance in which they would all be. For example, an enormous amount of legal emphasis is put on the minimization of externalities. The underlying principle, with respect to negative externalities, is that if an action has consequences that are damaging for some other person then there should also be a cost for the person who is engaged in it, and that, to the degree possible, the cost to the person doing it should reflect the magnitude of the damage done. (The principle for positive externalities is just the reverse.) When prices reflect social cost in this way, it ensures that the overall production of costs and benefits will be Pareto efficient. The most important legal mechanism that we have, when striving to achieve this outcome, is the system of property rights itself, which can be thought of as an all-purpose mechanism for internalizing externalities. Through ownership, the individual is able to “capture” much of the value that she produces through her labor, thereby minimizing positive externalities (and thus, increasing the incentive to produce value in the first place). And by asserting her property rights (e.g., against trespass or unauthorized use), she is able to “deflect” many of the negative externalities that others might like to impose on her. The tort system represents an extension of this basic mechanism into areas that are less clearly structured (e.g., generating legal constraint on various forms of nuisance behavior that are not in direct violation of anyone’s property rights, or that violate these rights in unanticipated ways). And finally there is regulation, which attempts to control the production of negative externalities in myriad ways (e.g., restricting the production of atmospheric pollution, excessive noise, toxic and dangerous substances, etc.), without requiring private individuals to step forward and assert their rights.

And yet, despite all of this effort, there are still many circumstances in which the legal system is powerless to stop the production of negative

externalities. In some cases this is simply because of cost considerations, which are quite high for criminal, tort, and regulatory law. The resources consumed in the effort to stop someone from doing something will often be of greater value than the losses imposed upon society by the behavior. In many cases the behavior is undetectable, or the victims unidentifiable. In other cases, particularly ones with an international dimension, no state has the legal authority or power to effectively control the behavior. What the market failures perspective suggests is that in such cases, economic actors have moral obligations that extend beyond their legal obligations, but that these moral obligations are merely an extension of the basic rationale underlying the law—understood broadly, to include the system of property rights, the tort system, and body of regulatory law. So if it is possible to increase revenue by displacing costs, rather than creating value, this may be morally prohibited, regardless of whether it is legal. In this way, the ideal of promoting Pareto efficiency generates a more specific deontology that constrains the behavior of economic actors. Furthermore, this deontology is generated by a normative theory that provides a unified account of the foundations of the market economy, the purpose of regulatory intervention and state ownership, as well as the basic “beyond compliance” obligations of firms (Norman 2012: 48–49).

1.2. The Goal of Business Ethics

Many business ethicists, perhaps responding to the pressures of modesty, deny that they have any ambition to make people behave more ethically. I actually consider that objective to be central to my task. I strongly agree that the point of philosophy is not just to understand the world, but to change it. Furthermore, as I suggest in chapter 11, if persuading our students to behave more “ethically” in later life seems like too lofty an ambition, perhaps making them less likely to behave criminally would be a useful contribution, and a more achievable one. After all, most of the classic “cases” that business ethicists like to discuss, from the Ford Pinto to the Enron scandal to the Deepwater Horizon disaster, are not really “ethics” cases at all, but rather examples of occupational or corporate crime. In fact, one of the distinguishing features of business ethics, as a domain of applied ethics, is that it deals with an area of social life in which crime is a very serious problem.

This is one of the reasons why, in my view, explaining the complementarity of morality and law is very important. I take it to be one of the central tasks of business ethics to articulate the normative foundations of regulation and to explain why managers should adopt a moral attitude toward compliance. Indeed, if business ethicists were to forget about “moral dilemmas” entirely and just focus their energies on trying to articulate the moral reasons for firms

to comply with existing laws, they would stand a much better chance of producing some benefit for society.

It seems to me that one of the useful tasks that business ethicists can perform, in the service of this ambition, is to combat two extraordinarily pernicious views (or, one is tempted to say, “ideologies”) that are, unfortunately, quite widely held. The first is the idea that, in a market economy, corporations have no obligation beyond respect for the law. Perhaps the most high-profile proponent of this doctrine is *The Economist* magazine, where it is reiterated often enough to suggest that it is part of the magazine’s official editorial stance. The following is a typical articulation of the view: “A company’s job is to make money for its shareholders legally. Morality is the province of private individuals and of governments,” and so if politicians want to change business behavior, “they should pass laws, not make speeches” (*The Economist* 2011: 18). One can easily find examples of businesspeople expressing variants of the same view (Carr 1968: 146–147).

Many will recognize this as a somewhat unnuanced rearticulation of the view that Milton Friedman (1970) expressed, when he claimed that the only social responsibility of business is to increase its profits. Friedman actually qualified this slightly, claiming that managers should typically try “to make as much money as possible while conforming to the basic rules of the society, both those embodied in law and those embodied in ethical custom” (1970). It is clear, however, that he did not imagine “ethical custom” imposing very significant constraints on business behavior. With respect to reducing pollution, for example, Friedman was clearly of the view that managers are obliged to do the minimum required by law—“ethical custom” generated no additional constraints that could be appealed to (or at least none that he mentions). A manager who voluntarily reduced emissions was, in effect, usurping a public power, and unjustly imposing a “tax” on shareholders.

