



EDITED BY

GEORGE F.

DeMARTINO

DEIRDRE N.

McCLOSKEY

≡ The Oxford Handbook of
PROFESSIONAL
ECONOMIC ETHICS

THE OXFORD HANDBOOK OF

**PROFESSIONAL
ECONOMIC
ETHICS**

THE OXFORD HANDBOOK OF

PROFESSIONAL

ECONOMIC

ETHICS

Edited by
GEORGE F. DEMARTINO
and
DEIRDRE N. MCCLOSKEY

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 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, United States of America

© Oxford University Press 2016

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
The Oxford handbook of professional economic ethics / edited by George F. DeMartino
and Deirdre N. McCloskey.

pages cm

Includes bibliographical references and index.

ISBN 978-0-19-976663-5 (alk. paper)

1. Economics—Moral and ethical aspects. 2. Economists—Professional ethics.

I. DeMartino, George F. II. McCloskey, Deirdre N.

HB71.O854 2015

174'.4—dc23

2015005462

1 3 5 7 9 8 6 4 2

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

*For the many virtuous economists who stay up nights worrying
about the nature and consequences of their professional practice.*

CONTENTS

<i>Foreword</i>	xiii
WILLIAM EASTERLY	
<i>Acknowledgments</i>	xvii
<i>List of Contributors</i>	xix

PART I INTRODUCTION

1. Introduction, or Why This Handbook?	3
GEORGE F. DEMARTINO AND DEIRDRE N. MCCLOSKEY	

PART II UNCERTAINTY, RISK AND PROFESSIONAL ECONOMIC ETHICS

2. The <i>Skin-in-the-Game</i> Heuristic for Protection Against Tail Events	13
NASSIM NICHOLAS TALEB AND CONSTANTINE SANDIS	
3. The Ethics of Economic Decision Rules	29
SVEN OVE HANSSON	
4. In Praise of Imperfect Commitment: An Ethic of Power, Professionalism, and Risk	55
SHARON D. WELCH	
5. “Econogenic Harm”: On the Nature of and Responsibility for the Harm Economists Do as They Try to Do Good	71
GEORGE F. DEMARTINO	

PART III THE ETHICAL NATURE OF ECONOMIC PRACTICE

6. About Doing the Right Thing as an Academic Economist	101
ERWIN DEKKER AND ARJO KLAMER	

7. The Social Responsibility of Economists 116
PETER J. BOETTKE AND KYLE W. O'DONNELL
8. The Ethical Economist: Duty and Virtue in the Scientific Process 137
JONATHAN B. WIGHT

PART IV THE ETHICAL ENTAILMENTS OF ECONOMIC THEORY

1. *GENERAL ISSUES*

9. Ethics in Relation to Economics, Ecology, and Eschatology 169
HERMAN DALY
10. Poisoning the Well, or How Economic Theory Damages
Moral Imagination 184
JULIE A. NELSON
11. Economists' Odd Stand on the Positive-Normative
Distinction: A Behavioral Economics View 200
JOHN B. DAVIS
12. The Complex Ethical Consequences of "Simple" Theoretical Choices 219
ROBERT H. FRANK
13. Good, Evil, and Economic Practice 230
TOMÁŠ SEDLÁČEK

2. *ECONOMIC THEORY AND THE GREAT RECESSION*

14. Alternative Ethical Perspectives on the Financial Crisis:
Lessons for Economists 249
IRENE VAN STAVEREN
15. Economists' Ethics in the Build-Up to the Great Recession 268
ROBERT H. WADE

PART V ETHICAL ISSUES IN ECONOMIC RESEARCH

1. *EXPERIMENTAL ECONOMICS*

16. Ethics and Advances in Economic Science: The Role of Two Norms 297
JINGNAN CHEN, ANGELINA CHRISTIE, AND DANIEL HOUSER
17. The Meaning of *Deceive* in Experimental Economic Science 317
BART J. WILSON

2. *ECONOMETRICS*

18. Honesty and Integrity in Econometrics 329
THOMAS MAYER
19. Lady Justice Versus Cult of Statistical Significance: Oomph-less Science and the New Rule of Law 352
STEPHEN T. ZILIAK AND DEIRDRE N. MCCLOSKEY

3. *FIELD RESEARCH*

20. Balancing Risk and Benefit: Ethical Tradeoffs in Running Randomized Evaluations 367
RACHEL GLENNERSTER AND SHAWN POWERS
21. Conducting Ethical Economic Research: Complications from the Field 402
HAROLD ALDERMAN, JISHNU DAS, AND VIJAYENDRA RAO
22. The Unprincipled Randomization Principle in Economics and Medicine 423
STEPHEN T. ZILIAK AND EDWARD R. TEATHER-POSADAS

4. *CONFLICT OF INTEREST*

23. Professional Disequilibrium: Conflict of Interest in Economics 455
DENNIS F. THOMPSON
24. Considerations on Conflict of Interest in Academic Economics 475
JESSICA CARRICK-HAGENBARTH AND GERALD EPSTEIN

PART VI ETHICAL ISSUES IN APPLIED ECONOMICS

1. DEVELOPMENT

- | | |
|---|-----|
| 25. Ethics, Economic Advice, and Economic Policy
JOSEPH E. STIGLITZ | 495 |
| 26. Neoclassical Economics as the New Social Engineering:
The Debacle of the Russian Post-Socialist Transition
DAVID ELLERMAN | 520 |
| 27. The Ethics of Economic Development and Human Displacement
DES GASPER | 534 |
| 28. How Can We Better Address the Gaps in Our Knowledge About
Development Effectiveness?
MARTIN RAVALLION | 558 |

2. ECONOMIC ADVISING IN GOVERNMENT AND BEYOND

- | | |
|--|-----|
| 29. Confessions of a Policy Analyst
ROBERT H. NELSON | 587 |
| 30. Ethics and the Government Economist
SUSAN OFFUTT | 618 |
| 31. The Ethics Problem: Toward a Second-Best Solution to the Problem
of Economic Expertise
DAVID M. LEVY AND SANDRA J. PEART | 635 |
| 32. First Tell No Untruth
ALAN FREEMAN | 651 |

3. FORENSIC ECONOMICS

- | | |
|---|-----|
| 33. Ethical Issues in Forensic Economics
ROBERT J. THORNTON AND JOHN O. WARD | 671 |
|---|-----|

PART VII ETHICAL ISSUES IN ECONOMIC EDUCATION

34. Exposure and Dialogue Programs in the Training of Development
Analysts and Practitioners 697
RAVI KANBUR
35. Ethics and Learning in Undergraduate Economics Education 714
ROBERT F. GARNETT, JR.

PART VIII LOOKING AHEAD

36. Creating Humble Economists: A Code of Ethics for Economists 737
DAVID COLANDER
37. Codes of Ethics for Economists, Pluralism, and the Nature
of Economic Knowledge 750
SHEILA C. DOW
- Author Index* 765

FOREWORD

DISCUSSING ethics makes economists uncomfortable, for some good reasons and some bad ones. Economics is ill-suited to discuss ethics in some ways, but very well-suited in others.

The best reason not to discuss ethics in economics is that we do not want to turn our debates into dueling ad-hominem attacks about unethical motivations for our arguments. Like other professionals such as physicians, we economists view ourselves as motivated by professional norms and not by crude self-interests that could be unethical. Moreover, we know that even if there are unethical motivations for an argument, this is not sufficient to disprove the argument—it still needs the kind of debate on theory and evidence that would have happened anyway if we had ignored ethics.

Nor do we want to allow declarations in economic debates that one position is morally superior to another. Economics is traditionally viewed as ill-suited to discuss competing ethical norms.

Yet economics is well-suited in other ways to address the very same concerns about discussing ethics. Whether we have a reputation for professionalism is not under our control—it is up to our audience to decide. We can do a better job convincing our audience of our professionalism by discussing self-interest openly rather than ignoring it. Our models of human behavior give insights into the need for professional norms. For example, sellers of a good whose quality is imperfectly observable will get a better price if they can convince buyers that the sellers obey ethical norms of not cheating on quality. The existence of a strong ethical norm that sellers don't cheat, along with some enforcement mechanism, would actually lead to a better outcome for sellers. It is economics that gives the paradoxical insight that your self-interest sometimes requires convincing everyone else that you are not responding too much to self-interest.

The belief that economics cannot adjudicate competing moral values is also overstated. Our normative analysis of what makes people better off is impossible without presuming some ethical positions. Take, for example, the concept of revealed preference: that if you chose A over B, then A must make you better off than B. Revealed preference analysis presumes that the individual does and should have the right to make her own choices and that she is the best judge of her own well-being. Most of our models make the same assumptions. Individual choice and autonomy are ethical values over which there has long been and still is a huge worldwide debate. If economists tried to do revealed preference analysis in a situation where there was actually coercion rather than

choice, we should certainly take notice—and so we cannot avoid the global debate on individual choice.

So how should we economists confront our own ethical behavior and values? On our core ethical values of individual autonomy and consent, it seems straightforward. Whenever we apply models or analysis based on individual choice, shouldn't we also analyze whether individuals are indeed free to choose? There may be disagreement on how or whether to do this, but let's air those disagreements rather than avoid the topic of ethical values altogether.

The question of whether there is a profound tension between our professional norms and our self-interest deserves careful attention. Conflict of interest in economics gained much (unwanted) attention after the documentary *Inside Job* accused some finance economists of doing analysis favorable to financial industry interests while receiving undisclosed pay from those same interests. Even if you believe, as I do, that *Inside Job* was unfair to some of its targets, it did fuel a crisis of confidence in economists that we all have a strong interest in correcting. The response has been to strengthen the norms that we disclose possible conflict of interests in our research and policy recommendations; this is surely a good thing. An example from my own field of development is that researchers on foreign aid should disclose whether they are employees of or consultants to agencies dispensing foreign aid (or conversely, recipients of funding from anti-aid interests).

Yet the issue of conflict of interest is too complex to be so quickly dismissed by a simple disclosure requirement. A lot of attention has been focused on conflict of interest in academic research, where the norms were already fairly strong and were at least partly self-enforcing through reputational incentives. Conflict of interest is potentially more severe and the existing norms weaker in some of economists' other roles as government officials, as advisors to governments or international organizations or political candidates, as critics or advocates of political ideologies, or as public intellectuals shaping policy debates. Sometimes the conflict of interest is obvious and easily resolved, other times not so much. For example, it will be obvious that an economist advising a political party or serving as a political appointee has a conflict of interest in research or policy recommendations that support or oppose that party's positions. The economist should obviously disclose any such relationship to political parties or causes if it is not already public knowledge.

But the political role of economists brings with it a subtler ethical need for accountability. Again, research in our own field gives us some insight on this. We have a lot to say about how bad it is for agents to be able to play with other people's money. To return to finance examples, we know a lot about how bad it is for banks to socialize risk through taxpayer bailouts while keeping any rewards from those risky investments. Analogously, economists making risky policy recommendations allow the cost of our advice to fall on everyone else, with only weak individual accountability for economists ourselves. There are still hardly any positive or negative individual consequences of good or bad advice for the dispensers of advice. In this environment, we should consider the ethical implications of our overpromising on the payoffs to our recommendations or overstating the certainty with which our promises will be realized.

There is the hope, of course, that economists affiliated with opposing parties will tend to subject each other's overpromising and excess certainty to critical scrutiny on the traditional grounds of theory and evidence. Just as we think competition is a positive force in markets, competition among partisan economists should be a positive ethical force in keeping the debate healthy.

Yet sometimes competition is not as sufficient as it is assumed to be. Both sides may have a common interest in overstating the importance of economists' recommendations in general, such as overstating the importance of current policy decisions in any direction. Both sides may want to downplay competing explanations of outcomes from the economics literature, ones that suggest outcomes are also dependent on other forces far beyond the control of the current policy-makers, such as culture or history. So we are left with a possible self-interested bias on the importance of economists as policy advisors, even as we disagree with each other over particular policy proposals.

Economists on opposing partisan sides may also have a common interest in limiting debate to what is politically feasible for the current mainstream parties, avoiding viewpoints that pose a more fundamental challenge to the status quo.

These subtler biases are evident in my own field of economic development. On the status quo bias, development economists as public intellectuals seem to be expected to direct all their efforts to advising the current development establishment of rich country governments, aid agencies, philanthropic foundations, and non governmental organizations on how to end poverty. Even critics (such as this author) focus attention on the mistakes of this establishment and how to correct them. Public intellectuals who simply dismiss this establishment as irrelevant or inherently counterproductive—rightly or wrongly—have far more trouble getting a hearing. So self-interest of public intellectuals militates against fundamental challenges to the status quo, and, indeed, there are few such public intellectuals in the development field.

On the importance bias, public intellectuals in development (including this author) seem to have a common interest in exaggerating the importance and certainty of our own advice on how to end poverty even if we disagree on how to do so. Less “constructive” views that challenge how much we economists really do know are a threat to the status and power of economists as public intellectuals in the development field.

It seems to take a lot of ethical commitment to admit publicly to something so nihilistic and self-destructive as the limits to our knowledge and expertise as economists. Yet such ethical effort could really pay off for the field in unexpected ways. To take finance again, the audience for finance economists wanted them, above all, to give advice on how to find surefire high returns on investments. It took some courage to embrace theories that bluntly state finance economists are useless in helping you beat the market. Yet this freed the most courageous finance economists to develop principles under which financial markets work well for all their participants, such as the earlier admonition against banks socializing risks through unlimited taxpayer bailouts.

In development, there is a principle similar to “you can't beat the market”: if there really is an easy and tangible solution *X* to end global poverty, why didn't *X* end global poverty already? Suppose development economists were ethically obliged to admit to

their audiences an inability to find such easy answers if we do indeed have such inability. This could free us to develop and advocate the principles of a whole problem-solving system in which many different private and public entrepreneurs will discover the more difficult answers to the problems of poverty—as indeed economists (usually the ones not facing so much pressure for immediate and direct solutions) have already been doing for a long time. It could be surprisingly liberating for economists to overcome our importance bias to admit how little everyone actually needs us economists to run the economy.

This brief discussion of ethical issues for economists has chosen an idiosyncratic grab-bag of issues, and this selection and discussion probably reflects this author's own self-interested biases. But what this fascinating volume shows is that the open discussion of professional ethics in economics has great potential to provoke debate on many exciting questions and that there is no reason to continue inexcusable and counterproductive silence on the topic.

William Easterly

ACKNOWLEDGMENTS

IN addition to the contributors to this volume, many individuals provided invaluable support. At the Josef Korbel School of International Studies at the University of Denver, several students, including Matthew Farrell, Colin Gaiser, Eugenia Goh, Minsun Ji, Laura Roberts, Patrick Sutherland, Cecil Shapland and, in particular, Noelle Schaef and Michelle Woodruff, provided the editors with extensive research, editorial, and administrative assistance. The Earhart Foundation awarded a generous grant that allowed contributors to present their work at conferences and funded administrative support and honoraria. Emily Hinga of the Josef Korbel School of International Studies kept the books and managed to ensure that salaries and expenses were paid. Ilene Grabel provided consistent encouragement and advice, while Oxford University Press's Joe Jackson, Scott Parris, Terry Vaughn, and Cathryn Vaulman shepherded the project expertly from its inception through to its completion. Finally, Preethi Sundar at Newgen Knowledge Works provided outstanding editorial and production services.

The editors note with sadness the passing of contributor Thomas Mayer as this book was in press. Tom was an extraordinary person and a wonderful economist who advocated ethical economic practice and, more importantly, modeled it in his own work.

CONTRIBUTORS

Harold Alderman is Senior Research Fellow at the International Food Policy Research Institute (IFPRI), Washington, D.C.

Peter J. Boettke is University Professor of Economics and Philosophy at George Mason University (GMU), as well as the BB&T Professor for the Study of Capitalism, Vice President for Research, and Director of the F. A. Hayek Program for Advanced Study in Philosophy, Politics, and Economics at the Mercatus Center at GMU.

Jessica Carrick-Hagenbarth is a doctoral candidate in the Department of Economics at the University of Massachusetts, Amherst, and a Visiting Instructor in Economics at Middlebury College.

Jingnan Chen is Lecturer in Economics at the University of Exeter Business School.

Angelina Christie is Financial Economist in the Enterprise Risk Analysis Division, Economics Department at the Office of the Comptroller of the Currency, Washington D.C.

David Colander is Christian A. Johnson Professor of Economics at Middlebury College.

Herman Daly is Emeritus Professor in the School of Public Policy at the University of Maryland.

Jishnu Das is Lead Economist, Development Research Group, World Bank, Washington D.C., and Visiting Fellow at the Centre for Policy Research, New Delhi.

John B. Davis is Professor of Economics at Marquette University and Professor of Economics at the University of Amsterdam.

Erwin Dekker is Postdoctoral Fellow with the F. A. Hayek Program for Advanced Study in Philosophy, Politics, and Economics at the Mercatus Center and the Department of Economics at George Mason University.

George F. DeMartino is Professor of Economics and Co-Director of the Global Finance, Trade, and Economic Integration program in the Josef Korbel School of International Studies at the University of Denver.

Sheila C. Dow is Emeritus Professor of Economics at the University of Stirling, and Adjunct Professor of Economics at the University of Victoria.

William Easterly is Professor of Economics at New York University (joint with Africa House) and Co-Director of the Development Research Institute.

David Ellerman is a Visiting Scholar at the University of California, Riverside, and a Fellow of the Center on Global Justice at the University of California, San Diego.

Gerald Epstein is Professor in the Department of Economics and Co-Director of the Political Economy Research Institute (PERI) at the University of Massachusetts, Amherst.

Robert H. Frank is Henrietta Johnson Louis Professor of Management and Professor of Economics at Cornell's Johnson Graduate School of Management.

Alan Freeman is Co-Director of the Geopolitical Economy Research Group at the University of Manitoba, Canada.

Robert F. Garnett, Jr. is Chair and Professor in the Department of Economics at Texas Christian University.

Des Gasper is Professor of Human Development, Development Ethics, and Public Policy in the International Institute of Social Studies at Erasmus University Rotterdam, The Netherlands.

Rachel Glennerster is Executive Director of the Abdul Latif Jameel Poverty Action Lab (J-PAL) at Massachusetts Institute of Technology, as well as Scientific Director of J-PAL Africa and Co-Chair of J-PAL's education sector.

Sven Ove Hansson is Professor of Philosophy and Head of the Division of Philosophy at the Royal Institute of Technology, Stockholm.

Daniel Houser is Chairman and Professor in the Department of Economics and Director of the Interdisciplinary Center for Economic Science at George Mason University.

Ravi Kanbur is T. H. Lee Professor of World Affairs, International Professor of Applied Economics and Management, and Professor of Economics at Cornell University.

Arjo Klammer is Chair and Professor of Cultural Economics at Erasmus University Rotterdam, The Netherlands.

David M. Levy is Professor of Economics, Research Associate at the Center for Study of Public Choice, and Co-Director of the Summer Institute for the Preservation of the History of Economics at George Mason University.

Thomas Mayer was Emeritus Professor of Economics at the University of California, Davis, until his death in January 2015.

Deirdre N. McCloskey is UIC Distinguished Professor of Economics and History, Professor of English, and Professor of Communication, Emerita, at the University of Illinois at Chicago.

Julie A. Nelson is Professor of Economics at the University of Massachusetts, Boston, and a Senior Research Fellow at the Global Development and Environment Institute, Tufts University.

Robert H. Nelson is Professor in the School of Public Policy at the University of Maryland.

Kyle W. O'Donnell is a doctoral candidate at the Department of Economics at George Mason University, and a Visiting Lecturer of Economics at James Madison University.

Susan Offutt is retired Government Accountability Office (GAO) Chief Economist, US Congress Washington, D.C.

Sandra J. Peart is Dean of the Jepson School of Leadership Studies at the University of Richmond and Co-Director of George Mason University's Summer Institute for the Preservation of the History of Economics.

Shawn Powers is a Young Professional in the Education Global Practice at the World Bank, and was formerly a Senior Policy Manager at the Abdul Latif Jameel Poverty Action Lab (J-PAL).

Vijayendra Rao is Lead Economist of the Development Research Group at the World Bank.

Martin Ravallion is the Edmond D. Villani Professor of Economics, Georgetown University, and President of the Society for the Study of Economic Inequality.

Constantine Sandis is Professor of Philosophy at the University of Hertfordshire.

Tomáš Sedláček is a Czech Economist and University Lecturer at Charles University, Prague. He is the Chief Macroeconomic Strategist at ČSOB, a former member of the National Economic Council of the Czech Republic, and an economic advisor to former President Václav Havel.

Joseph E. Stiglitz is an American economist and a professor at Columbia University. He is also the co-chair of the High-Level Expert Group on the Measurement of Economic Performance and Social Progress at the OECD, and the Chief Economist of the Roosevelt Institute. A recipient of the Nobel Memorial Prize in Economic Sciences (2001) and the John Bates Clark Medal (1979), he is a former senior vice president and chief economist of the World Bank and a former member and chairman of the (US president's) Council of Economic Advisers. In 2000, Stiglitz founded the Initiative for Policy Dialogue, a think tank on international development based at Columbia University.

Nassim Nicholas Taleb is Distinguished Professor of Risk Engineering at New York University's Polytechnic Institute.

Edward R. Teather-Posadas is a doctoral candidate in the Department of Economics at Colorado State University.

Dennis F. Thompson is Professor of Government and Alfred North Whitehead Professor of Political Philosophy in the Faculty of Arts and Sciences, and Professor of Public Policy in the John F. Kennedy School of Government, Harvard University. He is also the founding Director of the University Center for Ethics and the Professions (now the Edmond J. Safra Center for Ethics).

Robert J. Thornton is Charles William Macfarlane Professor of Economics at Lehigh University and Associate Director of the Martindale Center for the Study of Private Enterprise.

Irene van Staveren is Professor of Pluralist Development Economics in the International Institute of Social Studies at Erasmus University Rotterdam, The Netherlands.

Robert H. Wade is Professor of Political Economy and Development in the Department of International Development at London School of Economics.

John O. Ward is Professor Emeritus in the Department of Economics, University of Missouri-Kansas City and Editor Emeritus, *Journal of Forensic Economics*.

Sharon D. Welch is Provost and Professor of Religion and Society at Meadville Lombard Theological School (Unitarian Universalist), Chicago.

Jonathan B. Wight is Professor of Economics and International Studies in the Robins School of Business at the University of Richmond.

Bart J. Wilson is Donald P. Kennedy Endowed Chair and Professor of Economics and Law at the Economic Science Institute, Chapman University.

Stephen T. Ziliak is Professor of Economics at Roosevelt University, Chicago, and Faculty Member of the Angiogenesis Foundation, Cambridge, MA.

PART I

INTRODUCTION

CHAPTER 1

INTRODUCTION, OR WHY THIS HANDBOOK?

GEORGE F. DEMARTINO
AND DEIRDRE N. McCLOSKEY

Caminante, no hay camino. Se hace el camino al andar.
(Searcher, there is no road. We make the road by walking.)

—Antonio Machado, *Selected Poems*

ACADEMIC handbooks typically showcase the range of existing work defining mature fields of research. Handbooks catalogue what is taken to be representative, the best, or otherwise most important work in the field while pointing toward new areas of inquiry. They provide an opportunity to take stock, to map the field, and to identify the outstanding questions that deserve further attention.

Then why a handbook of professional economic ethics, of all things, when in fact there is no such field? You can look it up. Since its founding in 1885 and up to the present, the American Economic Association (AEA)—perhaps the most influential professional association of economists in the world—has never cultivated inquiry into professional ethics: what does it mean to be an ethical economist, or, what would it mean for economics to be an ethical profession? Never . . . as in, not once. The AEA's founders were on a mission to promote the influence of economists in public policy. Infected by the spirit of American Progressivism, they believed that professionals had a duty to promote the social good through the application of expertise to social and economic problems. Their ethical obligations began and ended with that duty. *Influence for economic experts*—that was the goal. The associated ethical burdens were taken to be trivial and obvious.

By the end of the 19th century, leading economists had concluded that achieving influence required the transformation of economics into a positivist science. They also believed, as many economists believe down to the present, that science required excision of all ethics talk from the field. Out with Richard T. Ely and his band of foggy-headed

social gospel ministers who had founded the AEA; in with Charles F. Dunbar and the men of objective science who would achieve for the profession the standing it deserved.

The AEA is not alone in neglecting the professional ethics of economists. Of the several dozen economic associations existing today in the United States and among the leading associations in Europe and beyond, only a handful have promoted investigation of the ethical aspects of economic practice; indeed, very few even have rules covering disclosure of conflict of interest for their journals.¹ Beyond the associations, it is safe to say that it has never occurred to the majority of economists and certainly to most academic economists that their work involves complex ethical entailments warranting attention.

The goal of the Handbook surely can't be cartographic, then, since there is no terrain out there waiting to be mapped. The purpose is more audacious. We seek to lay down first tracks across a wilderness that we hope will evolve into a network of pathways through a cultivated field of inquiry. We have not sought to restrict our authors' destinations or the direction they take in getting there. Instead, we invited contributors to the volume to help us make the road by walking, *al andar*. Our method was simple: we tasked leading thinkers within economics and beyond to ruminate on those aspects of economic practice that they view as ethically fraught. We gave contributors free rein, holding no preconceptions about what kinds of conversations, questions, arguments, concerns, methods, and styles do and do not count as professional ethics. We hope readers will engage the essays herein with the same spirit of open-mindedness so that, together, we can generate the new field of professional economic ethics. The field is vital to our profession. But it is also and even more acutely vital to those whom economists purport to serve.

THE ECONOMIST'S ERROR: PROFESSIONAL ETHICS IS NOT A CODE OF CONDUCT

As a rule, economists are not enamored with the idea of The Profession. Chicago School economists in particular, and others, too, think of professions as forms of labor unions to protect the privileges associated with monopoly power (Friedman, 1962; and the Chicago oral tradition). The expectation is that as soon as an occupation claims for itself the status of a "profession," it begins to seek government protection in the form of certification, licensing, and other mechanisms restricting access to the market—all the while

¹ The notable exceptions include, in the United States, the National Association of Forensic Economics (NAFE) and the American Academy of Economic and Financial Experts (AAEFE); and, in Sweden, two associations of applied economists: Akademikerförbundet SSR and Civilekonomerna. Since 2011, a number of leading economic associations in Britain, France, Germany, and the United States have considered and some have adopted new conflict-of-interest disclosure rules for the journals they publish. But the new rules generally do not reach beyond disclosure to consider other ethical issues.

purporting to protect the public interest. In Friedman's words, "registration, certification, or licensure, almost inevitably becomes a tool in the hands of a special producer group to obtain a monopoly position at the expense of the rest of the public. There is no way to avoid this result" (1962: 148).

The economist's take on professional *ethics* is equally unflattering. The received view is that professional ethics entails a code of conduct that imposes on professionals and their clients the values that The Profession deems best. A professional code of conduct turns consumer sovereignty on its head—it puts the suppliers of a service in charge of just what they will sell (and who can sell it and under what terms), thereby circumscribing the freedom of consumers to satisfy their preferences for the particular services they think they need. The competition among suppliers that would otherwise protect consumers' interests is replaced by a producer cartel that dictates the terms of the transaction. Making matters worse, codes may induce unwarranted public trust in the profession, paving the way for scoundrels to profit from public naiveté. Leland Yeager (1976: 569) put it best:

The more prevalent and well-based is the belief that people are generally decent and honest, the greater is the chance that culprits have to benefit from the presumption that they too have these virtues. They will enjoy a free ride on, while posing unfair competition with, the warranted credibility of other people.

There is some value in the economist's skepticism about the professions and codes of conduct. As good and loyal economists, we share it—although we would note, too, that a supplier has a legitimate interest in maintaining the quality of her product. But whenever an organized interest claims to know and to do what's best for you, watch out! Yet economists have made a basic yet consequential error that has stunted the development of ethical awareness in the profession. The error is to think of professional ethics as a tool to prevent the morally compromised from doing bad things. In the economist's view, professional ethics must entail ethical *legislation*, especially legalistic codes of conduct that are backed by sanctions in the form of certification or licensure. In other words, the economist thinks of codes as budget constraints, in the style of the "rules of the game" that Douglass North speaks of.

Properly understood, professional ethics is nothing of the sort. It is a conversation rather than a constraint, a dance rather than a pose. Professional ethics refers to a *tradition of critical inquiry* into the myriad questions that surround professional engagement. It entails reflection and debate, literatures, curriculum, texts, journals, and bulletins probing the nature of professional practice for its ethical substance. Professional ethics is intended to assist well-meaning professionals, and the profession as a whole, to do good—not to prevent charlatans from doing harm (DeMartino, 2011).

Professional ethics is an *activity*, not a rigid institution. It occurs in a profession only when reflective practitioners, ethicists, clients, and other members of the public wrestle with questions such as: just what ethical duties, obligations, and challenges attend

the profession's practice? What virtues ought its practitioners and their clients cultivate in themselves and their peers? How can individual practitioners and the profession as a whole undertake their work responsibly? Defined in this way, very few professions today can claim professional ethics—although almost all proudly boast on their websites about codes of conduct. Economics must follow law, medicine, journalism, and the handful of other professions that struggle continuously with the ethicality of their practice. *Education, articulation, elucidation, and aspiration, not regulation, legislation, or condemnation*—these are the watchwords of the field that the Handbook seeks to announce and launch.

THE CASE FOR PROFESSIONAL ECONOMIC ETHICS

But why the need for professional economic ethics when the profession has done quite well over the past century, thank you very much, without it? Surely there are risks associated with establishing the new field. Given the profession's proclivity to separate "us" from "them"—the mainstream from the backwater, the orthodoxy from the heterodoxy, Chicago from Cambridge, George Mason from the University of Massachusetts—isn't there a real danger that professional ethics might become the next weapon in the fight over who is and who is not a legitimate economist? *Preach free trade (or the raising of the minimum wage) or be damned as unethical!*

The case for economic ethics is simple and, we think, undeniable. Economists enjoy tremendous influence today over the life chances of others—innumerable others. That is the heart of the matter. The influence of economists arises from their expertise in a field vital to social well-being, freedom, and other valued goals. As economists know better than anyone, when you monopolize a resource that others need, you exert power over them. Moreover, in recent years, economists' influence has been amplified by institutional developments. Independent central banks, the multilateral development banks, and other international financial institutions are often in a position to set economic policy and even engage in social engineering without much oversight by elected officials or the public. Economists are at the helm of such institutions and occupy staff positions in the departments where the actual work gets done. Combined with its intellectual monopoly, institutional power enhances the ability of the economics profession to alter the course of human affairs—for the better, of course, but also, sometimes, for the worse.

Why also for the worse? In all but the simplest economies, even well-designed economic interventions impose harm on some agents. Sadly, Pareto improvements rarely present themselves, at any rate not with actual side payment. Most reform instead involves benefits to some and foreseeable costs to others. Indeed, as one of us argues (DeMartino, 2016), "econogenic harm," or the harm that follows from

economists' interventions, is ubiquitous. The grander the economic intervention, the greater the scale of the project, the larger the number of losers and the greater the size of the aggregate losses. Large-scale development projects are a notable example (Kanbur, 2003; see Gasper, 2016). The regrettable fact is that the practice of economics harms even as it helps. That is the tragedy of economics. Economists are in the harm business. Economists usually know it and teach it every day. But what they don't consider carefully enough, or teach adequately, are the ethical entailments of the harm their practice induces.

There is a response available to our concern about econogenic harm. In the style of John Harsanyi, followed by Buchanan and Tullock, and then outside economics by John Rawls, one can argue that the better level of decision is "constitutional." That is, the question can be what sort of society we want to live in—a society in which every Kaldor-Hicks gain is second-guessed as entailing econogenic harm or a society in which market-tested innovations that promote the social good go forward, even though blacksmiths and buggy-makers are thereby unemployed? But such a reply should not be taken as a smart-aleck banishing of ethics. It, too, requires us to "struggle continuously with the ethicality of our practice."

Foreseeable harms can be damaging enough, to be sure, as foreseeable benefits can be uplifting. But in economics (and, indeed, in other professions), there are the unforeseen and unpredictable costs and benefits complicating professional practice. Economists operate on a terrain of *epistemic insufficiency* (McCloskey, 1990; DeMartino, 2013). They don't today and probably can't ever know enough to ensure the success of the policies and other interventions they advocate. Economies are complex whereas knowledge is partial, tacit, and widely dispersed. Economic policy entails predicting what will be the impact tomorrow (and next year and in 10 years) of an intervention undertaken today—and that is something that defies the capabilities of even the brightest economists in possession of the best-calibrated models. Making matters worse, there can be and generally is substantial slippage between the heady, pristine moment of policy design and the messiness of implementation. The poet said, "Between the conception and the creation/Falls the shadow." Economists are just one input in the drama of governance. In the policy domain, it is politics and not economic science that is in charge. The intrepid economist who risks advocacy quickly learns that, at best, she has influence but little ultimate control. And, in the absence of control lies the risk that the economist can do much harm as she tries to do good.

Influence over the lives of others, which can be immense, coupled with the risk of doing even substantial foreseeable and unforeseeable harm, implies that economic practice is ethically fraught. And yet the profession largely manages to ignore the attending burdens. Perhaps because economists understand that harm is universal in economics, the Hippocratic tradition appears to offer no insight into how economists should comport themselves. What does "do no harm" mean in a world where there are no free lunches and where all actions (including doing nothing) entail tradeoffs? And perhaps because economists often paint on big canvases, where they affect the lives of thousands or even millions of people all at once rather than individual clients one by one, clinical

ethics seems largely irrelevant. The scale of economic interventions generates among economists a fear that serious and open engagement with professional ethical issues would paralyze them with doubt in those moments of human need when what is called for instead is focused audacity.

In our view and the view of many of the contributors to the Handbook, the implicit professional pact among economists to repress rather than engage the ethical concomitants of their work is ethically indictable. When the lives of so many are at stake, professional privilege and convenience cannot trump professional duties. Economists have secured the influence we've sought; hurrah for us! The price to be paid for our privilege is open acknowledgment and careful examination of the ethical challenges of our work—in academia, government service, and the myriad other institutions and forums where we ply our craft.

BUT ISN'T ETHICS JUST A MATTER OF SUBJECTIVE PREFERENCE AND POWER?

Economists and calculators have long led the attack by the new clerisy on preaching the virtues.² The late Marc Blaug (1980: 132–133), for example, in many other respects a surprisingly sensible member of his profession, asserted that “There are no . . . methods for reconciling different normative value judgments—other than political elections and shooting it out at the barricades.” By “methods for reconciling,” he appears to mean “air-tight proofs such as the Pythagorean Theorem.” Neatness reigns. The sort of amiable, casuistic reasoning—together that many in the rhetorical tradition recommend and that are practiced in courts of law and other fields of discussion, the trading of more or less good reasons, such as the stories of good or bad lives, ranging from the Hebrew Bible and Plutarch to the latest movie, are spurned by such a theory.