According to this view, when faced with a demand to discontinue a particular business practice on grounds of “social responsibility,” the appropriate response from the business manager is to say: “As long as it’s legal, we are going to do it. If there’s a problem with that, then the government should pass a law to make us stop.” Yet although this claim is often made, it is difficult to take seriously when one stops to think about it. After all, who could possibly want that level of legal regulation of economic activity? Because law is the most intrusive and costly form of social control, it is typically appealed to as an intervention of the “last resort” (Simpson 2002: 112). Furthermore, there are well-known difficulties associated with trying to regulate the behavior of firms that adopt a broadly uncooperative orientation (i.e., that exhibit no moral constraint in their attitude toward compliance). For example, there are clear trade-offs involved in determining the level of specificity at which regulations should be framed: write the rules too broadly and it creates legal uncertainty, as well as higher enforcement costs; write the rules more narrowly

and it encourages circumvention (or “gamesmanship”), as well the inefficiency caused by overly rigid specifications (Braithwaite 1981–82: 483–484).

Indeed, there would appear to be a strong element of bad faith in *The Economist’s* espousal of this doctrine, since the magazine is also a vocal critic of government “overregulation.” This makes it difficult to believe that they actually support the dramatic extension of regulation that would be required if firms were to abandon all self-restraint in the pursuit of their objectives. Thus one begins to suspect that a shell game is being played, where moral constraints are rejected on the grounds that they should be juridified, but then legal constraint is rejected on the grounds that it is too costly, passing it back to moral constraint.

The position is also strangely unmotivated. Why would a particular institutional actor be exempt from moral constraint? The idea that we are all as individuals obliged to act morally, but when a few of us get together and sign articles of incorporation, we are suddenly licensed to do anything at all in pursuit of our interests, subject only to the constraint of law, lacks even *prima facie* plausibility. It is not clear that any social institution has the power simply to absolve people of moral responsibility for their actions. So to the extent that anyone thinks that firms are outside the scope of moral constraint, the most likely explanation is undoubtedly some form of muddled “invisible hand” reasoning. Rather than thinking that articles of incorporation offer *carte blanche* to act immorally, the view is more likely that morality becomes unnecessary in a business context, because the invisible hand of the market guarantees that individuals are only able to pursue their self-interest in ways that will also, as a byproduct effect, maximize social welfare. So the aims of morality are, as it were, guaranteed, without any need for constraint. This is, for example, the view articulated by David Gauthier, who argues that there is simply “no need for morality” in the competitive market: “Where earlier thinkers saw in the unbridled pursuit of individual interests, the ultimate source of conflict in human affairs, the defenders of *laissez faire* see in it rather the basis of the true harmony that results from the fullest compossible satisfaction of those interests. The traditional moralist is told that his/her services are not wanted” (1982: 47).

The problem with this view is not so much that it is bad ethics but that it is bad economics. It vastly overestimates the success of market institutions (in effect, the legal framework that structures economic activity) at achieving this reconciliation. One need only consider the “efficiency conditions” required for the first fundamental theorem of welfare economics to obtain (Schultz 2001). It is easy to see that the reconciliation of public and private interests is never guaranteed by any set of existing market institutions. And when the alignment of private and public interests is not achieved automatically by the market, some attempt must be made to do it consciously and explicitly, through moral constraint. Gauthier does not disagree with this: “Where the Invisible Hand

fails to direct each person's actions to the public interest—or, as I shall prefer to say, to mutual benefit—the Visible Foot takes over. Hand and Foot share a common aim” (1982: 41).

The central disagreement between myself and Gauthier is over how much work actually gets done by “hand,” and how much is left over for the “foot” of morality.¹ Here the progress of economic thinking over the past thirty years has done much to clarify the issues. At the time that Gauthier was dismissing the need for business ethics, it is worth recalling, there was still very little awareness that pollution could cause serious quality of life issues. This is why Gauthier, like Friedman before him, despite being aware of the problem of externalities, glibly dismisses them. Gauthier was also writing at a time when “information economics” was still in its infancy, and as a result simply ignored the possibility of market failure caused by information asymmetries. So he had nothing useful to say about the market for insurance, professional services, or even branded goods, where information dynamics play a dominant role in determining whether transactions will occur and at what price.

Thus the first pernicious view—that firms are obliged to respect the law, but have no “beyond compliance” obligations—is often put forward insincerely, but when it is sincerely held, it is usually based on an overestimation of the effectiveness of the invisible hand at promoting social welfare. The market failures approach to business ethics arises, then, as a natural consequence of combining a normative commitment to Paretianism with a modern economic understanding of the conditions under which mere compliance in a market is unlikely to promote social welfare.

The second pernicious view constitutes a further attempt to expunge ethics from the realm of the marketplace interaction. Many social scientists have, over the years, been uncomfortable with “ethics,” on the grounds that moral motivation is thought to be something mysterious and obscure. External incentives, by contrast, seem much more solid, quantifiable, and verifiable in their effects. This general feeling of unease was, notoriously, elevated to the level of a strict methodological precept by economists (for discussion, see chapter 10). One of the consequences of this theory of action was that it made the law seem much more intelligible than morality. Indeed, many economists affected a certain cynicism toward morality, believing that it could not actually motivate or constrain agents, but was more likely a rationalization of self-interest. Law, by contrast, seemed like a more respectable social-scientific construct—you get punished if you break it, so there is no reason to doubt that legal rules actually constrain the way that economic agents behave. One can believe in both legal regulation and individual utility-maximization without cognitive dissonance.