Schumpeter of Vienna and Harvard had earlier expressed an ethical philosophy similar to Blaug's: “We may, indeed, prefer the world of modern dictatorial socialism to the world of Adam Smith, or vice versa, but any such preference comes within the same category of subjective evaluation as does, to plagiarize Sombart, a man's preference for blondes over brunettes” (Schumpeter, 1954: 330). Hurrah-boo. Thus also, Lionel Robbins of the London School of Economics: “If we disagree about ends it is a case of thy blood against mine—or live and let live, according to the importance of the difference, or the relative strength of our opponents. . . . If we disagree about the morality of the taking of interest. . . . then there is no room for argument” (Robbins, 1932: 134).³ And a fount

² Parts of what follow come from various of McCloskey's writings, such as *Knowledge and Persuasion in Economics* (1994).

³ Amartya Sen (1987) says that such a view was “quite unfashionable then.” Not, we think, among the reigning fashionistas of 1932.

of this attitude, again, Bertrand Russell (1977: 136; quoted in Perry): “As to ultimate values, men may agree or disagree, they may fight with guns or with ballot papers, but they cannot reason logically.”

Such a theory, to which many economists subscribe, is called by philosophers “emotivism.” Emotivism was believed by very many early 20th-century people, some under the influence of logical positivism, some under the influence of a falling away from religious faith. It is “the doctrine that all evaluative judgments and more specifically all moral judgments are *nothing but* expressions of preference” (MacIntyre, 1981: 11, italics in original). Or, as Hobbes wrote in 1651, “‘Good’ and ‘evil’ are names that signify our appetites and aversions (1982/1651: Part I, ch. 15, 82).” Emotivism, observe again, taken as a doctrine one *should* believe is, of course, self-contradictory because preaching against preaching is preaching. But logic is not the strong point of logical positivism or of those who have fallen away from religious faith.

Undergraduates and many of their professors become uneasy and start giggling when an ethical question arises. They regard such questions as having mainly to do with sex—thank you fundamentalists of the late 20th century—or with unargued authority, such as the Baltimore Catechism and the nuns who enforce it. The agreement to disagree that ended the wars of religion in Europe can be traced in their unease and in their stock remarks expressing it: “That’s just a matter of opinion,” “Religion should not be mentioned in polite conversation,” “If we disagree about ends it is a case of thy blood against mine,” “The only methods for reconciling different normative value judgments are political elections or shooting it out at the barricades.” According to the emotivists, to be caught making ethical statements is to be caught in meaningless burbling. Shame on you.

We can do better than such obsolete philosophy. We can have a serious conversation about the goods and bads of our practice. We need to. Let the emotivists then make their case in the new agora we and our fellow contributors to the Handbook intend to construct. But be warned: in that new intellectual space, sneering won’t do.

WE KNOW . . .

Yes, of course: paying attention to the ethical issues might substantially complicate the professional life of the economist. Yes, there are opportunity costs here. And yes, there are no guarantees that the profession would do much better in its mission to promote the public good were it to invest the time and energy necessary to construct the field of professional ethics. But objections such as these do not undermine the case for professional ethics. What they do is warn us against ethical simple-mindedness, and they alert us that good intentions will not suffice. They prepare us for what will be a daunting and risky undertaking. But what gain is there, we ask our fellow economists, without risk?

REFERENCES

- Blaug, M. 1980. *The Methodology of Economics: Or, How Economists Explain*. Cambridge: Cambridge University Press.
- DeMartino, G. F. 2016. “‘Econogenic Harm’: On the Nature of and Responsibility for the Harm Economists Do as They Try to Do Good.” In *The Oxford Handbook of Professional Economic Ethics*, edited by G. F. DeMartino and D. N. McCloskey, pp. 71–97. New York: Oxford University Press.
- DeMartino, G. F. 2011. *The Economist’s Oath*. New York: Oxford University Press.
- DeMartino, G. F. 2013. “Epistemic Aspects of Economic Practice and the Need for Professional Economic Ethics.” *Review of Social Economy* 71(2): 166–186.
- Friedman, M. 1962. *Capitalism and Freedom*. Chicago: University of Chicago Press.
- Gasper, D. 2016. “The Ethics of Economic Development and Human Displacement.” In *The Oxford Handbook of Professional Economic Ethics*, edited by G. F. DeMartino and D. N. McCloskey, pp. 534–557. New York: Oxford University Press.
- Hobbes, T. 1982/1651. *Leviathan*. New York: Penguin Classics.
- Kanbur, R. 2003. “Development Economics and the Compensation Principle.” *International Social Science Journal* 2003(55): 175. First published online July 1, 2004. doi: 10.1111/1468-2451.5501019_2
- MacIntyre, A. *After Virtue*. 1981. South Bend, IN: University of Notre Dame Press.
- McCloskey, D. N. 1990. *If You’re So Smart: The Narrative of Economic Expertise*. Chicago: University of Chicago Press.
- McCloskey, D. N. 1994. *Knowledge and Persuasion in Economics*. Cambridge: Cambridge University Press.
- Perry, J. F. “The Dream Hypothesis, Transitions, and the Very Idea of Humanity.” Available at <https://www.bu.edu/wcp/Papers/Teac/TeacPerr.htm>
- Robbins, L. 1932. *An Essay on the Nature and Significance of Economic Science* (second edition, 1935; third edition, 1983). London: Macmillan.
- Russell, B. 1977. *Education and the Social Order*. London: Unwin Paperbacks.
- Schumpeter, J. A. 1954. *History of Economic Analysis*. New York: Oxford University Press.
- Sen, A. 1987. *On Ethics and Economics*. Malden, MA: Blackwell.
- Yeager, L. B. 1976. “Economics and Principles.” *Southern Economic Journal* 42(4): 559–571.

PART II

.....

UNCERTAINTY, RISK
AND PROFESSIONAL
ECONOMIC ETHICS

.....

CHAPTER 2

THE *SKIN-IN-THE-GAME* HEURISTIC FOR PROTECTION AGAINST TAIL EVENTS

NASSIM NICHOLAS TALEB
AND CONSTANTINE SANDIS

AGENCY PROBLEMS AND TAIL PROBABILITIES

THE chances of informed action and prediction can be seriously increased if we better comprehend the multiple causes of ignorance.¹ The study of ignorance, then, is of supreme importance in our individual and social lives, from health and safety measures to politics and gambling (Rescher, 2009). But how are we to act in the face of all the uncertainty that remains after we have become aware of our ignorance? The idea of *skin in the game* when involving others in tail risk exposures is crucial for the proper functioning of a complex world. In an opaque system fraught with unpredictability, there is, alas, an incentive and easy opportunity for operators to hide risk: to benefit from the upside when things go well without ever paying for the downside when one's luck runs out.

The literature of risk, insurance, and contracts has amply dealt with the notion of information asymmetry (see Ross, 1973; Grossman and Hart, 1983*a*, 1983*b*; Tirole, 1988; Stiglitz, 1988) but not with the consequences of deeper information opacity (in spite of getting close, as in Hölmstrom, 1979) by which tail events are impossible to figure out from watching time series and external signs: in short, in the “real world” (Taleb, unpublished), the law of large numbers works very slowly or does not work at all in the time horizon of operators, hence statistical properties involving tail events are completely opaque to the observer. And the central problem that is missing behind the abundant research on moral hazard and information asymmetry is that these rare, unobservable events represent the bulk of

¹ A mathematical version of this paper has appeared in *Review of Behavioral Economics*, 2014, 1: 1–21.

the properties in some domains. We define a fat-tailed domain as follows: a large share of the statistical properties come from the extremum; for a time series involving n observations, as n becomes large, the maximum or minimum observation will be of the same order as the sum. Excursions from the center of the distributions happen brutally and violently; the rare event dominates. And economic variables are extremely fat tailed (Mandelbrot, 1997). Furthermore, standard economic theory makes an allowance for the agency problem but not for the combination of the agency problem, informational opacity, and fat-tailedness. It has not yet caught up to the fact that tail events are not predictable and are not measurable statistically unless one is *causing* them or is involved in increasing their probability by engaging in a certain class of actions with small upsides and large downsides. (Both parties may not be able to gauge probabilities in the tails of the distribution, but the agent knows which tail events do not affect him.) Sadly, the economics literature's treatment of tail risks or "peso problems" has been to see them as outliers to mention *en passant* while hiding them under the rug or removing them from analysis rather than as a core center of modeling and decision making, or to think in terms of robustness and sensitivity to unpredictable events. Indeed, this pushing under the rug of the determining statistical properties explains the failures of economics in mapping the real world, as witnessed by the inability of the economics establishment to see the accumulation of tail risks leading up to the financial crisis of 2008 (Taleb, 2009). The parts of the risk and insurance literature that have focused on tail events and extreme value theory, such as Embrechts (1997), accept the large role of the tails, but then the users of these theories (in the applications) fall for the logical inconsistency of assuming that they can be figured out somehow: this is naive. Because they are rare, what do we know about them? The law of large numbers cannot be of help. Nor do theories have the required robustness. Alarmingly, very little has been done to make the leap that small calibration errors in models can change the probabilities (such as those involving the risks taken in Fukushima's nuclear project) from 1:1,000,000 to 1:50.

Add to the fat-tailedness the asymmetry (or skewness) of the distribution by which a random variable can take very large values on one side but not the other. An operator who wants to hide risk from others can exploit skewness by creating a situation in which he has a small or bounded harm to him while others face large harm, thus exposing others to the bad side of the distributions by fooling them with the tail properties.

Finally, the economic literature focuses on incentives as encouragement or deterrent but not on disincentives as potent filters that remove incompetent and nefarious risk takers from the system. Consider that the symmetry of risks incurred on the road causes the bad driver to eventually exit the system and stop killing others. An unskilled forecaster with skin in the game would eventually go bankrupt or out of business. But if shielded from potentially (financially) harmful exposure, he would continue contributing to the buildup of risks in the system.²

² The core of the problem is as follows. There are two effects: "crooks of randomness" and "fools of randomness" (Nicolas Tabardel, private communication). Skin in the game eliminates the first effect in the short term (standard agency problem), the second one in the long term by forcing a certain class of harmful risk takers to exit from the game.

Hence, there is no possible risk management method that can replace skin in the game in cases where informational opacity is compounded by informational asymmetry, as in the principal-agent problem that arises when those who gain the upside resulting from actions performed under some degree of uncertainty are not the same as those who incur the downside of those same acts.³ For example, bankers and corporate managers get bonuses for positive “performance” but do not have to pay out reverse bonuses for negative performance. This gives them an incentive to bury risks in the tails of the distribution, particularly the left tail, thereby delaying blowups.

The ancients were fully aware of this incentive to hide tail risks and implemented very simple but potent heuristics (for the effectiveness and applicability of fast and frugal heuristics both in general and in the moral domain, see Gigerenzer, 2010). But we find the genesis of both moral philosophy and risk management concentrated within the same rule.⁴ About 3,800 years ago, Hammurabi’s code specified that if a builder builds a house and the house collapses and causes the death of the owner of the house, that builder shall be put to death. This is the best risk-management rule ever.

What the ancients understood very well was that the builder will always know more about the risks than the client, and the builder can hide sources of fragility and improve his profitability by cutting corners. The foundation is the best place to hide such things. The builder can also fool the inspector because the person hiding risk has a large informational advantage over the one who has to find it. The same absence of personal risk is what motivates people to only appear to be doing good rather than actually do it.

Note that Hammurabi’s law is not necessarily literal: damages can be “converted” into monetary compensation. Hammurabi’s law is at the origin of the *lex talionis* (“eye for eye,” discussed later), which, contrary to appearances at first glance, is not literal. *Tractate Bava Kama* in the Babylonian Talmud⁵ builds a consensus that “eye for eye” must be figurative: what if the perpetrator of an eye injury were blind? Would he have to be released of all obligations on grounds that the injury has already been inflicted? Wouldn’t this lead him to inflict damage to other people’s eyesight with total impunity? Likewise, the Quran’s interpretation equally gives the option of the injured party to pardon or alter the punishment.⁶ This nonliteral aspect of the law solves many problems of asymmetry under specialization of labor because the deliverer of a service is not required to have the same exposure in kind but must incur risks that are costly enough to be a disincentive.

³ Note that Pigovian mechanisms fail when, owing to opacity, the person causing the harm is not easy to identify.

⁴ Economics seems to be born out of moral philosophy (mutating into the philosophy of action via decision theory) to which was added naive and improper 19th-century statistics (Taleb, 2007, unpublished). We are trying to go back to its moral philosophical roots, to which we add more sophisticated probability theory and risk management.

⁵ *Tractate Bava Kama*, 84a, Jerusalem: Koren Publishers, 2013.

⁶ Quran, *Surat Al-Ma’idat*, 45: “Then, whoever proves charitable and gives up on his right for reciprocation, it will be an atonement for him.” (our translation).

The problems and remedies are as follows:

First, consider policy makers and politicians. In a decentralized system, say municipalities, these people are typically kept in check by feelings of shame from harming others with their mistakes. In a large centralized system, the sources of error are not so visible. Spreadsheets do not make people feel shame. The penalty of shame is a factor that counts in favor of governments (and businesses) that are small, local, personal, and decentralized versus ones that are large, national or multinational, anonymous, and centralized. When large organizations or governments fail, everybody except the culprit ends up paying the cost, leading to national and international measures of indebtedness against future generations or “austerity” programs.⁷ These points against “big government” models should not be confused with the standard libertarian argument against states securing the welfare of their citizens, but only against doing so in a centralized fashion that enables people to hide behind bureaucratic anonymity. Much better to have a communitarian municipal approach: in situations in which we cannot enforce skin in the game, we should change the system to lower the consequences of errors.

Second, we misunderstand the incentive structure of corporate managers. Counter to public perception, corporate managers are not entrepreneurs. They are not what one could call agents of capitalism. Between 2000 and 2010, in the United States, the stock market lost (depending how one measures it) up to \$2 trillion for investors, compared to leaving their funds in cash or treasury bills. It is tempting to think that since managers are paid on incentive, they would be incurring losses. Not at all: there is instead an irrational and unethical asymmetry. Because of the embedded option in their profession, managers received more than \$400 billion in compensation. The manager who loses money does not return his bonus or incur a negative one.⁸ The built-in optionality in the compensation of corporate managers can only be removed by forcing them to eat some of the losses.⁹

Third, there is a problem with applied and academic economists, quantitative modelers, and policy wonks. The reason that economic models do not fit reality (fat-tailed reality) is that economists have no disincentive and are never penalized for their errors. So long as they please the journal editors or produce cosmetically sound “scientific” papers, their work is fine. So we end up using models such as portfolio theory and similar methods without any remote empirical or mathematical reason. The solution is to prevent economists from teaching practitioners, simply because they have no mechanism to exit

⁷ See McQuillan (2013) and Orr (2013); cf. the “many hands” problem discussed by Thompson (1987).

⁸ There can be situations of overconfidence by which the CEOs of companies bear a disproportionately large amount of risk by investing in their companies, as shown by Malmendier and Tate (2008, 2009), and end up taking more risk because they have skin in the game. But it remains that CEOs have optionality, as shown by the numbers quoted. Furthermore, the heuristic we propose is necessary but may not be sufficient to reduce risk, although CEOs with a poor understanding of risk have an increased probability of personal ruin.

⁹ We define “optionality” as an option-like situation by which an agent has a convex payoff (i.e., has more to gain than to lose from a random variable) and thus has a positive sensitivity to the scale of the distribution (i.e., can benefit from volatility and dispersion of outcomes).

the system in the event of causing risks that harm others. Again, this brings us to decentralization by a system where policy is decided at a local level by smaller units and thus has no need for economists.¹⁰

Fourth, the predictors. Predictions in socioeconomic domains don't work. Predictors are rarely harmed by their predictions. Yet we know that people take more risks after they see a numerical prediction. The solution is to ask—and only take into account—what the predictor has done (what he has in his portfolio) or is committed to doing in the future. It is unethical to drag people into exposures without incurring losses. Furthermore, predictors work with binary variables (Taleb and Tetlock, 2013)—that is, “true” or “false”—and they play with the general public misunderstanding of tail events. They have the incentives to be right more often than wrong, whereas people who have skin in the game do not mind being wrong more often than they are right, provided the wins are large enough. In other words, predictors have an incentive to play the skewness game (more on this below). The simple solution is as follows: predictors should be exposed to the variables they are predicting and should be subjected to the dictum “do not tell people what you think, tell them what you have in your portfolio” (Taleb, 2012: 386). Clearly, predictions are harmful to people because, by the psychological mechanism of anchoring, they increase risk taking.

Fifth, to deal with warmongers, Ralph Nader has rightly proposed that those who vote in favor of war should subject themselves (or their own kin) to the draft.

We believe *skin in the game* is a heuristic for a safe and just society. It is even more necessary in fat-tailed environments. Opposed to this is the unethical practice of taking all the praise and benefits of good fortune while disassociating oneself from the results of bad luck or miscalculation. We situate our view within the framework of ethical debates relating to the moral significance of actions whose effects result from ignorance and luck. We demonstrate here how the idea of skin in the game can effectively resolve debates about (a) moral luck and (b) egoism versus altruism while successfully

¹⁰ A destructive combination of false rigor and lack of skin in the game. The disease of formalism in the application of probability to real life by people who are not harmed by their mistakes can be illustrated as follows, with a very sad case study. One of the most “cited” documents in risk and quantitative methods about “coherent measures of risk” set strong principles on how to compute the “value at risk” and other methods. Initially circulating in 1997, the measures of tail risk—although coherent—have proved to underestimate risk at least 500 million times (the number is not a typo). We have had a few blowups since, including Long Term Capital Management, and we had a few blowups before, but departments of mathematical probability were not informed of them. As we are writing these lines, it was announced that J. P. Morgan made a loss that should have happened every 10 billion years. The firms employing these “risk minds” behind the “seminal” paper blew up and ended up bailed out by the taxpayers. But we now know about a “coherent measure of risk.” This would be the equivalent of risk managing an airplane flight by spending resources to make sure the pilot uses proper grammar when communicating with flight attendants in order to “prevent incoherence.” Clearly, the problem is that tail events are very opaque computationally and that such misplaced precision leads to confusion. The “seminal” paper: Artzner, P., F. Delbaen, J. M. Eber, and D. Heath. 1999. Coherent measures of risk. *Mathematical Finance* 9(3), 203–228.

bypassing (c) debates between subjectivist and objectivist norms of action under uncertainty by showing how their concerns are of no pragmatic concern.

Reputational Costs in Opaque Systems

Note that our analysis includes costs of reputation as skin in the game, with future earnings lowered as the result of a mistake, as with surgeons and people subjected to visible malpractice who have to live with the consequences. So our concern is situations in which cost hiding is effective over and above potential costs of reputation, either because the gains are too large with respect to these costs or because these reputation costs can be “arbitraged” by shifting blame or escaping it altogether because harm is not directly visible. The latter category includes bureaucrats in nonrepeat environments where the delayed harm is not directly attributable to them. Note that in many domains the payoff can be large enough to offset reputational costs, or, as in finance and government, reputations do not seem to be aligned with effective track records. (To use an evolutionary argument, we need to avoid a system in which those who make mistakes stay in the gene pool, but throw others out of it.)

Forecasters

We can see how forecasters who do not have skin in the game have the incentive for betting on the low-impact high-probability event and ignoring the lower probability ones even if these are high impact. They can thus game their reputation. There is confusion between “digital payoffs” (i.e., payoffs linked to probability) and full distribution, called “vanilla payoffs” (i.e., payoffs in which the operator cares about expectation; see Taleb and Tetlock, 2013).¹¹ The crux is that if one earns money 99 percent of the time, but in pennies, and loses dollars 1 percent of the time, according to records he would be an excellent forecaster. But he would be bankrupt if he risked his own money by having skin in the game.

Application of the Heuristic

The heuristic implies that one should be the first consumer of one’s product: a cook should test his own food, helicopter repairpersons should be ready to take random

¹¹ Money managers do not have enough skin in the game unless they are so heavily invested in their funds that they can end up in a net negative from the event. The problem is that they are judged on frequency, not payoff, and tend to cluster together in packs to mitigate losses by making them look like an “industry event.” Many fund managers beat the odds by selling tails, say covered writes, by which one can increase the probability of gains but possibly lower the expectation. They also have the optionality of multitime series; they can manage to hide losing funds in the event of failure. Many fund companies bury hundreds of losing funds away in the “cemetery of history” (Taleb, 2007).

flights on the rotorcraft they maintain, hedge fund managers should be maximally invested in their funds. But it does not naively imply that one should always be using one's product: a barber cannot cut his own hair, the maker of a cancer drug should not be a user of his product unless he is ill. So one should use one's products *conditionally* on being called to use them. However, the rule is far more rigid in matters entailing systemic risks: simply that some decisions should never be taken by a certain class of people.

Heuristic Versus Regulation

A heuristic, unlike a regulation, does not require state intervention for implementation. It is simple contract between willing individuals ("I buy your goods if you use them" or "I will listen to your forecast if you are exposed to losses if you are wrong") and would not require the legal system any more than do simple commercial transactions. It is bottom-up.

The ancients and more-or-less ancients effectively understood the contingency and probabilistic aspect in contract law, and asymmetry under opacity, as reflected in the works of Pierre de Jean Olivi. Moreover, the foundation of maritime law has resided in skin-in-the-game unconditional sharing of losses even as far in the past as 800 B.C. with the *Lex Rhodia*, which stipulates that all parties involved in a transaction have skin in the game and share losses in the event of damage. The rule dates back to Phoenician commerce and caravan trades among Semitic people. We note that the idea is still present in Islamic finance commercial law (see Wardé, 2010).

OPACITY AND RISK HIDING

Here, we summarize in verbal form the mathematical argument concerning payoff skewness and lack of skin in the game presented elsewhere (Taleb, unpublished).

A. If an agent has the upside of the payoff of the random variable, with no downside, and is judged solely on the basis of past performance, then the incentive is to hide risks in the left tail using a negatively skewed (or, more generally, asymmetric) distribution for the performance. This can be generalized to any payoff for which one does not bear the full risks and negative consequences of one's actions.

B. Furthermore, even if it is not intentional (i.e., the agent does not aim at probabilistic rent at the expense of the principal; which is at variance with the way agents are treated in the economics literature), by a survival argument, those agents without skin in the game who tend to engage in strategies that hide risk in the tail tend to fare better and longer and populate the agent population. So the argument is not one of incentive driving the agents but one of survival.

We can sketch a demonstration of these statements with the following reasoning. Assume that an agent has a payoff as a proportional cut of his performance or the benefits to the principal and can get a percentage at year end, his compensation being tied to the *visible* income. The timing of the compensation is periodic, with no total claw back (subsequent obligation to completely return past compensation). The expected value to the agent is that of a stream, a sum of payoffs over time, extending indefinitely (or bounded by the life of the agent). Assume that a loss will reduce his future risk taking or even terminate it, in terms of shrinking of such contracts, owing to change in reputation. A loss would hurt the track record, revealing it, so to speak, and making such a stream of payoffs stop. In addition, the payoff of the agent is compounded over time as the contracts get larger in response to the track record.

Critically, the principal does not observe statistical properties, only realizations of the random variable. However, the agent has an edge over the principal; namely, that *he can select negatively skewed payoffs*. All he needs to do is to figure out the shape of the probability distribution, not its expected returns—nothing else. More technically, the expectation for the agent does not depend on the size of the loss: a small loss or a large loss are the same to him. So the agent can benefit by minimizing the probability of the loss, not the expectation. Minimizing one and not the other results in the most possibly negatively skewed distribution.

This result can be extended to include any situation in which the compensation or reward (in any form) to the agent depends on the probability rather than on the true expectation.

In an evolutionary setting, downside harm via skin in the game would create an absorbing state, without the system failing to be ergodic; hence, it would eliminate this class of risk takers.

SYMMETRICAL CONSTRAINTS IN MORAL PHILOSOPHY

We now turn to a philosophical approach to the problem. The skin-in-the-game heuristic is best viewed as a rule of thumb that places a pragmatic *constraint* on normative theories. Whatever the best moral theory (consequentialism, deontology, contractualism, virtue ethics, particularism, etc.) or political ideology (socialism, capitalism, libertarianism) might be, the “rule” tells us that we should be suspicious of people who appeal to it to justify actions that pass on the cost of any risk-taking to another party while keeping the benefits for themselves. At the heart of this heuristic lies a simple moral objection to negative asymmetry that lies at the heart of some of the oldest and most famous moral ideas, as illustrated in Table 2.1.

Of course, the clearest examples of any rule are likely to stem from a deontological approach, but the skin-in-the-game constraint is not committed to deontology. Indeed,

Table 2.1 Moral Symmetry

1 <i>Lex Talionis:</i> "An eye for an eye, a tooth for a tooth" (Exodus 21:24)	2 <i>15th Law of Holiness and Justice:</i> "Love your neighbor as yourself" (Leviticus 19:18)	3 <i>Silver Rule:</i> "Do not do unto others what you would not have them do unto you" (Isocrates and Hillel the Elder)	4 <i>Golden Rule:</i> "Do unto others as you would have them do unto you" (Matthew 7:12)	5 <i>Formula of the Universal Law:</i> "act only in accordance with that maxim through which you can at the same time will that it become a universal law" (Kant 1785: 4:421)
---	---	---	--	--

moral symmetry is one of the key ideas behind many forms of social contract theory (e.g., "I scratch your back, you scratch mine"), and different emphases on symmetry may also be found in consequentialism (which places the overall good above that of the agent) and virtue ethics (which looks for an ethical mean between excess and deficiency).

As worded, all of the principles in Table 2.1 are problematic. Take, for example, the fourth principle of reciprocity in Table 2.1. This "golden rule" seems to suggest that if I would like you to come up and kiss me, then I should go up to you and kiss you (regardless of whether *you* would like this). But whereas the precise principles may be faulted, the spirit of symmetry behind them (and arguably every moral tradition)¹² contains much insight. Indeed, the very plausibility of Derek Parfit's recent attempt to demonstrate that the best versions of the most popular normative theories converge (Parfit, 2012) must ultimately hang upon a common spirit of this kind. As we shall see, however, there can be positive asymmetries in our behavior as well as negative ones.

ALTRUISM VERSUS EGOISM

Psychological egoists claim that we always do what we most desire (Mandeville, 1714). Those who believe in the possibility of altruism tend to either deny this (Nagel, 1970) or to distinguish between self-centered desires and the desire to benefit others (Butler, 1726). So, although it is not false to think that whether or not we ever act altruistically is an empirical question (Slote, 1964), its answer will partly depend on a priori distinctions among notions such as "desire," "motivation," "reason," and so on. It is such distinctions, rather than experimental research, that allow us to recognize that whereas anyone who

¹² See Blackburn (2001: 101).

Table 2.2 Egoism Versus Altruism

No skin in the game	Skin in the game	Skin in someone else's game
Selfish/egoistic	Neither egoistic nor altruistic	Selfless/altruistic
Negative asymmetry	Symmetry (neutral)	Positive asymmetry
Individualistic Morality	Conventional morality	Other-based morality

is not a sociopath will feel contentment in helping others, it would be perverse to help others *in order* to acquire this feeling (Sandis, 2012: 75; cf. Broad, 1930).

The most pragmatic way of distinguishing between egoists and altruists is to ask whether someone has ever voluntarily (a) paid a cost for someone else's benefit or (b) been willing to reap the rewards of risk while passing the cost to another. The first, altruistic, action is one where the agent has skin in another person's game (Taleb, unpublished), including the lives of future generations.¹³ The second, egoistic act, is one where the person has no skin in the game. People we call "saints" are frequently disposed to act in the former way. Those who tend to act in the latter way we typically call "assholes." In reality, most of us are neither: we usually have skin in our own games and those of our loved ones, but nobody else's. On occasion, however, even the most average of people is liable to either slip up or rise to the occasion. Such moments are respectively marked by negative or positive asymmetries (see Table 2.2).

The middle column in Table 2.2 is the largest because most of the actions of the average person tend to fall within it. It is no wonder, then, that the "eye for an eye" reciprocity it epitomizes is—for better or worse—a conventional morality. To its left lies the sort of individualistic morality frequently associated with Nietzsche but most clearly ascribable to the "rational" normative egoism of Ayn Rand (1964) and many others who maintain that "greed is good." To its right lies the morality of self-sacrifice. This comes in all sorts of stripes: Christian, socialist, utilitarian, and so on. Needless to say, these divisions are never as sharp in practice as they are in theory. Rand's egoistic heroes, for example, subscribe to the symmetrical thought that one should never demand that others take a risk one wouldn't take oneself. Conversely, most welfare states are run by bureaucrats with no skin in the game. Both sides are fooling themselves.

The symmetrical constraint entails that we act wrongly when we open ourselves to great harm that could have reasonably been foreseen and avoided, but the wrongness isn't a moral one. We act immorally when we open *others* to great risk but are only willing to be considered as responsible for our actions if the risk turns out not to harm anyone. Such actions involve the malignant transfer of fragility and antifragility from one party to another with the aim of getting any possible benefits of our actions without

¹³ Such altruism includes cases in which one voluntarily removes oneself from the social pool (e.g., through suicide or self-imposed exile) so as not to harm it. These should be distinguished from the agency problem in evolutionary theory.

being liable for any possible harms (Taleb, 2012). This agency problem is that of a negative asymmetry.

Those who are responsible for such transfers (most predictive analysts, economists, bankers, bureaucrats, consultants, editors, politicians, risk vendors, and sophists) attempt to justify their hypocrisy by appealing to bad luck and uncertainty. They offer excuses of the “we acted on information we believed was correct at the time” or “obviously it fell way short of expectations” variety but refuse to accept any liability for their actions and protest wildly at the mere thought that they should pay the cost. These may be contrasted with those who have skin in the game; that is, those who take risks for themselves and keep their downside. Typical examples are activists, artisans, citizens (as opposed to “*idiotes*”), entrepreneurs, traders, and writers. The greatest contrast, however, is with those who put their own skin in the game for the sake of others. We call such people heroes and saints, but they include not only knights and warriors but also some maverick artists, journalists, scientists, and writers who put their livelihood reputations on the line for the sake of others (Taleb, 2012). This all brings us to the so-called “problem of moral luck.”

MORAL LUCK

Consider the case of two equally reckless drivers, only one of whom kills a pedestrian. According to Bernard Williams, the unlucky driver is morally guilty of something worse than the other driver (namely, manslaughter). Kantians, by contrast, maintain that both drivers would only be liable for reckless driving. Both views are confused. What we should say is that, from the moral point of view, a certain kind of reckless driving is as bad as manslaughter. When a person drives recklessly, he takes upon himself the risk of manslaughter and is accordingly responsible for it if it happens and for opening himself up to it (which is just as bad from a purely ethical point of view) if it doesn’t (see Sandis, 2010). Hegel got it right, then, when he wrote not only that “[t]he laurels of mere willing are dry leaves that never were green” but also:

It happens of course that circumstances may make an action miscarry to a greater or lesser degree. In the case of arson, for instance, the fire may not catch or alternatively it may take hold further than the incendiary intended. In spite of this, however, we must not make this a distinction between good and bad luck, since in acting a man must lay his account with externality. The old proverb is correct: “A flung stone is the devil’s.” To act is to expose oneself to bad luck. Thus bad luck has a right over me and is an embodiment of my own willing. (Hegel, *Philosophy of Right*, 119A)

We are not only responsible for the *known* of our actions and their effects but also for those that we *ought* be aware of (even if we are not). Our ignorance does not always relieve us of responsibility for things we have done because others can claim that, as

rational beings, we *should* know what we were doing even if we did not.¹⁴ Such is the knowledge involved in putting other people's lives at risk with no skin (of our own) in the game. Hegel's solution famously offers two aspects of any given act: *Tat* (deed) corresponding to the objective (which I am causally responsible for) and *Handlung* (action) corresponding to the subjective (which can be morally imputed to me); rights relating to the latter in turn divide into ones relating to various elements of the self, such as knowledge, intention, and purpose (PR 115, 117, 120; see also 118A).

Bad luck is no excuse when it could have been reasonably foreseen. Foresight should not be restricted here to a particular event. If I know that 1 in 1,000 actions of type A will have a tragic result, it is not acceptable to perform thousands of these actions on the grounds that for each one there is only a probability of 1/1,000 that something will go wrong. The greater the potential disaster, the smaller the probability has to be for an act that could bring it about to be immoral. There is an inverse symmetry between the acceptable probability of risk and the weight of the potential damage being assessed.

All action is, to varying degrees, *exposition* luck and must be judged accordingly. When we take a risk, we cannot wash our hands of the consequences for others and hide behind masks of expectation, intention, ignorance, luck, uncertainty, and so on. The central point bears repeating here: *asymmetry in taking risks without having skin in the game is unethical*. Any system deemed "too big too fail" not only encourages but demands that we live according to such skinless asymmetry. The real black swan event of the 21st century is not that any financial crisis occurred (which was predictable) but that there was no full-blown revolution against the governments that continue to encourage "idiotes" to gamble with other people's lives and money.

OBJECTIVISM VERSUS SUBJECTIVISM

The ethics of risk is frequently thought of as a branch of moral philosophy concerned with abstract principles that tell us how we ought to act when we lack (or do not know whether or not we lack) information that is relevant to our choice.¹⁵ Far from being infrequent, such scenarios are the norm and can only be excluded in controlled thought experiments. In an important sense, then, all acts are performed under uncertainty, which is not to say that we never know what the consequences of our actions will be (see Prichard, 1932/2002: 233). This raises the problem of how we ought to act in the face of known ignorance. The skin-in-the-game ethic bypasses the issue, revealing it to be pragmatically irrelevant.

The worry is that of whether a person's obligation to perform (or omit from performing) some action depends "on certain characteristics of the situation in which he

¹⁴ For a related point, see Thompson (1983).

¹⁵ For instance, see Altham (1984), who makes a technical distinction between mere risk and general uncertainty.

is, or on certain characteristics of his thought about the situation” (Prichard, 1932: 84). Objectivists (such as Sidgwick and Parfit) claim that we ought to do whatever is *in fact* best, even when we cannot be reasonably expected to know what this is. By contrast, subjectivists (including Ross, 1939, and Prichard, 1932) claim that we ought to do whatever we *believe* will be best.

The difficulty of choosing between these positions is supposed to stem from two considerations that are in tension. On the one hand, we want to leave room for the thought that we can be unaware of what we ought to do. The fact that what we *believe* we ought to do and what we *actually should* do can come apart in this way seems to lend credence to objectivism. On the other hand, there is the procedural obstacle of the impossibility of stepping out of one’s own mind in order to compare reality with one’s impressions of it. Thus, the objective view appears to entail the absurd view that “although we may have duties, we cannot know but can only believe that we have; and therefore we are rendered uncertain whether we, or anyone else, has ever had, or will ever have a duty” (Prichard, 1932: 88–89).¹⁶ A parallel absurdity is implied in this rhetorical question posed by Jonathan Dancy: “Suppose that, unknown. . . to me, someone has been buried alive in my garden during the night. Could this make it wrong of me to go away for a fortnight’s holiday?” (Dancy, 2000: 57). Prospectivists, most prominently Michael Zimmerman, attempt to avoid this dilemma by arguing that we ought to perform whichever action it is most reasonable to *expect* will be the best.

Such academic debates have little pragmatic weight. All three views share the common mistaken assumption that they are each motivated by the same notions of “what one ought to do” when there are actually three different concepts at play:

1. *Objectivists* equate *what* we ought to do with whichever action turns out to be best. This is what we should *aim* at when we act.
2. *Subjectivists* equate what we ought to do with whatever we judge to be best. This the only way *through which* we can aim at what is best.
3. *Prospectivists* equate what we ought to do with what we can rationally expect to be best. This view attempts to reconcile objectivist and subjectivist intuitions that are only in tension because of the aforementioned assumption.