¹ Indeed, if it were not irremediably obscure, the phrase “Hand and Foot share a common aim” could easily serve as the slogan for the market failures approach to business ethics.

The problem with this view is not just that it is dismissive of morality, but that it also generates a terrible misunderstanding of how compliance with legal rules is actually achieved. To the extent that this misunderstanding is encouraged by the adoption of an “economic” model of human behavior, it is important to emphasize that the specific thesis—“people obey the law because they fear punishment”—is an *a priori* deduction from the model, it is not a fact that anyone has discovered about the world. When one turns to the empirical literature, it turns out, as a generalization, to be false. While deterrence (i.e., the threatened application of sanctions) is an important element of compliance, it is very far from being the entire story. The most obvious reason is that the state lacks the resources and the information required to deter most crime. Given existing legal infrastructure, if people actually sat down and calculated, in a hard-headed way, how best to advance their own individual interests, then crime would be the rule rather than the exception. As a result, we (as a society) rely very heavily on internal controls (such as moral constraint) and informal social controls (such as stigmatization) to achieve conformity with legal rules. Furthermore, it is generally recognized among criminologists that these mechanisms do most of the heavy lifting when it comes to controlling crime. One of the most influential schemas in the literature is John Braithwaite’s “enforcement pyramid,” which features internal self-regulation at the bottom, followed by informal social controls, legal persuasion, and finally deterrence, in an ascending hierarchy (Ayres and Braithwaite 1992: 38–39). The reason it is a pyramid is that self-regulation is sufficient to control most of the population most of the time. It is only where this fails that informal social controls become active, and it is only when these fail that official legal intervention is required. Even then, legal intervention seldom starts out with punishment, but usually begins with some form of engagement aimed at correcting the behavior (such as an order, or a warning). Thus deterrence is at the top of the enforcement pyramid simply because it is only for a very small percentage of the population that the threat of punishment is required in order to achieve compliance.

It is sometimes assumed that deterrence is going to be more important when it comes to controlling corporate crime, compared to street crime, because corporations are more instrumentally rational, and hence more likely to be doing explicit cost-benefit calculations in deciding what to do. Yet surprisingly, the exact opposite is true; society relies more heavily on voluntary compliance in controlling white-collar crime than it does with street crime. This is because the state has much greater difficulty detecting, prosecuting, and securing convictions with white-collar criminals. There are a number of reasons for this (Simpson 2002: 45–60; Coleman 1989: 185–194). The first has to do with the nature of the victims, who are usually a very diffuse group and often unaware that they have been victimized. This makes the detection of criminal behavior much more difficult, since the state relies very heavily on victim complaints as a cue to initiate investigation. With street crime, the state

also derives considerable advantage from the fact that most crimes are committed by individuals or loosely formed groups. This is why organized crime is so difficult to combat (and why significant enforcement effort is aimed at breaking up organized criminal groups, or preventing their formation). Many of the due process safeguards that have been put in place are aimed at protecting isolated individuals from the more organized power of the state. But when crimes are committed by a corporation, the state is dealing with a highly organized group (which, in the standard run of cases, the state has no desire to disrupt or break up), which nevertheless enjoys the same due process rights that protect individuals. Furthermore, corporations are often able to invest more resources in mounting a defense than the state can marshal for its enforcement efforts.

All of these factors, combined with the information asymmetries that make it very difficult to figure out when a crime has been committed, makes the probability of apprehension very low with most white-collar offences. Every so often a major scandal or prosecution will generate closer scrutiny, but when it does, it will often reveal systematically criminal behavior across an entire industry that has gone undetected for years. (To take just one example, under the 2012 mortgage fraud settlement in the United States, almost all the major banks operating in the American mortgage market admitted to using forged documents to illegally foreclose on homeowners [Dayen, 2013]. Commentators struggled to find adjectives that could fully convey the enormity, scale, and boldness of the crimes committed, especially in an industry that should have been expecting heightened scrutiny in the wake of the 2008 financial crisis.)

Simplistic cost-benefit reasoning suggests that these enforcement problems could be made up for by an increase in the associated penalties. There are, however, limits on this strategy in the case of corporate crime. Most obviously, limited liability puts a ceiling on how large fines and damage awards can be—if they become too large firms will simply declare bankruptcy. Furthermore, the threatened loss to managers who are making the decisions is often much lower than the potential loss to shareholders, so there are agency problems to be considered. This is in fact what underlies the trend toward criminal prosecution of individual managers, instead of monetary damages from the firm. But this creates all kinds of other problems, because of the internal complexity of firms and the difficulty of assigning responsibility. Again, the issue is that the criminal law is very strongly tailored toward dealing with individuals, not groups, and so when actions actually are planned and carried out by a group agent, like the firm, it can be very difficult to apply many of the essential categories of criminal law.