Whereas objectivists are concerned with the rightness of the things we do (typically thought to be universals), prospectivists and subjectivists are concerned with the rightness of our acts of *doing* these things (typically thought to be particulars). Yet it is possible that one rightly acts in doing something that results in negative value and, by the same token, that one acts wrongly in doing something that turns out positively.¹⁷

¹⁶ Ross 1930 rightly (but for the wrong reason) suggests that objectivists and subjectivists are talking at cross purposes. Cf. Zimmerman (2008: 1–2).

¹⁷ This point runs parallel to the distinction between a belief and a believing being justified (e.g., as introduced in the literature on Gettier cases).

Given that one can do the right thing for the wrong reason, the deontic question of *what* the right thing to do is should therefore be distinguished from the evaluative question of when one is acting *rightly*. The evaluative question is best answered via an account of how and when people and institutions are *liable* for choices they make under uncertainty. We have sought to answer the question (e.g., in the case of moral luck) via the skin-in-the game principle. Strictly speaking, this necessary (although insufficient) moral heuristic is not about action but about dispositions. Indeed, it relates directly to the virtue of being such that the system will not only *survive* uncertainty, randomness, and *volatility* but will actually benefit from it. Such a system is defined as *antifragile* (see Taleb, 2012).¹⁸ Skin-in-the-game heuristics follow directly from the principle of antifragility.

ACKNOWLEDGMENTS

The authors thank Jon Elster, David Chambliss Johnson, Gur Huberman, Raphael Douady, Robert Shaw, Ralph Nader, Barkley Rosser, James Franklin, Russell Roberts, Marc Abrahams, Andreas Lind, and Elias Korosis.

REFERENCES

- Altham, J. E. J. 1984. "Ethics of Risk." *Proceedings of the Aristotelian Society*. New Series 84(1983–1984): 15–29.
- Blackburn, S. 2001. *Ethics: A Very Short Introduction*. Oxford: Oxford University Press.
- Broad, C. D. 1930. *Five Types of Ethical Theory*. London: Kegan Paul.
- Butler, J. 1726. "Sermons I and XI." In his *Fifteen Sermons Preached at the Rolls Chapel*, in *The Works of Bishop Butler*, edited by J. H. Bernard. London: Macmillan.
- Dancy, J. 2000. *Practical Reality*. Oxford: Oxford University Press.
- Dancy, J. 2002. "Prichard on Duty and Ignorance of Fact." In *Ethical Intuitionism: Re-evaluations*, edited by P. Stratton-Lake. Oxford: Oxford University Press.
- Embrechts, P. 1997. *Modelling Extremal Events: For Insurance and Finance*. Springer.
- Gigerenzer, G. 2010. "Moral Satisficing: Rethinking Moral Behavior as Bounded Rationality." *Topics in Cognitive Science* 2: 528–554.
- Grossman, S. J., and O. D. Hart. 1983a. "An Analysis of the Principal-Agent Problem." *Econometrica* 7–45.
- Grossman, S. J., and O. D. Hart. 1983b. "Implicit Contracts Under Asymmetric Information." *The Quarterly Journal of Economics* 123–156.

¹⁸ One can, of course, render the virtue of the previous sentence into a principle about action ("act in whatever way renders you and the systems that you inhabit antifragile"), but such a principle treats antifragility as the ultimate end-in-itself, whereas it is best to treat it as a property whose value is derived from its effects.

- Hegel, G. W. F. 1991. *Elements of the Philosophy of Right*, translated by H. B. Nisbet. Cambridge: Cambridge University Press.
- Hölmstrom, B. 1979. "Moral Hazard and Observability." *The Bell Journal of Economics* 74–91.
- Kant, I. 1785/2002. *Groundwork for the Metaphysic of Morals*, A. Zweig, edited by T. E. Hill, Jr. and A. Zweig. Oxford: Oxford University Press.
- Malmendier, U., and G. Tate. 2008. "Who Makes Acquisitions? CEO Overconfidence and the Market's Reaction." *Journal of Financial Economics* 89(1): 20–43.
- Malmendier, U., and G. Tate. 2009. "Superstar CEOs." *Quarterly Journal of Economics* 124(4): 1593–1638.
- Mandelbrot, Benoit B. 1997. *Fractals and Scaling in Finance: Discontinuity, Concentration, Risk*. New York: Springer-Verlag.
- Mandeville, B. 1714. *The Fable of the Bees*. London: J Roberts.
- McQuillan, M. 2013. "The Severity of Austerity." *Times Higher Education* 30–35.
- Nagel, T. 1970. *The Possibility of Altruism*. Princeton, NJ: Princeton University Press.
- Orr, D. 2013. "It Wasn't Labour Who Spent Too Much, It Was the Banks. How Did We Forget This?" *The Guardian*, Saturday 5 January: 39.
- Parfit, D. 2012. *On What Matters*, vols. I & II. Oxford: Oxford University Press.
- Prichard, H. A. 1932/2002. "Duty and Ignorance of Fact." In *Moral Writings*, edited by J. MacAdam. Oxford: Oxford University Press.
- Rand, A. 1964. *The Virtue of Selfishness: A New Concept of Egoism*. New York: Signet.
- Rescher, N. 2009. *Ignorance: On the Wider Implications of Deficient Knowledge*. Pittsburgh: University of Pittsburgh Press.
- Ross, S. A. 1973. "The Economic Theory of Agency: The Principal's Problem." *The American Economic Review* 63(2): 134–139.
- Ross, W. D. 1930. *The Right and the Good*, rev. ed. 2002. Edited by P. Stratton-Lake.
- Ross, W. D. 1939. *The Foundations of Ethics*. Oxford: Clarendon Press.
- Sandis, C. 2010. "The Man Who Mistook His Handlung for a Tat: Hegel on Oedipus and Other Tragic Thebans." *Bulletin of the Hegel Society of Great Britain* 62: 35–60.
- Sandis, C. 2012. *The Things We Do and Why We Do Them*. Palgrave Macmillan.
- Slote, M. A. 1964. "An Empirical Basis for Psychological Egoism." *Journal of Philosophy* 61: 530–537.
- Stiglitz, J. E. 1988. *Principal and Agent*. In *The New Palgrave: A Dictionary of Economics*, vol. 3. London: Macmillan.
- Taleb, N. N. 2007. *The Black Swan*. Penguin and Random House.
- Taleb, N. N. 2009. "Errors, Robustness, and the Fourth Quadrant." *International Journal of Forecasting* 25(4): 744–759.
- Taleb, N. N. 2012. *Antifragile: Things That Gain From Disorder*. Penguin and Random House.
- Taleb, N. N. unpublished. *Silent Risk*, manuscript.
- Taleb, N. N., and P. E. Tetlock. 2013. On the Difference Between Binary Prediction and True Exposure with Implications for Forecasting Tournaments and Decision Making Research (June 25, 2013). Available at <http://ssrn.com/abstract=2284964> or <http://dx.doi.org/10.2139/ssrn.2284964>
- Thompson, D. F. 1983. "Ascribing Responsibility to Advisers in Government." *Ethics* 93(3): 546–560.
- Thompson, D. F. 1987. *Political Ethics and Public Office*. Cambridge, MA: Harvard University Press.

- Tirole, J. 1988. *The Theory of Industrial Organization*. Cambridge: MIT Press.
- Wardé, I. 2010. *Islamic Finance in the Global Economy*. Edinburgh University Press.
- Williams, B. 1993. *Shame and Necessity*. Cambridge: Cambridge University Press.
- Zimmerman, M. J. 2008. *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge: Cambridge University Press.

CHAPTER 3

THE ETHICS OF ECONOMIC DECISION RULES

SVEN OVE HANSSON

INTRODUCTION

ACCORDING to the received view, moral philosophy and the decision sciences are distinct disciplines with clearly demarcated subject areas. Moral philosophy is concerned with the choice of values to be used as guides in practical decision making. It does not tell us much about the means for realization of these values. The decision sciences take values as given and add no new values. Instead, they show us what to do to maximize the attainment of the given values, using for that purpose criteria of rationality that are assumed to be devoid of moral implications. The common decision-making rules that are used in economics are taken to be based on such value-free rationality criteria, which means that they tell us how to reach the goals we have chosen but they have nothing to say about what goals we should choose.

This conception of decision rules is based on the premise that instrumental rationality can at least in principle be perfectly separated from moral values and, more generally, from the ends and purposes of our actions. This is often called the Humean view of practical reasoning, despite the fact that it does not seem to have been endorsed by Hume himself (Hampton, 1995; Millgram, 1995). It is an important part of the self-understanding of economics and the decision sciences. If we can fully separate ends and means from each other, then we can presumably also give an objective answer to the question how best to further a given set of goals.

But can we? Or does our use of economic decision rules have implications that reach into the realm of values? Are the supposedly value-free, entirely rationality-driven decision rules in fact imbued with moral values? To answer that question we need to take a closer look at the various decision rules that arise in economic practice, and we also need to clarify what we use them for and what impacts they have on our decision making.

Let us begin with a very simple situation: Suppose that there is only one thing we want to achieve,¹ and we know to what extent each of our alternative actions will lead to its achievement. In such a situation we have no need for a decision rule—apart from the trivial rule “choose an option that will maximize goal achievement.” Economic decision rules are intended to enhance our ability to deal with situations that differ from this uncomplicated case in at least one of the following two respects: Either we have *competing goals*, or we have *uncertainty* about how the goals can be achieved. Therefore, the question of whether the decision rules are value-free or value-laden has to be seen in relation to how we deal respectively with competing goals and uncertainty about goal achievement. In the second section the focus is on decision rules for competing goals and in the third section on rules for dealing with uncertainty.²

DECISION RULES FOR GOAL CONFLICTS

Three Types of Goal Conflicts

For our purposes it is useful to distinguish between three major types of competition or conflict among goals. The first are *intercategorical* goal conflicts, conflicts among goals that belong to different categories, that is, they are different in nature. This is perhaps the type of goal conflict that comes most easily to mind. Everyday life is full of conflicts of this type, many of which arise because our time and our money are limited. For instance, the goal of keeping down public expenditures may run into conflict with many, if not most, other social welfare goals.

The second category consists of conflicts among goals that may be similar in nature but concern the interests of different persons. They can be referred to as *interindividual* (interpersonal) goal conflicts. They are often presented as problems of distribution. Some examples are the distribution of wages in a company, of the time of physicians and nurses among patients in a hospital, of risks among soldiers, and of social resources in a society at large.

The third category comprises *intertemporal* goal conflicts, conflicts among goals to be satisfied at different points or periods in time. Some intertemporal problems refer to the same person or persons, for instance, a person's choice of whether to do tedious work now and have the next weekend free, or take free time now and defer the drudgery to next weekend. Many intertemporal goal conflicts refer to goals that involve partly but not completely overlapping sets of concerned persons. One example is the conflict between investment and future consumption. Investments will reduce current consumption, but may be expected to result in significantly higher consumption within a

¹ Or several goals, so constructed that the actions needed to achieve them coincide completely.

² This chapter is devoted to the use of decision rules to guide decisions. Decision rules can also be used as descriptive or analytical tools, for example, when a decision is described as being compatible with a certain rule. From this it does not follow that the rule was actually used.

decade or two. Many, but not all, of those who forego consumption today will benefit from the resulting higher future consumption. Obviously, intertemporal conflicts over longer time intervals, such as the conflict between the costs of climate change mitigation measures and the burdens of a less inhabitable planet some hundred years ahead, are interpersonal in addition to being intertemporal. (They can also be called intergenerational conflicts.) The interpersonal aspect of such a combined conflict may be at least as important as the intertemporal aspect. It would be a serious mistake to take it for granted that such problems can be solved with the same decision rules that we use for intertemporal conflicts in a single person's life.

The Reductive Approach

Let us begin with intercategory goal conflicts, conflicts between different types of goals. Such conflicts are common enough. Many if not most of the decisions that we find to be difficult, both in private and public life, are difficult largely because we cannot easily weigh the different goals against each other. A parent considering divorce may face difficulty in comparing the effects on the children's welfare with those on her own freedom and quality of life. Politicians and public officials preparing decisions about infrastructure investments may have to weigh the advantages of improved transportation against the negative environmental impacts of a new road or railroad. But although these and many other comparisons may be very difficult to make, somehow we have to make them. A complete stalemate in all decisions that hinge on intercategory comparisons is not an option.

Many proposals have been made to facilitate decisions involving difficult comparisons. A common recommendation is to systematize the decision by identifying all the advantages and all the disadvantages of each of the options that are open to us. The next, more difficult step is to compare the advantages and disadvantages, or weigh them against each other. A practical way to do this was proposed in lively terms by Benjamin Franklin in 1772 in a letter to the chemist Joseph Priestley:

When these difficult Cases occur . . . my Way is, to divide half a Sheet of Paper by a Line into two Columns, writing over the one *Pro*, and over the other *Con*. Then during three or four Days Consideration I put down under the different Heads short Hints of the different Motives that at different Times occur to me for or against the Measure. When I have thus got them all together in one View, I endeavour to estimate their respective Weights; and where I find two, one on each side, that seem equal, I strike them both out: If I find a Reason *pro* equal to some two Reasons *con*, I strike out the three . . . and if after a Day or two of farther Consideration nothing new that is of Importance occurs on either side, I come to a Determination accordingly. (Franklin, 1970: 437–438)

Franklin struck out items or groups of items with equal weight. From this it is not a big step to assign to each item a number representing its weight, and to add up these

numbers in each column. This is the decision procedure proposed by the moral philosopher Jeremy Bentham (1748–1832):

Sum up all the values of all the *pleasures* on the one side, and those of all the pains on the other. The balance, if it be on the side of pleasure, will give the *good* tendency of the act upon the whole, with respect to the interests of that *individual* person; if on the side of pain, the *bad* tendency of it upon the whole.

Take an account of the *number* of persons whose interests appear to be concerned; and repeat the above process with respect to each. *Sum up* the numbers . . . Take the *balance* which if on the side of *pleasure*, will give the general *good tendency* of the act, with respect to the total number or community of individuals concerned; if on the side of pain, the general *evil tendency*, with respect to the same community. (Bentham, 1780: 27–28)

Bentham used the word “utility” for “that property in any object, whereby it tends to produce benefit, advantage, pleasure, good, or happiness” (Bentham, 1780: 2). Therefore, moral theories based on this type of calculus are called “utilitarian.” But neither Bentham nor any of his successors have been able to come up with a compelling method to actually measure the moral values of options. Modern utilitarian philosophers often refer to a hypothetical unit called “util” or “utile,” and in the literature on utilitarianism we can read phrases such as “Suppose person A receives 3 units of utility and person B loses 2 units . . .” However, this is just a fictitious measure, not a real one. Its relationship to meters and liters is the same as that of a fire-spitting dragon to lizards and turtles.

But in spite of this, utilitarian calculations are performed in economics, namely in the discipline of cost–benefit analysis. This is an economic discipline that employs a collection of decision-facilitating techniques in which numerical estimates of advantages and disadvantages are established and weighed against each other. In a typical cost–benefit analysis, two or more options in a public decision are compared to each other in a calculation that is in a sense quite Benthamite. The most conspicuous difference is that in cost–benefit analysis the common currency is not a fictional util but instead a most real currency, such as dollars or euros. The basic methodology of cost–benefit analysis is to value all the different types of advantages and disadvantages in monetary terms, determine the amount of these advantages and disadvantages for the different alternatives, and then sum up the monetary values for each alternative. The result is a total value for each option under consideration, summing up consequences that may be very different in nature, such as economic costs, risks of disease and death, environmental damage, and so forth. The associated decision rule is to choose an option for which the highest total “economic” value has been calculated.

Arguably, cost–benefit analysis does exactly what Franklin achieved by identifying pairs of advantages and disadvantages that intuitively cancel out each other. But cost–benefit analysis has a major advantage over Franklin’s method in terms of cognitive economy: the intuitive comparisons between different types of advantages and disadvantages do not have to be repeated again for each new decision. An example will clarify the difference: Suppose that in an analysis of traffic projects we have concluded,

after considerable deliberation, that 5 minutes' daily reduction in travel time for 100,000 persons cancels out 100 cases of asthma. Presumably, it took us a lot of time and trouble to reach that conclusion. The next day we are confronted with a different but similar problem: Can 20 minutes' reduction in travel time for 1000 persons cancel out 4 cases of asthma? Applying Franklin's method would require us to start the whole intuitive process again. We are in a much better position if we have assigned numerical values in some unit (such as money) to a minute's time gain and to a case of asthma. All we have to do when a new case arrives is to apply these values in a new calculation.

But cost-benefit analysis is controversial. Probably the most common criticism is that some negative outcomes, in particular the loss of human lives, cannot be measured in money (Zelizer, 1978; Ashby, 1980; Baram, 1981; Kelman, 1981). It is indeed a rather odd activity to assign economic values to assets that do not have a market value. Prices arise on a market, and since values such as the "human life values" of cost-benefit analysis have no relation to a market, they are not prices. Obviously the assignment of a sum of money to the loss of a human life does not imply that someone can buy another person, or the right to kill her, at that price. Such a value means something else, namely either (1) that society tends to pay up to that sum to save a human life, or (2) that society ought to pay up to that sum to save a human life.