Thus the dominant view among criminologists is that it is often impossible to come up with a threatened punishment that is large enough, credible enough, and sufficiently well targeted to serve as an effective deterrent against

corporate crime (Braithwaite 1981–82). So to the extent that corporations do respect the law, it must be due to a higher level of moral self-restraint than is operative in the area of street crime. This is, of course, not implausible when one looks at the profile of a typical businessperson. As Edwin Sutherland pithily observed, “businessmen are generally not poor, are not feeble-minded, do not lack organized recreational facilities, and do not suffer from the other social and personal pathologies” (1968: 58). So there is certainly no reason to expect them to be immune to the force of moral constraint. Unfortunately, there is a certain line of thinking—again, strongly influenced by economics—which denies that corporations have any moral obligation to obey the law. Perhaps the most high-profile exponents of this view are Frank Easterbrook and Daniel Fischel, who argue that managers should adopt a purely instrumental orientation toward regulation, and comply with it only when it is in the firm’s interests to do so. Any punishments the firm may incur as a result of breaking the law should be regarded as just another cost of doing business. According to Easterbrook and Fischel, “Managers do not have an ethical duty to obey economic regulatory laws just because the laws exist. They must determine the importance of these laws. The penalties Congress names for disobedience are a measure of how much it wants firms to sacrifice in order to adhere to the rules; the idea of optimal sanctions is based on the supposition that managers not only may but also should violate the rules when it is profitable to do so” (1982: 1177).

This latest refinement of “Chicago school” doctrine is far more audacious than what Friedman proposed. Friedman thought that firms should maximize profits using all legally available means. Easterbrook and Fischel claim that illegal means are perfectly acceptable as well, so long as the firm’s actions are profit-maximizing *ex ante*, taking into consideration the probability of apprehension and severity of punishment. Indeed, they go further, arguing that managers are positively obliged to break the law when it is profitable to do so, anything less would be a dereliction of duty toward shareholders. They arrive at this alarming conclusion by taking the doctrine of proportionality—the idea that the severity of a punishment should reflect the severity of the offense—and interpreting it to mean that, by making a particular sort of behavior illegal, the legislature is not really prohibiting it, but just pricing it. This is complemented by the economic view that the law can control behavior only by providing external incentives, and so it would be unreasonable to expect firms to do anything other than respond to the incentives provided. It follows that, when faced with a demand to discontinue a particular business practice on grounds that it is illegal, the appropriate response from the business manager is to say: “As long as it’s profitable, we are going to do it. If there’s a problem with that, then the government should increase the penalties for non-compliance.”

It is difficult to express just how irresponsible this view is, particularly when promulgated by a sitting judge. And yet these are not ideas from the

fringe. Easterbrook and Fischel are among the most influential theorists of corporate law of the twentieth century, and their position has become something close to Chicago school orthodoxy. Yet anyone acquainted with the empirical literature can see that what they are proposing is a recipe for utter chaos. Furthermore, there is good reason to think that very few people would be attracted to the idea of street criminals doing these sorts of calculations when deciding whether to mug a pedestrian or break into one's home. This suggests, in turn, that their analysis is not so much a distinctive normative position as it is an ideology (i.e., a system of motivated false belief), one that is based on a deep misunderstanding of how market economies actually function.

In this social context, it seems to me that the central role of business ethics is to provide an “immanent critique” of corporate conduct (Benhabib 1986). Its objective is not to bring in “outside” moral considerations to condemn the latest outrage, but to clarify and to correct the self-understanding of participants in the market economy, who are being bombarded—both by the business press and a certain segment of the academy, who appear not to have recovered from the epiphany they experienced in their first-year economics class—by a seductive but ultimately false suggestion that the institutions of the market free them from all forms of moral constraint. In order to do so, it has no need to appeal to normative standards beyond those that are already implicit in the institutions of a market economy.

I.3. Further Directions

By way of conclusion, I would like to mention a few areas in which the market failures approach to business ethics stands in need of much further theoretical development—issues that are touched on in the papers collected here, but are by no means given a satisfactory treatment. The doctrine as it is presented is, first and foremost, intended as an approach to the question of corporate social responsibility, in the sense that it deals with the obligations that managers have to individuals outside the firm, in relations that are mediated by the price mechanism (or should be mediated by the price mechanism, in the case of certain externalities). It says very little about relations inside the firm, among those engaged in administered, rather than market, transactions (e.g., between line managers and those who are subject to their authority). The central difference is that administered transactions are cooperative, in a way that competitive relations in the market are not. Because these relations are not intrinsically adversarial, the norms that govern behavior within the firm more closely resemble the norms of everyday morality. For example, wages are much more egalitarian within firms than across different firms (Frank 1985: 35–57). Although a number of different factors play into this, it is not difficult to

imagine that concerns about fairness have much greater force within firms simply because people need to work together cooperatively.

However, the fact that norms within the firm more closely resemble everyday morality should not mislead one into thinking that they are the same. For instance, the firm can only go so far in satisfying people's intuitions about "fair wages," because it is constrained by the need to keep its internal transaction costs lower than what the market alternative would be. (For example, the more egalitarian the wage structure is within a firm, the more attractive it will be to "contract out" certain functions.) So while firms may cultivate their own internal values or ethos, as part of the corporate culture, which is much "thicker" than the extremely minimal framework that governs market transactions, it is still constrained by a general set of efficiency imperatives that arise out of the need to keep its administrative transaction costs lower than those of a corresponding set of market transactions. Because of this, while I think there is much to be said for Allen Buchanan's "agency risk minimization" (1996) analysis of the ethics of bureaucratic organizations, I also think there is a lot more to the story than what agency vocabulary is able to reveal. (To take just one example, reciprocity is extremely important when it comes to establishing cooperative relations between individuals in face-to-face interaction, yet the principal-agent framework is designed to handle only unilateral relationships of influence. It is therefore likely to overlook many of the devices that, say, managers use to motivate employees, devices that may in turn generate moral obligations.) Thus my discussion in chapter 4, particularly figure 4.3, provides too simplified a picture of these relations. At the moment, however, the most that I feel I can say about intrafirm relations is that they are complex and also extremely contextual. Analyzing them from an efficiency perspective may not provide the most perspicuous understanding; it may simply be the only thing that can be said about them at a high level of generality.

The second major issue that I would someday like to address more thoroughly concerns the non-ideal aspect of the theory (in the sense described in chapter 7). In particular, I am interested in trying to delimit more carefully the circumstances in which competitive pressure actually does offer economic actors a legitimate excuse for acting immorally, as well as the subsidiary obligations that might accompany that excuse. A related question arises about the circumstances in which firms might be justified in breaking the law. Both Wayne Norman and I have a tendency to talk about "regulation" as though it was mostly aimed at correcting market failure, which is why the "spirit of the law" is taken to be consonant with the general thrust of business ethics. But this is not always a fair characterization of regulation. There is, of course, the famous thesis advanced by "public choice" theorists, who see regulation almost in its entirety as a consequence of rent-seeking by various actors. In this context though, it is worth observing that this account of regulation is not based on empirical evidence, it is an a priori deduction from an economic

model of action combined with the assumption that legislators always act in a self-interested fashion, assigning no value at all to the public interest. In other words, public choice theory and the attendant view of regulation is part and parcel of the economistic ideology that must be rejected in order for business ethics as an enterprise to make any sense at all. It cannot therefore be appealed to as an objection to any particular view in business ethics.

A more plausible challenge to the market failures approach would be one that points to specific regulations, such as a professional licensing requirement that serves an obvious cartelizing function, then highlights the conflict between the inefficiency of the law and the efficiency imperative arising from an ethical perspective. Should a firm then be prepared to break the law? It seems to me that two points are in order with respect to such a case. The first is the familiar one from the literature on civil disobedience, which observes that a law can be unjust or immoral, yet still be legitimate, and therefore worthy of obedience. The threshold at which civil disobedience becomes justified is typically set higher than the level at which a law is determined to be unjust. The second is the observation that, as long as corporate crime remains such a serious blight on society, the problem of excessive business compliance is likely to remain somewhat academic, in the pejorative sense of the term. Given a choice between the world we live in and some other possible world in which businesses automatically and unthinkingly obeyed every law, including those that are unjust, I would not hesitate to choose the latter.

The details of this, however, would need to be worked out, since there are a number of controversial issues—including the distinction between excuses and justifications, which I discuss briefly in some of the papers here, but make no attempt to defend systematically.

Finally, readers are sometimes a bit flummoxed by the fact that I present myself as a critic of “economism,” while at the same time making extensive use of economic concepts and accepting the essential correctness of several contemporary economic theories. The reason for this is never directly articulated in any of these papers, but it is due to the fact that I reject the narrow instrumental conception of rationality at the heart of utility-maximization theory, while nevertheless accepting several “meso” level economic theories—such as the transaction cost theory of the firm, as well as something like the traditional supply-demand model of price determination. This is simply because I do not believe that the correctness of these theories depends on the correctness of the traditional utility-maximization model; they can be given other micro-foundations. (Those who are interested in my official views on action theory, and the way in which moral considerations can be integrated into a formal model of action can find some indications in chapter 10, but also are encouraged to consult my book *Following the Rules*, where the position is laid out in greater detail.)

PART } I

The Corporation and Society

A Market Failures Approach to Business Ethics

“Business ethics” is widely regarded as an oxymoron. The only way to be a good soldier in an unjust war is to disobey orders, or maybe even to desert. Many people believe, along similar lines, that the only way to maintain one’s ethical integrity in business is not to go into business. The reasons for this are not hard to find. Students are still routinely taught in their introductory economics classes that in a market economy, when engaged in market transactions, individuals act out of self-interest—whether it be by maximizing profits as producers, or by maximizing satisfaction as consumers. This sets up an almost indissoluble link in people’s minds between “profit-maximization” and “self-interest.” As a result, anyone who thinks that the goal of business is to maximize profits will also tend to think that business is all about self-interest. And since morality is widely regarded as a type of constraint on the pursuit of individual self-interest, it seems to follow quite naturally that business is fundamentally amoral, if not immoral.

The problem is that the association between profit-maximization and self-interest so often taken for granted is based upon a naïve and inadequate theory of the firm. Profit-maximization and self-interest are not the same thing, and the failure to distinguish adequately between the two can be a source of enormous confusion. Business ethics, as a subject, is essentially concerned with the moral responsibilities of managers. Managers often find themselves placed in circumstances in which the imperative to “maximize shareholder value” conflicts with their self-interest. Thus there are many cases in which profit-maximization should be viewed as a managerial obligation, not as an expression of self-interest.

Because of this somewhat elementary confusion, there has been a marked tendency in the business ethics literature to dismiss out of hand views that take the profit motive seriously. In particular, Milton Friedman’s classic article “The Social Responsibility of Business is to Increase its Profits,” is more often treated as a piece of apologetic than as a serious piece of moral reasoning (Friedman 1970). This is unfortunate, since the moral laxity on display

in Friedman's work is not so much a symptom of an inadequate normative framework as it is a consequence of specious economic reasoning. Or so I will attempt to show.