It is important to note that the incommensurability (unmeasurability) of lives in terms of money is only one of many incommensurabilities that cost-benefit analysis has to deal with. There is no definite answer to how many cases of juvenile diabetes correspond to one death, or what amount of human suffering or death corresponds to the extinction of an antelope species. Since such comparisons are technically effectuated in a cost-benefit analysis by assigning monetary values, the problem of incommensurability appears to be a problem of monetization. But even if money were removed from the analysis it would still be necessary to deal with comparisons between deaths, diseases, environmental damage, losses of cultural heritage, and so forth. Hence, the fundamental problem is not that we compare these disparate consequences of our actions in terms of money. Instead, the fundamental problem is that we compare them at all. But of course we have no choice, since such "impossible" comparisons are inherent in all major social decisions. Cost-benefit analysis has the advantage of bringing them to light and making the reduction from a multidimensional to a unidimensional decision problem extremely transparent. If we removed money from the analysis we would still have to deal with comparisons between deaths, diseases, and environmental damage. How should we, for instance, compare the value of saving the Indian tiger to the "tens to hundreds" (Nyhus et al., 2003) of people who are killed by wild tigers every year? This comparison is difficult enough in itself, and whether we make it by translating both values into money does not seem to have much impact on its difficulty.

The reason why monetary units are used in these calculations is of course that they are highly convenient for the purpose. Economic costs and gains are parts of almost all social decisions. Therefore, money is the only unit that is present in almost all decisions, irrespective of social sector. Furthermore, many of the goods and services that need to be taken into account in decisions are bought and sold on markets, and for them we can

use market values in the calculations.³ We need to add fictional values “only” for those items that do not have a market price.

But in spite of its practical advantages, the choice of money as a universal currency for all types of values may have an undesirable impact on our moral deliberations.⁴ It has often been argued that certain types of goods should not be valued as market commodities. Instead they should be valued “in accordance with a particular set of norms” appropriate to their nature that differs from how we treat marketable commodities (Anderson, 1990: 72; cf. Anderson, 2000). Putting a monetary price on such objects can then incite us to treat them inappropriately (Titmuss, 1970). In particular, the assignment of an economic value to a human life has been conceived as sending a message that desecrates the value of life. Stuart Hampshire (1972: 9) has warned that the habits of mind engendered by this type of comparisons may lead to “a coarseness and grossness of moral feeling, a blunting of sensibility, and a suppression of individual discrimination and gentleness.”

Irrespective of what type of unit is used for the comparisons, cost-benefit analysis and its associated decision rule have another feature that may be morally problematic. As already mentioned, the assignment of values to various aspects of decision outcomes is performed before the actual analysis of the different decision options, and it is assumed to be the same in different cost-benefit analyses.⁵ Therefore, cost-benefit analysis is in fact a *two-tiered* process. In the first tier the money equivalents of outcomes such as lost lives, diseases, and different types of environmental damage are determined for cost-benefit analysis in general. In a second tier, these general values are applied to a specific problem, and a complete balance is obtained for each of the options in that decision. This is very different from Franklin’s method, in which the comparative weighing of disparate effects took place on a case-by-case basis.

Most decisions involving incommensurable values are not performed in accordance with cost-benefit analysis. Instead they follow an *integrated one-tiered* process, in which the comparative impact of different aspects of the outcome is open to renegotiation in each particular decision. This is exemplified by a national budget process, in which policy goals in areas as diverse as natural security, crime prevention, healthcare, education, research, and social security have to be weighed against each other. This is done in an intuitive fashion, in a process driven by argumentation and negotiation from the

³ The use of market prices is not morally neutral. Market prices may be influenced by morally indefensible preferences, and there are markets for goods and services whose sale is outright immoral (such as dangerous recreational drugs). However, this is not a specific problem for cost-benefit analysis but one that it shares with real markets. It will therefore not be treated here.

⁴ The only common alternative to monetary units is QALYs, quality-adjusted life years, which is used in health care to evaluate the effects of treatments and other interventions.

⁵ This appears to be in part due to practical considerations (the first tier is resource intensive) and in part due to a conviction that the calculation values should be the same in all applications. According to Kip Viscusi we should “spend up to the same marginal cost-per-life-saved amount for different agencies” (Viscusi, 2000: 855)

beginning to the very end. The outcome is reached without any explicit translation of one type of values into another.

The two-tiered method of cost–benefit analysis has created incentives to invent “objective” or otherwise uncontroversial ways to perform the first tier. One of the most common first-tier methodologies is contingent valuation. This is a survey-based method to create monetary values for nonmarket goods. It is commonly applied to environmental resources such as the preservation of a species or a natural reserve. People are asked for instance how much they would be willing to pay to preserve the giant panda. This methodology has been subject to severe criticism (Diamond and Hausman, 1994; Hausman, 2012). One of its most serious problems is that the sums mentioned in responses to these questions do not seem to differ in a credible way between questions referring to different objectives. For instance, in one study “willingness to pay to clean one lake is approximately equal to stated willingness to pay to clean up five lakes—including the one asked about individually” (Hausman, 2012). Some studies indicate that many respondents tend to report an amount that would not seriously disturb their normal expenditure and savings patterns, in other words, a sum corresponding to what they would normally be willing to pay to a charity (Beattie et al., 1998).

Another method, particularly common with respect to estimates of the value of life, is to derive values from willingness to pay for related goods. There is a long tradition in cost–benefit analysis of deriving the value of life from estimates of how much male workers are paid extra for working in occupations with a high risk of fatal accidents. These calculations are based on the assumption of a constant marginal value. However, it does not take much reflection to see that this is an untenable assumption. If a worker accepts a risk of 1 in 1000 of being killed in a mining accident next year, it cannot be inferred that he would take a 10 times higher risk against a 10 times higher premium, and it most certainly cannot be inferred that he would accept certainty of being killed against an amount that is 1000 times higher. Furthermore, such marginal values cannot necessarily be transferred to another population or another situation. It is difficult to justify the use of life values derived from marginal risk taking by male workers in cost–benefit analysis for another group that is exposed to a quite different type of risk. This was pointed out by Heinzerling (2000) when she criticized the use of these values to monetize the benefits for women of lifesaving screening to avoid breast cancer. Instead, women could have been asked how much they were prepared to pay for mammography, given realistic assumptions about the risk reduction it gives rise to. Their willingness to pay for reduced risks could then be used in a cost–benefit analysis. Although the use of such values would not have been unproblematic, it would at least have been much closer to the relevant context than the life value that was actually used (Hansson, 2007a).

The two-tiered structure of cost–benefit analysis ensures a sort of consistency that can be summarized as “the same life value for all purposes.” The same monetary life value is used in transport planning, health-care priority-setting, environmental regulation, and all other areas in which such a value is asked for. However, it is far from evident that this type of consistency is desirable in all cases. We tend to pay more to save a life in some contexts than in others, and some of these differences may be morally justifiable. Many

of us would be willing to pay more per life saved in a law enforcement program that reduces the frequency of manslaughter than what we would pay for most other lifesaving activities. There are plausible reasons to do so: We consider homicide a particularly unacceptable form of fatality; and we put a particular value on security against violent crime. Cost-benefit analysis as it is currently performed leads to a disregard of such values, which is of course another proof that it is a far from value-neutral tool. (It is possible to modify cost-benefit analysis so that it treats different deaths differently. That, of course, would not make it value free, only value laden in a different way.)

Since cost-benefit analysis aims at numerical calculations, it also tends to leave out aspects of future developments that can be predicted only in nonquantitative terms. This applies for instance to risks of cultural impoverishment, social isolation, and increased tensions between social strata (Hansson, 1989). Furthermore, owing to their aggregative structure, cost-benefit analyses often leave out social justice and other distributional aspects from the analysis even when they are accessible to quantitative treatment. In response to criticism of such limitations, cost-benefit analysts often point out that the factors that were left out could easily be included in the analysis. It is, for instance, not difficult to put a price on inequality and include it in a cost-benefit analysis, and the same applies to many other aspects that are commonly left out (Johansson-Stenman, 2005; Sen, 2000). However, such all-encompassing cost-benefit analyses are much more seldom performed than they are referred to in defense of the methodology.

Efficiency

Up to now we have discussed attempts to measure all aspects of a decision in one and the same value currency, in order to sum them up and maximize the outcome in terms of the aggregated value. But we have also seen that there are considerable problems with that method. The other major way to deal with plurality of goals is to refrain from reducing the goals to one shared value, and instead restrict the policy advice to what can be said without such a reduction. There are indeed cases when meaningful advice can be given without goal reduction. For a simple example, suppose that you need to borrow an extra table for a party. Your neighbor offers you to borrow one of three tables. You classify their properties as follows:

Table 1: right size, unstable, ugly

Table 2: right size, stable, ugly

Table 3: too small, stable, ugly

If these are all the aspects you wish to take into account then it would be difficult to deny that Table 2 is your best choice. The reason for this is that it *dominates* the other two options. An option dominates another if and only if it is better in at least one aspect and not worse in any aspect. If an option dominates all the other options, then it is the obvious choice.

You also need to borrow a large bowl, and fortunately your neighbor has three that you may choose among. You classify their relevant properties as follows:

Bowl 1: breakable, beautiful

Bowl 2: unbreakable, ugly

Bowl 3: breakable, ugly

Since Bowl 3 is dominated by each of the other two, it should not be your choice. However, neither Bowl 1 nor Bowl 2 dominates the other. To choose between the two you would have to weigh the two criteria against each other. And this is precisely the procedure from the previous section that we have now set out to do without. All we can say in this case is that Bowl 1 and Bowl 2 are the two undominated options. Given your choice of these two decision criteria, you should therefore choose one of these two bowls, and not choose Bowl 3.

Another, more common way to express this is that the first and second alternatives in this case are the two *efficient* alternatives. An alternative is efficient in relation to a given set of goals if and only if none among the other options is an improvement in terms of one of the goals without also being a change for the worse in terms of at least one of the other goals (Sen, 1975; Le Grand, 1990, 1991). Efficiency is, of course, a central concept in economics. The different types of efficiency discussed in economics differ in the number and nature of the goals that they refer to. *Single-goal efficiency* is also called effectiveness. The concept of effectiveness is often used in technical contexts. For instance, the most effective cleaning agent is the one that removes most of the dirt. *Two-goal efficiency* can be of different types, the most common of which are cost efficiency and productivity. In cost efficiency, one of the goals is to minimize costs, and the other is usually the maximization of production. (In practice strivings for cost efficiency can take the form of minimizing costs relative to a fixed level of output, or maximizing production relative to some fixed cost.) If a health care unit is cost efficient, then we cannot produce the same health care for less money or improved health care for the same money. Productivity is usually conceived as efficiency with respect to the two goals of maximal production and minimal labor input. (In practice, productivity can be achieved by maximizing production volume relative to a fixed labor input, or minimizing labor input relative to a fixed level of output.⁶) Finally, *multi-goal efficiency* is most commonly referred to in discussions of the distribution of economic advantages among individuals. The welfare of each

⁶ If the two goals are both measurable in numerical terms, then two-goal efficiency can be expressed as the (single) goal of maximizing a single number. In particular, if one of the goals is to maximize something and the other is to minimize something else (usually expenditure or effort), then we can refer to the ratio between the two numbers instead of using the general definition of two-goal efficiency. For instance, productivity can be characterized as maximization of the output per workhour. In cases when at least one of the goals cannot be expressed numerically, the more general definition cannot be replaced in this way. Such cases are rarely discussed in economics, but they are nevertheless common in social contexts. A young academic may have two goals: to qualify herself as a researcher and as a teacher. Even if these goals are not measurable in numerical terms we may talk about efficiency in the more general two-goal sense, but there is of course no way to reduce this instance of two-goal efficiency to the maximization of a single number.

individual can then be treated as a goal in itself. (See the subsection Interindividual Goal Conflicts).

It is essential to keep in mind that there is no efficiency per se, only efficiency in relation to specified goals. Claims that a certain activity is efficient or inefficient must always be specifiable in term of the goals referred to. Ambiguity in this respect has often led to confusion. For instance, around 1990 a lively debate took place on the efficiency of the slavery-based agricultural economy of the American South (Saraydar, 1989, 1990, 1991; Nye, 1990; Singh and Frantz, 1991; Schmidt, 1994). The plantation system appears to have been efficient in terms of cotton production per unit of labor input, that is, in terms of the two goals to maximize cotton production and to minimize labor input. However, it was far from efficient in terms of welfare production; much more welfare could have been produced with the same labor input (namely in an economy with a free workforce). To clarify this distinction it is necessary to replace unspecified efficiency claims by a precise analysis in which the goals to which efficiency refers are specified.

The injunction to be as efficient as possible should be counted among the economic decision rules, but it is often taken to be such a self-evident rationality requirement that its status as a decision rule is not recognized. Once we have realized that efficiency is always goal relative it becomes obvious that the requirement to be efficient is reasonable only to the extent that the particular efficiency concept to which it refers has sufficient coverage of goals. A good case can be made that we should be efficient with respect to the full set of goals for our endeavors. However, this does not imply that we should be efficient with respect to various subsets of the full set. Consider the following examples:

- ECONOMIST: You are building this shieling in a very inefficient way. If you used modern saws you could have obtained the same result in less than half the time.
- CARPENTER: Yes, but we also strive to maintain knowledge about old craft methods.
- ECONOMIST: Running this factory with only 40 work hours a week is not at all efficient. Shift work would lead to much better use of the expensive investments in machinery.
- MANAGER: We have refrained from introducing shift work because of its negative effects on worker welfare.
- ECONOMIST: This farm is inefficient. You could produce more wheat per acre if you used more fertilizers.
- FARMER: I know, but that would lead to nitrogen leaching and groundwater contamination.

In all these cases (and many others), claims that an activity is inefficient are based on a questionable delimitation of the goals that the activity should satisfy. Furthermore, in cases like these the choice of goals for a social activity is at least in part a moral choice. This brings us to an important way in which seemingly value-free statements about efficiency can in fact be value laden: A claim that an activity is inefficient is value laden to the extent that it is based on the exclusion of goals that might have led to another conclusion.

One possible response to this problem is to explicitly delimit the range of goals, for instance, to those that we consider to be central for economic analysis. If goals such as employee health, environmental quality, and preservation of cultural heritage are regarded as noneconomic, then we can describe the three cases above as cases in which (economic) efficiency clashes with other social goals. Under this approach, (economic) efficiency is not always desirable; in fact this is how laypeople often think about economic efficiency. However, this approach has the disadvantage of not helping us to do as much good as we can for health and the environment with the resources that we have at our disposal. An analysis that includes these goals might, for instance, reveal that we could achieve more for the environment without any sacrifice in terms of our other goals.

Such considerations would seem to lead us to include all relevant goal dimensions in efficiency analysis. But unfortunately, such a practice would create problems because of a phenomenon that may be called the *efficiency dilemma*. It consists in a conflict between goal coverage and decisiveness. On one hand, we want to base our efficiency analysis on a selection of goals that is as exhaustive as possible, so that all important considerations are included. This often leads to the inclusion of a rather long list of goals. On the other hand, we also want our search for efficient options to be of real help in decision making. This requires that the set of efficient options is as small as possible; if it (almost) coincides with the set of all available actions, then the whole exercise is (almost) of no help. Unfortunately, the more aspects we include, the larger can we expect the set of efficient alternatives to be. The reason for this is a very simple mathematical fact: Including an additional aspect in the analysis can make a previously inefficient alternative efficient, but not the other way around. To exemplify this, consider the following set of alternative ways to construct a new road: (Time gain refers to average gain in travel time for road users.)

Option 1: Time gain 5 minutes, cost €100,000,000

Option 2: Time gain 4 minutes, cost €110,000,000

Option 3: Time gain 3 minutes, cost €170,000,000

Here, Option 1 is the only cost-efficient option, and there seems to be no doubt that it should be chosen. But let us add a third dimension, environmental impact. It turns out as follows:

Option 1: Time gain 5 minutes, cost €100,000,000, large environmental damage

Option 2: Time gain 4 minutes, cost €110,000,000, medium environmental damage

Option 3: Time gain 3 minutes, cost €170,000,000, small environmental damage

When the additional dimension has been included, all three options are efficient, which means that efficiency analysis does not help us to make a choice. Unfortunately, this is a common situation. Owing to the efficiency dilemma, decisiveness in efficiency analysis

is often achievable only at the price of limiting the number of aspects that are taken into account, and doing so threatens the moral adequacy of the conclusion.

Interindividual Goal Conflicts

In the previous two subsections we had intercategorical goal conflicts as examples. Most of what was said in those sections can be applied to interindividual goal conflicts as well, but the latter are in need of a special discussion because in practice, we treat them differently. The major difference is the assumption that there is one goal per person, usually referring to the welfare or the interests of that person. In a social context, this often means that a large number of goals will be involved.

Interindividual goal conflicts can be illustrated with the help of vectors representing distributions of the good. (We can leave open what the good is, but it may for instance be economic resources.) Thus in a three-person case we can compare vectors such as $\langle 16, 2, 2 \rangle$ and $\langle 6, 6, 6 \rangle$, where the numbers represent the amount of good assigned to each of the three persons. In this particular example, if the numbers are utilitarian “utils,” a utilitarian would prefer $\langle 16, 2, 2 \rangle$ to $\langle 6, 6, 6 \rangle$, for the simple reason that the former vector has a higher sum (20) than the latter (18).

The other extreme in this case is to apply the concept of efficiency. Neither $\langle 16, 2, 2 \rangle$ nor $\langle 6, 6, 6 \rangle$ dominates the other, since the first individual is worse off in $\langle 6, 6, 6 \rangle$ and the other two are worse off in $\langle 16, 2, 2 \rangle$. If these two are the only options that we can choose between, then they are both efficient. This type of efficiency is called Pareto efficiency. (The original term “Pareto optimality” is misleading and should be avoided.) Two major decision rules are based on this concept (Sen, 1987).⁷ The *weak Pareto rule* recommends that we perform Pareto efficient changes (Pareto improvements) if we have a chance to do so, but it has no recommendations on what to do otherwise. Thus suppose that the current situation is $\langle 10, 3, 3 \rangle$ and we have a chance to change either to $\langle 21, 3, 3 \rangle$, $\langle 9, 9, 9 \rangle$, or $\langle 6, 3, 3 \rangle$. The weak Pareto rule recommends a change to $\langle 21, 3, 3 \rangle$ since it is a Pareto improvement (someone gains and no one loses). It advises against a change to $\langle 6, 3, 3 \rangle$ since that would represent a Pareto loss. However, this rule is agnostic about a change to $\langle 9, 9, 9 \rangle$ since such a change is neither a Pareto improvement nor a Pareto loss.