The more serious consequence of this confusion is the widespread perception that, in order for business ethics to be genuinely ethical, it must extend managerial responsibility to groups other than shareholders. This is, I believe, often the intuition underlying "stakeholder" theories of managerial responsibility. In this paper, I will argue that such efforts are misguided. Profit-maximization, understood as an obligation, rather than as an expression of self-interest, provides a perfectly legitimate platform for the development of a robust moral code. However, if profit-maximization is an obligation, the question naturally arises where this obligation stems from. It is in seeking to justify the profit motive that we discover that the appropriate form of managerial responsibility is not to maximize profits using any available strategy, but rather to take advantage of certain specific opportunities for profit. In many cases, the set of conditions under which profit-seeking is permissible is reflected in the legal environment in which firms operate. I will argue that business ethics is best understood as a set of additional constraints that preclude legally permissible, but not normatively justifiable, profit-maximization strategies.

1.1. The Profit Motive

Andrew Stark's controversial 1993 *Harvard Business Review* article, "What's the Matter with Business Ethics?" argued that conventional business ethics was "largely irrelevant for most managers," because it failed to offer them any "practical" advice (Stark 1993). "Moral philosophy," he argued, "tends to value *altruism*, the idea that an individual should do good because it is right or will benefit others, not because the individual will benefit from it" (Stark 1993: 40). As a result, business ethicists have had too little to say about "the potential conflict between ethics and interests," and in particular, how managers should handle such conflicts when they arise.

This article had many people nodding their heads in agreement. But to see just how peculiar the claim is, suppose that the subject had been medical ethics instead of business ethics. Substitute "doctors" for "managers" throughout. Now imagine criticizing medical ethics on the grounds that it fails to offer doctors any "practical" advice on what to do in cases where the imperatives of patient care conflict with their self-interest. Suppose the patient doesn't really need an operation, but the doctor could make a lot of money by performing it anyway. What to do, what to do?

I would suggest, *pace* Stark, that we do not need professional ethicists to tell us where our obligations lie in such cases. Everyone knows that when there is a straightforward conflict between our self-interest and our moral obligations,

the moral obligations win, at least from the moral point of view. This is not “ethical absolutism,” as Stark maintains, it is simply the logic of moral justification. The question of when we may be forgiven for disregarding our moral obligations (i.e., acting immorally) is a separate one and is in no way specific to the domain of business ethics.

So why does Stark’s argument sound even remotely plausible, whereas a comparable argument in medical ethics would be dismissed out of hand? The confusion has two distinct sources. The first arises from the way that introductory economics is usually taught. The standard microeconomics textbook starts out with the assumption that individuals maximize utility. When it comes to particular goods, these utility functions can be represented as a set of indifference curves. These indifference curves are then taken to provide the supply and demand curves. The thesis that individuals maximize utility is interpreted to mean that consumers will seek to maximize satisfaction, and suppliers will seek to maximize profits. Finally, in order to make the model more “realistic” consumers get aggregated together into “households,” and suppliers into “firms”—each of which is thought to maximize some joint utility function.

While everyone understands that “the firm” is something of a black box in this analysis, the result is still an unhelpful blurring of the boundaries between the pursuit of self-interest and the maximization of profits. Stark, for instance, variously describes the conflict that managers face as one between “self-interest and altruism,” “ethics and interests,” “ethical demands and economic realities,” “moral and financial costs,” “profit motives and ethical imperatives,” and even “consumer’s interests” versus the “obligation to provide shareholders with the healthiest dividend possible” (Stark 1993: 44). Here we see a clear blurring of the distinction between self-interest, profit-maximization, and the obligation to shareholders.

We understand implicitly that the professional conduct of doctors is to be entirely governed by their obligations to their patients, and thus that they are not permitted to let considerations of self-interest intrude. Profit-maximization has precisely the same status for managers. To my knowledge, no one has ever tried to defend the managers of RJR-Nabisco, or Enron, on the grounds that they were simply acting in their own self-interest. Of course, if the incentive systems have been properly designed, managers will find it to be in their interest to maximize shareholder value (in the same way that doctors generally find it to be in their interest to cure their patients). But this is accidental and irrelevant from the moral point of view. In the case of a conflict, the obligations simply trump the relevant set of interests. Where things get interesting is when multiple obligations conflict, as in the case of a doctor who can improve a patient’s chances of survival by lying to him about his condition, or of a manager who finds herself able to please investors by initiating an unnecessarily severe downsizing.

The second major source of confusion stems from the moral status of the objective sought by managers—profit-maximization. The doctor's obligations to the patient flow quite naturally from the objective, which is to restore the patient to health. Health is widely regarded as a good thing, and thus the doctor's actions serve to promote a state of affairs that is morally desirable. This makes the doctor's actions directly justifiable, even intrinsically altruistic. Things are more complicated in the case of business. It is not clear that profits are intrinsically good. Furthermore, when a manager makes a decision that disadvantages workers in order to benefit owners, the profit-maximization imperative generates a distributive transfer that is by no means morally sanctioned. In fact, under the typical set of circumstances, the transfer will be regressive, and thus problematic from the moral point of view.