⁷ There is also a third, the so-called Potential Pareto Criterion or Kaldor-Hicks Criterion. According to this criterion, a policy change is commendable if it would have been possible for those whom it makes better off to compensate those who were initially worse off, so that in the end no one would lose and at least one would gain from the combination of the measure and the posterior redistribution. This criterion is often invoked in defense of cost-benefit analyses that are purely aggregative, that is, performed as if the distribution of benefits and detriments did not matter. It was originally justified by Nicholas Kaldor as follows: “In all cases, therefore, where a certain policy leads to an increase in physical productivity, and thus of aggregate real income, the economist’s case for the policy is quite unaffected by the question of the comparability of individual satisfactions; since in all such cases it is *possible* to make everybody better off than before, or at any rate to make some people better off without making anybody worse off” (Kaldor, 1939: 550). However, the justificatory power of a transaction that does not take place is obviously nil (Rateman, 2012).

The *strong Pareto rule* recommends that we only perform Pareto improvements, and no other changes. It differs from the weak Pareto rule in advising against changes that are neither Pareto improvements nor Pareto losses. Thus, in our example it would advise in favor of a change to $\langle 21, 3, 3 \rangle$ and against changes to $\langle 9, 9, 9 \rangle$ as well as $\langle 6, 3, 3 \rangle$. Opposition to change such as that from $\langle 10, 3, 3 \rangle$ to $\langle 9, 9, 9 \rangle$, which increases both equality and total welfare, is of course not a morally neutral standpoint. To the contrary, it has immediate moral and political implications.

The weak (but not the strong) Pareto rule is plausible, given that we distribute “the good” to individuals. But if the ultimate goal is well-being, then we cannot distribute the ultimate good. Instead, we can (sometimes) distribute something that promotes wellbeing. Let us call that which is distributed “resources.” The usefulness of weak Pareto efficiency in economic theory is based on a simple but remarkably powerful assumption on the relationship between resources and well-being: It is assumed that each individual’s well-being increases when her resources increase. Nothing needs to be assumed about how much it increases. To see how this assumption works, again consider a small society with three individuals. Owing to our assumption there is for each of them a function from resources to well-being. Individual 1’s function f_1 satisfies the simple criterion that for any two amounts x and y of resources:

$$f_1(x) < f_1(y) \text{ holds if and only if } x < y$$

In other words, more resources always generate more well-being, and fewer resources always generate less well-being. The same holds for the corresponding functions f_2 and f_3 of the other two persons. Now consider any two distributions of resources, represented by the vectors $\langle x, y, z \rangle$ and $\langle x', y', z' \rangle$. It is easy to show that $\langle x, y, z \rangle$ dominates over $\langle x', y', z' \rangle$ if and only if $\langle f_1(x), f_2(y), f_3(z) \rangle$ dominates over $\langle f_1(x'), f_2(y'), f_3(z') \rangle$. From this it follows that from the viewpoint of Pareto efficiency it makes no difference if we consider a set of resource vectors or the corresponding set of well-being vectors; the same changes will be Pareto improvements on both counts. This simple property of Pareto efficiency is a large part of the explanation of why it is such a useful tool in economics.

But the formal elegance of this conclusion is based on an assumption that is quite problematic both empirically and morally, namely the assumption that each individual’s well-being depends exclusively on her or his own material resources. In actual life, each individual’s well-being depends also on other factors, among them material resources that belong to other persons. Such dependencies can come in many variants. For a simple example, suppose that you live in a small village with three inhabitants. You have eight units of resources.

Case 1: The resource vector of the village is $\langle 8, 8, 8 \rangle$.

Case 2: The resource vector of the village is $\langle 32, 30, 8 \rangle$.

We might expect the well-being that you derive from your eight units of resources to be much lower in the second case, owing to the negative effects of the inferior comparative position (Hirsch, 1976; Hansson 2004). This contradicts the assumption just referred to that your well-being is strictly a function of (i.e., only depends on) the resources at your own disposal. Your relative position, that is, how the resources at your disposal relate to those of the other inhabitants, also has a role. Taking such “positional” effects into account would seem to be necessary not only descriptively but also from the viewpoint of moral philosophy, since the effects of being in an inferior position has moral implications (Wilkinson and Pickett, 2009). As soon as we allow for positional effects, Pareto efficiency on the two levels will no longer coincide. In particular, the well-being vector associated with $\langle 32, 30, 8 \rangle$ will not be an improvement over that associated with $\langle 8, 8, 8 \rangle$ since the well-being of the third individual is expectedly worse in the former case owing to positional effects.

Intertemporal Conflicts

Sometimes when summarizing the advantages and disadvantages of an option we find that they materialize at different points in time. For the smoker, the most important positive effect of smoking is immediate: she avoids the symptoms of nicotine withdrawal. The most important negative effect is the risk of serious disease that will typically materialize decades later. (About half of smokers die prematurely due to smoking [Boyle, 1997].) In climate and environmental policies, we often consider measures that cost money today but have their positive effects much later. Nuclear waste management provides what is perhaps the most extreme example of such temporal discrepancies: on one hand energy is produced to be consumed now, while on the other hand the potential damages from nuclear waste may materialize hundreds of thousands of years hence.

Intertemporal conflicts seldom come alone. They are often combined with both intercategory and interindividual conflicts, and with uncertainty. We can see this from the two examples just mentioned. The smoker’s decision whether to (try to) stop smoking is not only intertemporal. It is in most cases to some extent also an interindividual decision (i.e., unless she always smokes at sufficient distance from others to prevent secondary exposure, and no one else is affected by her welfare). It is also a decision under uncertainty since not all smokers are affected by serious smoking-related disease. The social decision on nuclear waste is intertemporal, intercategory, and interindividual, and it takes place under uncertainty.

In a systematic treatment these components of decision problems should be discussed separately as far as possible. To begin with, let us consider purely intertemporal decisions, that is, decisions in which intercategory, interindividual, and uncertainty-related aspects are either absent or so unimportant that we are justified in neglecting them. Such situations are in fact common, and as children we were all told how to deal with them:

CHILD: Father, I want to eat all the ice cream today.

FATHER: But then there will nothing left for tomorrow. Think of how sorry you will be then. Eat half of it, and save the other half for tomorrow.

The typical parents' advice in issues like this is to divide the limited resource in equal parts between the instances of consumption. The presumption seems to be that tomorrow's enjoyment is equally valuable as today's. The child would typically not be enjoined, for instance, to "eat two thirds today and leave a third for tomorrow, since having a good time tomorrow is not worth as much as having a good time today."

Moral philosophers have not spent much effort on (purely) intertemporal issues, but instead have focused on interindividual ones. However, some moral philosophers have commented on the intertemporal issues within a person's own life, and they have essentially given the same advice as the parent in the above ice cream example. For instance, Thomas Hobbes (1588–1679) regretted that "men cannot put off this same irrational appetite, whereby they greedily prefer the present good (to which, by strict consequence, many unforeseen evils do adhere) before the future." (Hobbes, [1642] 1841: 48).⁸ Henry Sidgwick (1838–1900) wrote that proximity in time "is a property which it is reasonable to disregard except in so far as it diminishes uncertainty."

For my feelings a year hence should be just as important to me as my feelings next minute, if only I could make an equally sure forecast of them. Indeed this equal and impartial concern for all parts of one's conscious life is perhaps the most prominent element in the common notion of the *rational*—as apposed to the merely *impulsive*—pursuit of pleasure. (Sidgwick, 1907: 124; cf. 381)

More recently, Thomas Nagel proposed "a standpoint of temporal neutrality towards the events of our lives" (Nagel, 1970: 61). To the extent that other moral philosophers have expressed views on the matter, this is usually the view they have expressed.⁹

The major contribution of economics to the evaluation of future outcomes is the method of discounting future values. It was originally developed for money, and based on the assumption of a foreseeable, positive interest rate. For example, suppose that the interest rate is constantly 3 percent in real terms, and furthermore suppose that we want to have €100,000 in 10 years. Then it is sufficient to deposit €74,400 in the bank. We can therefore say that the "present cost" of having €100,000 10 years from now is €74,400. Alternatively, we can say that the "present value" of receiving €100,000 10 years from now is €74,400. With a similar argument, a loss €100,000 10 years from now corresponds to a loss of €74,400 today. In this way, we can "convert" future monetary value into current monetary value and vice versa.¹⁰

⁸ *De Cive*, Chapter 3, §32. In Latin: "homines exuere non possunt appetitum illum irrationalem, quo bona praesentia (quibus, arcta consequentia, multa adhaerent improvisa mala) prae futuris appetunt."

⁹ See also Rawls (1972: 422–423), Williams (1976), Ladmore (1999), Dickenson (1991), and Hansson (2007b, 2013).

¹⁰ More precisely, the value of a future good is assumed to be equal to the product of two factors. One of these is a time-independent evaluation of the good in question, that is, the value of having it now. The other factor represents the subject's "pure time preferences." It is a function of the length of the delay, and is the same for all types of goods. This can be summarized in the formula is $v_0(x) = v_t(x) \times 1/(1+r)^t$, where x is the object whose value we are converting, $v_0(x)$ its present value, $v_t(x)$ its value after t years and r the interest rate (Samuelson, 1937).

The practice of discounting has often been seen as at variance with the moral principle of temporal neutrality just referred to. That, however, is not necessarily true since the relationship between the two principles depends on how much money can buy at different points in time. If you can buy exactly as much of what you want to buy for €74,400 today as you can for €100,000 in 10 years, then discounting will be perfectly compatible with temporal neutrality in the moral sense. If you can buy more of what you want for €74,400 now than you can for €100,000 in 10 years, then discounting overrates future values, in comparison to what temporal neutrality prescribes. If it is the other way around, then discounting underrates future values. Which of these is the case is an empirical, not a conceptual question.

In cost–benefit analysis, discounting is the standard way to deal with intertemporal conflicts. It is routinely used for comparisons that are both intertemporal and interindividual. Suppose for instance that we discuss measures that would prevent an accident 15 years into the future in which 31 persons would die. With a 3 percent interest rate, the formula tells us to value the loss of 31 lives in 15 years the same way that we would value a loss of 20 lives today.¹¹

However, although discounting is a reasonable practice for real market prices, it is much more problematic when applied to the calculation values for nonmarket entities that are used in cost–benefit analysis. Obviously, there is no interest rate for lives, and we cannot convert human lives today directly into human lives in the future. The same applies to other nonmarket goods such as environmental preservation. The only feasible way to defend discounting of nonmarket goods is indirect: they are discounted because they can be converted to money, and money can be discounted. Therefore, a justification for discounting for instance future losses in lives presupposes a series of three conversions:

1. From human lives today to money today
2. From money today to money in the future
3. From money at some future point in time to lives at that same future point in time

Let us consider each of these steps in turn. The first step is problematic for the reasons discussed in the subsection *The Reductive Approach*. It can be performed as a way to summarize how we in practice settle conflicts between incommensurables, but the process of doing so is very different in nature from the market mechanisms that give rise to prices in the ordinary sense of the word. The second step is a standard economic procedure. In the time spans that economists usually work with, it is both justified and indispensable. In longer time spans, its value is more uncertain since we cannot take for granted that the economy will function in the same way as today thousands of years ahead. But in the third step, the whole conversion breaks down. There is no ground whatsoever for projecting the tradeoffs that we make today between human lives and

¹¹ $31 \times 1/1,03^{15} \approx 20$.

money far off into the future. We have no reason to believe that our present priorities in these respects will remain the same over a longer period of time. (This is corroborated by historical experience. Although difficult to measure, our willingness to accept losses in lives seems to have decreased substantially; see Smith, 2005.)

It should also be noted that discounting of lives yields absurd results if applied to long time periods. Consider, as a simple schematic example, a hypothetical choice between the following two actions:

1. Killing one person now.
2. Performing an action today that will lead to the death of the whole population of the earth, 10 billion people, in the year 2800.

If we apply a discount rate of 3 percent, then the first of these actions will be worse than the second. The example is unrealistic, but it illustrates that even very large disasters count as almost nothing if they take place a couple of hundred years from now. (Lowering the discount rate only delays this effect. With a discount rate of 0.5 percent it will still be worse that one person dies today than that 10 billion people die in 4620 years.)

Summing this up, the common economic decision rule for future outcomes, namely to discount them, is sound in the short or medium run if applied to money and to goods with a monetary price. However, discounting is unjustified for goods without a price (even if a nonmarket monetary value has been assigned to it). The uncertainty component of intertemporal decision problems has not been discussed here; it is taken up in the subsection Uncertainty about the Future.

DECISION RULES FOR UNCERTAINTY

Up to now we have abstracted from uncertainty, that is, we have assumed that we know the consequences of the different options in our decisions. That is of course an idealization. In real life, uncertainty about future consequences is ubiquitous, and it is often the most important source of difficulties, disagreements, and irresoluteness in our decisions. The most common way to structure these difficulties is to assign probabilities to the various outcomes under consideration. This usually makes decision problems much easier to tackle. Consider, for instance, a company's decision on whether to buy one of its competitors. Someone points out: "It is possible that this acquisition will be stopped by the antitrust authorities." It is difficult to know what weight to assign to such a supposition. It is much easier to act on more specified information such as "according to our experts, there is a 90 percent probability that the acquisition will be stopped by the antitrust authorities."

The standard terminology distinguishes between decision making under uncertainty and decision making under risk. A decision is said to take place under risk if the

probabilities of the relevant outcomes are assumed to be known by the decision maker. If these probabilities are unknown or only insufficiently known, then the decision is said to take place under uncertainty (Luce and Raiffa, 1957: 13). It should be noted that a decision “under risk” need not be one in which the probabilities are known. It is sufficient that they are assumed to be known. It should also be noted that the term “uncertainty” is ambiguous since it is used both in the sense just mentioned and as a general term also covering decision making under risk.

Expected Utility and Its Limitations

The standard rule for decision making under risk is expected utility maximization. The word “expected” refers here to probability-weighting, and the rule recommends the use probability-weighted values (statistical expectation values) for all outcome measures. If there is a probability of 0.7 to lose €1,000,000 then that is counted as a loss of €700,000. If there is a probability of 0.1 to gain \$20,000 then that is counted as an income of \$2000. If there is a probability of 0.01 of an accident that will kill 200 persons, then that is counted as equivalent with certainty that 2 persons will be killed. This is how both risks and uncertain benefits are dealt with in cost–benefit analysis.

The decision rule prescribing maximization of expected utility is often described as the “risk-neutral” decision rule. By this is meant that it puts neither too large nor too small an emphasis on undesired, uncertain outcomes. (Decision rules that put more emphasis on avoiding undesired outcomes than what expected utility maximization does are called risk averse, and those deviating in the opposite direction are called risk prone.) It is often assumed that deviations from risk-neutrality in this sense are irrational and that the expected utility rule is a requirement of instrumental reason, devoid of moral contents (Charnley, 1999, 2000; Durodié, 2003). However, that turns out to be far from true, and this for at least two major reasons.

The first reason is that probability estimates do not tell us all we need to know about the uncertainties we wish to take into account. To see that, let us first consider a simple, schematic example: A dime has been found among the property of a deceased cardsharp. We suspect that the coin may be unfair, but we have no clue to whether it is in that case biased toward heads or tails. Someone decides to toss the coin. If I have to assign a probability that it will yield heads, then I will say 0.5. This is the same answer that I would have given before someone tossed a coin that I knew to be fair. However, although the probabilities are the same, I am much more uncertain about the behavior of the cardsharp’s coin than about that of the ordinary coin.

Suppose that we throw the coin and get heads 10 times in a row. If we had very good reasons to believe that the coin is fair, then we would see this as an unusual random event, and we would assign the probability 0.5 to the next throw yielding heads. But since it is a coin from the cardsharp’s collection, we tend to see 10 heads in a row as a strong indication that the coin is biased in favor of heads. Therefore, we have good reasons to assign a probability higher than 0.5 to the event of this coin yielding heads when tossed an 11th time.

All this can be expressed in exact mathematical terms, but I will leave out that aspect here.¹² Instead, let us consider a highly important practical case, namely climate change. There is scientific uncertainty regarding the effects on the future climate of any possible future pattern of greenhouse gas emissions. We may specify this as uncertainty in the choice among several scientific models that yield different probabilities of climate events. The consensus or near-consensus of climate scientists is stated in the reports of the IPCC (Intergovernmental Panel on Climate Change). There is, however, a possibility that the IPCC's conclusions are too pessimistic. It is at least in principle possible that the greenhouse effect is completely counterbalanced by other mechanisms, so that no global warming disaster is approaching. On the other hand, there is also a possibility that global warming will accelerate much faster than what the IPCC predicts. Consider the probability of some specific potential effect of global warming (such as a large sea level rise within 20 years). For concreteness, we may assume that the probability of that event is 0.01 according to the IPCC, 0.9 according to the more pessimistic view, and 0 according to the more optimistic view. The standard approach in such situations is to use the appraisal that has the highest credibility. In our case, this means that decisions will be based on the IPCC's estimate that the probability is 0.01. However, although convenient, this approach is obviously flawed. In an adequate appraisal of the probability in question, the other two possibilities need to be taken into account. Unless the more pessimistic of the two is highly improbable, it could be seen as dangerously incautious not to take it into account (Hansson, 2006).