The asymmetry arises from the fact that profit-maximization is only indirectly justified. It is useful to note that this problem is one that business ethics shares with legal ethics. The adversarial trial system imposes upon lawyers an obligation to do whatever is in their power to defend or advance the interests of their client, even when these interests are highly refractory to the concerns of justice. Thus, the professional obligations of lawyers often conflict with the imperatives of everyday morality. What justifies their behavior is the fact that they operate in the context of an institution with differentiated roles. The desirable outcome is a product of the interaction between individuals acting in these roles, none of whom are actually seeking that outcome. Justice is best served when there is both vigorous prosecution and vigorous defense.

Thus the effective trial lawyer "promotes an end which is no part of his intention." The adversarial system may, for example, maximize acquittal of the innocent, even though neither the prosecution nor the defense adopts that as their objective. As a result, neither lawyer's conduct can be justified by the intended outcome. It is justifiable only through the consequences that the pursuit of this outcome leads to, when combined with the actions of the others.

The same can be applied to the case of managers. The manager should seek to maximize profits for the same reason that the defense lawyer should seek to have his client acquitted—not because the acquittal of his client would be a good thing, or even because his client wants to be acquitted and is paying the bill, but rather because the adversarial trial system as a whole is taken to be the best form of institutional arrangement to serve its appointed function. This is why one cannot do legal ethics without a broader appreciation of how the legal system as a whole functions, and what valuable tasks the various roles are thought to discharge. Similarly, one cannot do business ethics without some appreciation of what justifies the system of private enterprise.

Thus the straightforwardly moralizing critique of the profit motive is jejune (comparable to attacking lawyers for "defending rapists and murderers"). We need to understand why criminals should be entitled to the best possible defense, in order to understand the responsibilities of lawyers. Similarly, we

need to understand why corporations should be entitled to pursue profits, in order to understand the responsibilities of managers.

1.2. What Justifies Profit?

The right of corporations to earn profits is sometimes regarded as self-evident. This conviction usually stems from a set of broadly Lockean convictions, which suggest that individuals come naturally equipped with a set of property rights prior to the institution of government. Profit-maximization is then understood as the attempt to augment these holdings through labor input or voluntary exchange—neither of which the state has any obvious authority to restrict.

The problem with this Lockean view—apart from the fact that the underlying conception of rights is deeply problematic—is that corporations are not individuals, they are highly artificial legal constructs. Furthermore, the corporate organizational form provides individuals with a number of very tangible advantages that they do not enjoy as private citizens. The most significant among these is limited liability—the ability to insulate their own private resources from those of the corporation, so that they cannot be pursued by creditors in the event of default. Because of this, creating a corporation is widely regarded as a privilege, not a right. This makes it legitimate for the state to impose certain obligations, in return for the privileges granted.

Many of the corporations chartered by the state are nonprofit. They are specifically prohibited from showing more than a modest revenue surplus. So why permit an exception for other firms? To put it in Marxian terms, why should society tolerate the private appropriation of the social product?

The answer to this question is somewhat complex. Basically, it is that society wants to encourage competition between suppliers. This competition, when combined with competition between purchasers, will affect the prices at which goods trade. Under the correct circumstances, competition will push prices toward the level at which markets clear (i.e., suppliers will not be left with unsold merchandise, and consumers will not be left with any unsatisfied demands). When this occurs, it means that society has succeeded in minimizing the overall amount of waste in the economy. It means that fewer resources will have been spent producing goods that no one wants, at the expense of goods that people do want.

Thus the primary reason for introducing the profit motive into the economy is to secure the operation of the price mechanism. The price mechanism is in turn valued for its efficiency effects. It allows us to minimize waste. The formal proof of this is often referred to as “the first fundamental theory of welfare economics” (hereinafter FFT), or else, in a nod to Adam Smith, the “invisible hand theorem.” The central conclusion is that the outcome of a perfectly competitive market economy will be Pareto-optimal—which means that

it will not be possible to improve any one person's condition without worsening someone else's.

The importance of the price mechanism is often underestimated. Since the profit orientation of firms definitely has some adverse social consequences, this can sometimes make it difficult to see what the big gains are that justify our tolerance for the various abuses. In order to put things into perspective, it is helpful to consider the difficulties that we would face trying to make decisions in the absence of a set of prices. This is the situation that planners often confronted in the former Soviet Union. Imagine that one of your plants increases its production, so that you now have the capacity to produce an extra 500 tons of plastic. What to do with this material? You need to figure out where it is most needed. But how do you decide? Suppose, to simplify enormously, that there are two possible uses: to make toothbrushes or soup ladles. The question is: which do people need more of?

In a market economy, these needs will be expressed in the form of relative willingness to pay. If stores have too many ladles, and not enough toothbrushes, they will be willing to order more toothbrushes, and pay more for them. This in turn means that the toothbrush makers will be willing to pay more for the plastic. Thus, if all firms sell to the highest bidder, the resources will be channeled toward the use for which there is the greatest need. But if there is not a competitive market for all these goods, not only will firms not have the incentive to engage in the necessary transactions, but the absence of prices will make it difficult for anyone even to determine which transaction should be occurring. Planners in the former Soviet Union used to get around this problem by sometimes looking at commodity prices in Western Europe and North America, and using these figures to do calculations for their own economy. In fact, they used to joke that in the event of a global communist revolution, it might be worthwhile to keep Hong Kong capitalist, so that everyone else would know what prices their goods should be trading at.

The joke has a very serious underlying point. Without prices, you simply cannot organize a complex economy, whether it be capitalist, socialist, or communist. And not just any prices will do. There are an enormous number of price points at which exchanges can occur. In cases where there is only one supplier or one consumer, this gives one side considerable power to dictate terms. Under such conditions, there is no reason to expect that the price level chosen will be the price that clears the market. Thus the price system will not induce efficiency. But when there is more than one supplier, or more than one customer, each one is in a position to undermine the negotiating power of the other. If one supplier insists on a price that is too high, the customer can go to the competition. The competitor is then able to make a profit by undercutting the other one's price, making up for it through a larger volume of sales. The result is a race to the bottom among the suppliers, in which they competitively underbid one another until the market clears, and all profit disappears.

Thus the central rationale for having private profit-seeking firms is to establish competition among suppliers and consumers. This competition drives prices toward market-clearing levels, allowing society in turn to generate a more efficient allocation of its resources and labor time.

It should be noted that this concern with competitive markets, and market-clearing prices, is not simply an abstract philosophical theory about what might justify profit-maximization. The entire legal structure of the firm, along with the regulatory environment, has been organized in such a way as to promote not just competition, but the precise type of competition that is likely to generate market-clearing prices. This is true of everything from antitrust to consumer protection law. In the past decade in Russia, corporations have been known to maximize profit by blowing up each other's factories and assassinating each other's chief executives. Much of the massive legal apparatus that governs corporate behavior in more mature capitalist economies is designed to ensure that firms seek to maximize profits through a much more limited set of strategies—namely, those strategies that are likely to generate more efficient production, along with a more efficient allocation of goods and services in the economy.

Thus, if we ask what the obligations of managers are, the answer can be provided quite directly. The function of the market economy is to produce the most efficient use of our productive resources possible. This can be done, roughly speaking, by achieving the price level at which all markets clear. The role of the firm in that economy is to compete with other suppliers and purchasers for profits in order to drive prices to that level. Thus managers are obliged to do what is necessary in order for the firm to maximize profits in this way. Profits show that the balance of “needs satisfied” to “resources consumed” is positive, while losses show that the resources would have been put to better use elsewhere. Hence the old saying that if we penalize a man for making a profit, we should penalize him doubly for showing a loss.

1.3. Milton Friedman

The approach to business ethics that takes profit-maximization as a central concern is often viewed with suspicion, since it has traditionally been used more as an apologetic for irresponsible behavior than as a platform for a good-faith effort to develop a code of ethics.

As we have seen, in order to be plausible, the profit-maximization approach to business ethics cannot identify profit-maximization with individual utility-maximization on the part of managers. The naïve version of the “invisible hand” view, according to which markets miraculously transform private vices into public virtues, has clearly become obsolete in the era of professional management.

Thus when Milton Friedman argued that the social responsibility of business is to increase its profits, his primary emphasis was on the fiduciary relationship between managers and shareholders (Friedman 1962). The manager is in a similar position with respect to the shareholder that the lawyer is in with respect to a client—he is expected to advance the interests of the principal, not his own. This requires trust, and hence moral obligation, between the two parties. And, of course, there are many ways in which the lawyer can exploit this relationship for private gain, as can the manager.

This makes Friedman's view a genuine code of ethics, and not simply an apologia for self-interest. However, while Friedman is clear that managers are subject to genuine moral constraint, he is less than clear about the source of these obligations or constraints. At one point, he suggests that the manager is bound to assist the shareholder in the satisfaction of his or her desires, and that profits just happen to be what most shareholders want. This is clearly absurd—the manager is not the personal servant of the shareholder. The shareholder might like to have the manager do his laundry, and if he can supply appropriate incentives, he may even succeed in getting the manager to do it. But there is no sense in which the manager is morally obliged to do so, by the mere fact that the shareholder desires it. The manager's responsibility toward the shareholder is clearly restricted to the latter's investment returns. Or, as Friedman puts it when he is being careful, the responsibility of managers is "to make as much money for their stockholders as possible" (Friedman 1962: 133).

However, even this more restricted concept of managerial responsibility is not enough to explain the source of the obligation. Simply making a promise is not enough to generate an obligation, in cases where the end in view is itself not morally justifiable. Promising to help a friend rob a bank does not generate an obligation to rob the bank. Thus the manager's obligation to help the shareholder maximize profits must be derivative of the latter's entitlement to do so. And since it is the FFT that justifies this entitlement, Friedman's argument derives managerial responsibilities from the efficiency argument for capitalism on the whole.¹

This implicit dependence upon the FFT is discernible in a seemingly innocuous caveat that Friedman tacks onto the formulation of his central thesis. Here is what he says:

The view has been gaining widespread acceptance that corporate officials and labor leaders have a "social responsibility" that goes beyond serving the interests of their stockholders or their members. This view

¹ Friedman also has a parallel argument concerning the role of markets in promoting freedom. But this line of thinking is, in my view, so riddled with fallacies that it does not merit serious consideration. Furthermore, it seems fairly obvious that Friedman's preference for market solutions to almost every social problem came from his conviction that governments were inefficient and markets were efficient.