The second problem with the expected utility rule is that even if we actually know the probabilities, it does not follow that we have to base our decisions on probability-weighting and expected utility maximization. Suppose that an eccentric person offers you a bet that is based on repeated tosses of a fair coin of your own choice. If the coin yields heads n times in a row, then you have to give everything you own to the eccentric. Otherwise, he will pay you €100,000. According to the standard assumptions of expected utility maximization, there must be some n such that you are willing to play the game. However, it is not necessarily irrational to refuse the bet, simply on the ground that you are not willing to risk losing all that you own. In policy discussions the avoidance of very large catastrophes, such as a nuclear or chemical accident costing thousands of human lives, is often given higher priority than what is warranted by the statistically expected number of deaths. It has also been argued that as a matter of principle, serious events with low probabilities should be given a higher weight in decision making than what they receive in the expected utility model (O'Riordan and Cameron, 1994; O'Riordan et al., 2001; Burgos and Defeo, 2004).

In interindividual decision problems, there are additional moral reasons to deviate from the expected utility rule:

In an acute situation we have to choose between two ways to repair a serious gas leakage in the machine-room of a chemical factory. One of the options is to send in the

¹² See Hansson (2008, 2009) for a formal treatment.

repairman immediately. There is only one person at hand who is competent to do the job. He will run a risk of 0.9 to die due to an explosion of the gas immediately after he has performed the necessary technical operations. The other option is to immediately let out gas into the environment. In that case, the repairman will run no particular risk, but each of 10,000 persons in the immediate vicinity of the plant runs a risk of 0.001 to be killed by the toxic effects of the gas. (Hansson, 1993: 24, 2013: 27)

In this case, to maximize expected utility we would have to send in the repairman. However, that would be a morally repulsive decision. In this and many other cases, a rational decision maker may refrain from maximizing expected utility (minimizing expected damage) for a moral reason, namely to avoid being unfair to a single individual.

Alternative Decision Rules

One alternative to expected utility maximization is to completely disregard probabilities and focus entirely on avoiding the worst possibilities. This is done by using the maximin decision rule. It requires that we identify, for each alternative action, its security level, that is (the value of) the worst possible outcome that it can give rise to. We then have to choose an alternative with a maximal security level, thus maximizing the minimal outcome. This means that we completely disregard all but the worst possible outcomes of an action. Not surprisingly, the use of such a decision rule can have strange consequences.

You are offered the choice between two investments. Option A yields a 99 percent chance of gaining €2000 and a 1 percent chance of gaining only €50. Option B yields a gain of €60 for sure.

The maximin rule recommends you to choose option B. As this example shows, the total neglect of chances and of all outcomes but the worst renders this rule unworkable in practice.¹³

An intermediate solution seems more promising. We can decide on the degree of cautiousness that we wish to apply. Mathematically it can take the form of a number c between 0 and 1, called the *cautiousness index*. 0 represents minimal and 1 maximal cautiousness. For each option x in the decision, let $\text{SEC}(x)$ be its security value (the value of its worst possible outcome) and $\text{EU}(x)$ its expectation value (expected utility). The rule requires that we choose an option with a maximal value of

$$c \times \text{SEC}(x) + (1 - c) \times \text{EU}(x)$$

¹³ The maximax rule (choosing an option whose best outcome is at least as good as that of any other option) has the same problem. Although maximax decision making is seldom recommended, it appears to be common in practice (DeMartino, 2011; cf. Hansson, 2006: 234n).

For $c = 0$ this coincides with the expected utility rule and for $c = 1$ with the maximin rule. We can of course select different values of c for different decision problems, depending on what degree of cautiousness we consider to be appropriate.¹⁴

The repertoire of available methods for decision making becomes larger if we replace the single probability function of the traditional model by a credal set, by which is meant a set of probability functions. The credal set is interpreted as consisting of the plausible probability functions. For a simple example, consider the following:

There are two jars in the room. One of them contains 5 red and 95 black balls. The other contains 95 red and 5 black ones. Someone puts one of the jars—you do not know which—in front of you and asks: “If you draw a ball from this jar, what is the probability that it is red?”

A quite natural answer would be “It is either 0.05 or 0.95.” More precisely, you hesitate between two probability functions, p_1 and p_2 , such that p_1 assigns the probability 0.05 and p_2 the probability 0.95 to the ball you draw being red. Your credal set is then $\{p_1, p_2\}$. Similarly, in the above example of the cardsharp’s coin, the credal set may, for instance, consist of all probability functions that assign a value between 0.3 and 0.7 to the coin yielding heads.¹⁵

Even more information can be encoded if we introduce second-order probabilities, that is, probabilities over probabilities. In the example of the two jars, it would be quite plausible to say that p_1 and p_2 are equally probable. We can then assign the (second-order) probability 0.5 to each of them. This and other, more elaborate representations of uncertainties allow for more sophisticated decision rules than the expected utility rule (Walley, 2000; Halpern, 2003; Hansson, 2016) but none of these constructions have gained much influence. Expected utility maximization is still the dominating economic decision rule for decision making under uncertainty.

Uncertainty about the Future

Finally, let us return to the issue that was left open in the subsection Intertemporal Conflicts, namely how to deal with intertemporal decisions in which there is uncertainty about the future. According to the standard approach to economic discounting, one of the functions of the discount rate is to reflect these uncertainties. However, a

¹⁴ This type of index is usually called an optimism–pessimism index, but that is a misleading name. A person is pessimistic with respect to a possible undesirable event if she believes that its probability is high or its consequences large, or both. A person is cautious with respect to that same event if she is willing to take significant trouble to decrease its probability or limit its consequences, or both. Pessimism and cautiousness need not coincide.

¹⁵ Higher or lower values do not seem to be technically feasible for coins (Gelman and Nolan, 2002).

few examples are sufficient to show that uncertainties about the future seldom have the structure of interest rates:

The choice between a benefit now and the same benefit in 10 years' time is not the same for a person in her late 80s as it is for a healthy person in her 20s. A constant discount rate cannot reflect these differences.

We know fairly well what effects a massive mercury spillage in a lake today will have on life in the lake a decade later. For spillages of most other chemical substances we do not have that information. Therefore the uncertainty of such events is much larger. A discount rate that is the same for all decisions cannot capture the difference.

We have limited knowledge about the persistence of a particular substance in the environment: We know that its half-life in the environment is between 2 and 20 years, but we cannot say anything more precise. Then our uncertainty is small about its environmental concentration one month from now, large about its concentration in 20 years and small about its concentration in 400 years. A constant discount rate cannot reflect these differences.

The methods introduced in the previous subsection are much more useful than a discount rate to deal with uncertainty about the future. A person choosing between taking a long vacation trip now or taking it in 10 years can, for instance, think in terms of a credal set consisting of two probability functions: one in which she is alive and well in 10 years and one in which she is not. If she assigns second-order probabilities to them, then these probabilities will depend on her age and health status (contrary to a discount rate that is assumed to be the same for all persons and all decision problems). Similarly, when we are uncertain about the future effects of some environmental pollution, we can describe these uncertainties with the help of a credal set containing several probability functions corresponding to the possibilities that scientists hesitate between. Again, second-order probabilities can be used to represent the credibility of these different scenarios.

In general, discounting is unsuitable to represent our uncertainty about the future. Such uncertainty is much better represented by credal sets and second-order probabilities. The latter representation also has the advantage of leaving a choice open between different decision rules depending, for instance, on what degree of cautiousness we want to apply.

CONCLUSION

The use of decision rules has obvious advantages. Rules can simplify decision making and make it more consistent. But the employment of decision rules can also have negative effects. When using them we achieve simplicity by restricting the amount of information that we take into account in any particular decision. This may lead to disregard

for aspects of the decision that should have been taken into account. In this way, the potential for a normative problem is inherent in the very notion of a decision rule. In a sense rules can be said to function as blinkers on a horse: They make it possible to move forward, but only at the price of leaving much unseen. This presentation has focused on some of the problems with common economic decision rules, and in particular on how their use makes value assumptions invisible that should instead have been brought forward and discussed.

So what should we do? There are two major strategies for improvement. One is to develop improved decision rules, for instance, rules that make the value assumptions visible (such as the cautiousness index of the subsection Alternative Decision Rules). The other strategy is to improve the way in which decision rules are used, for instance, by always clarifying to decision makers what aspects have been excluded from consideration and what the effects of including them might have been. In my view, these two strategies should be combined to develop a new style of policy advice that highlights the value issues instead of hiding them in the dark.

ACKNOWLEDGMENT

I thank George DeMartino for unusually helpful comments on an earlier version of this chapter.

REFERENCES

- Anderson, Elizabeth S. 1990. "Is Women's Labor a Commodity?" *Philosophy and Public Affairs* 19: 71–92.
- Anderson, Elizabeth S. 2000. "Why Commercial Surrogate Motherhood Unethically Commodifies Women and Children: Reply to McLachlan and Swales." *Health Care Analysis* 8: 19–26.
- Ashby, Eric. 1980. "What Price the Furbish Lousewort?" *Environmental Science & Technology* 14: 1176–1181.
- Baram, Michael S. 1981. "The Use of Cost-Benefit Analysis in Regulatory Decision-making Is Proving Harmful to Public Health." *Annals of the New York Academy of Sciences* 363: 123–128.
- Beattie, J., Covey, J., Dolan, P., Hopkins, L., Jones-Lee, M., Loomes, G., Pidgeon, N., Robinson, A., and Spencer, A. 1998. "On the Contingent Valuation of Safety and the Safety of Contingent Valuation: Part 1—Caveat Investigator." *Journal of Risk and Uncertainty* 17: 5–25.
- Bentham, Jeremy. 1780. *An Introduction to the Principles of Morals and Legislation*. London: T. Payne. (<http://gallica.bnf.fr/ark:/12148/bpt6k93974k/f2.image.r=.langEN>)
- Boyle, P. 1997. "Cancer, Cigarette Smoking and Premature Death in Europe: A Review Including the Recommendations of European Cancer Experts Consensus Meeting, Helsinki, October 1996." *Lung Cancer* 17(1): 1–60.

- Burgos, R., and Defeo, O. 2004. "Long-term Population Structure, Mortality and Modeling of a Tropical Multi-fleet Fishery: The Red Grouper *Epinephelus morio* of the Campeche Bank, Gulf of Mexico." *Fisheries Research* 66: 325–335.
- Charnley, G. 1999. "President's Message." *RISK Newsletter* 19 (2): 2.
- Charnley, G. 2000. "1999 Annual Meeting: Past President's Message: Risk Analysis under Fire." *RISK Newsletter* 20(1): 3.
- DeMartino, George F. 2011. *The Economist's Oath: On the Need for and Content of Professional Economic Ethics*. Oxford: Oxford University Press.
- Diamond, Peter A., and Hausman, Jerry A. 1994. "Contingent Valuation: Is some Number Better than No Number?" *Journal of Economic Perspectives* 8: 45–64.
- Dickenson, Donna. 1991. *Moral Luck in Medical Ethics and Practical Politics*. Aldershot, UK: Avebury.
- Durodié, Bill. 2003. "The True Cost of Precautionary Chemicals Regulation." *Risk Analysis* 23: 389–398.
- Franklin, Benjamin. 1970. *The Writings of Benjamin Franklin, Vol. V, 1767–1772*, edited by Albert Henry Smyth. New York: Haskell House.
- Gelman, Andrew, and Nolan, Deborah. 2002. "You Can Load a Die, but You Can't Bias a Coin." *American Statistician* 56: 308–311.
- Halpern, Joseph. 2003. *Reasoning about Uncertainty*. Cambridge, MA: MIT Press.
- Hampshire, Stuart. 1972. *Morality and Pessimism*. Cambridge, UK: Cambridge University Press.
- Hampton, J. 1995. "Does Hume Have an Instrumental Conception of Practical Reason?" *Hume Studies* 21: 57–74.
- Hansson, Sven Ove. 1989. "Dimensions of Risk." *Risk Analysis* 9: 107–112.
- Hansson, Sven Ove. 1993. "The False Promises of Risk Analysis." *Ratio* 6: 16–26.
- Hansson, Sven Ove. 2004. "Welfare, Justice, and Pareto Efficiency." *Ethical Theory and Moral Practice* 7: 361–380.
- Hansson, Sven Ove. 2006. "Economic (Ir)rationality in Risk Analysis." *Economics and Philosophy* 22: 231–241.
- Hansson, Sven Ove. 2007a. "Philosophical Problems in Cost-Benefit Analysis." *Economics and Philosophy* 23: 163–183.
- Hansson, Sven Ove. 2007b. "Hypothetical Retrospection." *Ethical Theory and Moral Practice* 10: 145–157.
- Hansson, Sven Ove. 2008. "Do We Need Second-order Probabilities?" *Dialectica* 62: 525–533.
- Hansson, Sven Ove. 2009. "Measuring Uncertainty." *Studia Logica* 93: 21–40.
- Hansson, Sven Ove. 2013. *The Ethics of Risk. Ethical Analysis in an Uncertain World*. London: Palgrave Macmillan.
- Hansson, Sven Ove. 2016. "Representing Uncertainty. In *A Handbook of Formal Philosophy*, edited by Sven Ove Hansson and Vincent F. Hendricks. New York: Springer.
- Hausman, Jerry. 2012. "Contingent Valuation: From Dubious to Hopeless." *Journal of Economic Perspectives* 26: 43–56.
- Heinzerling, L. 2000. "The Rights of Statistical People." *Harvard Environmental Law Review* 24: 189–207.
- Hirsch, F. 1976. *Social Limits of Growth*. Cambridge, MA: Harvard University Press.
- Hobbes, Thomas. [1642] 1841. *De Cive* in *English Works*, Vol. 2, edited by William Molesworth. London: Bohn.
- Johansson-Stenman, Olof. 2005. "Distributional Weights in Cost-Benefit Analysis—Should We Forget about Them?" *Land Economics* 81: 337–352.

- Kaldor, Nicholas. 1939. "Welfare Propositions of Economics and Interpersonal Comparisons of Utility." *Economic Journal* 49: 549–552.
- Kelman, Steven. 1981. "Cost-Benefit Analysis. An Ethical Critique." *Regulation* 5: 33–40.
- Ladmore, Charles. 1999. "The Idea of a Life Plan." *Social Philosophy and Policy* 16: 96–112.
- Le Grand, Julian. 1990. "Equity versus Efficiency: The Elusive Trade-off." *Ethics* 100: 554–568.
- Le Grand, Julian. 1991. *Equity and Choice. An Essay in Economics and Applied Philosophy*. London: Harper Collins Academic.
- Luce, Duncan R., and Howard, Raiffa. 1957. *Games and Decisions. Introduction and Critical Survey*. New York: John Wiley & Sons.
- Millgram, E. 1995. "Was Hume a Humean?" *Hume Studies* 21(1): 75–93.
- Nagel, Thomas. 1970. *The Possibility of Altruism*. Oxford: Clarendon Press.
- Nye, John Vincent. 1990. "The Conflation of Productivity and Efficiency in Economics and Economic History—A Comment." *Economics and Philosophy* 6: 147–152.
- Nyhus, P. J., Tilson, R. L., and Tomlinson, J. L. 2003. "Dangerous Animals in Captivity: Ex Situ Tiger Conflict and Implications for Private Ownership of Exotic Animals." *Zoo Biology* 22: 573–586.
- O'Riordan, T., and Cameron, J., eds., 1994. *Interpreting the Precautionary Principle*. London: Earthscan.
- O'Riordan, T., Cameron, J., and Jordan, A., eds. 2001. *Reinterpreting the Precautionary Principle*. London: Cameron May.
- Rateman, Ty. 2012. "Regulation, Compensation, and the Loss of Life: What Cost-Benefit Analysis Really Requires." *Ethics, Policy and Environment* 15: 97–118.
- Rawls, John. 1972. *A Theory of Justice*. Oxford: Oxford University Press.
- Samuelson, P. A. 1937. "A Note on Measurement of Utility." *Review of Economic Studies* 4: 155–161.
- Saraydar, Edward. 1989. "The Conflation of Productivity and Efficiency in Economics and Economic History." *Economics and Philosophy* 5: 55–67.
- Saraydar, Edward. 1990. "The Inefficiency of Some Efficiency Comparisons. A Reply to Nye." *Economics and Philosophy* 6: 153–155.
- Saraydar, Edward. 1991. "Productivity and X-efficiency." *Economics and Philosophy* 7: 91–92.
- Schmidt, Volker H. 1994. "Some Equity–Efficiency Trade-offs in the Provision of Scarce Goods: The Case of Lifesaving Medical Resources." *Journal of Political Philosophy* 2: 44–66.
- Sen, Amartya. 1975. "The Concept of Efficiency. In *Contemporary Issues in Economics*, edited by M. Parkin and A. R. Nobay, pp. 196–210. Manchester: Manchester University Press.
- Sen, Amartya. 1987. *On Ethics and Economics*. Oxford: Basil Blackwell.
- Sen, Amartya. 2000. "The Discipline of Cost-Benefit Analysis." *Journal of Legal Studies* 29(S2): 931–952.
- Sidgwick, H. 1907. *The Methods of Ethics*. London: Macmillan.
- Singh, Harinder och Roger Frantz. 1991. "The Conflation of Productivity and Efficiency in Economics and Economic History." *Economics and Philosophy* 7: 87–89.
- Smith, Hugh. 2005. "What Costs Will Democracies Bear? A Review of Popular Theories of Casualty Aversion." *Armed Forces and Society* 31: 487–512.
- Titmuss, Richard M. 1970. *The Gift Relationship: From Human Blood to Social Policy*. London: George Allen & Unwin.
- Viscusi, W. K. 2000. "Risk Equity." *Journal of Legal Studies* 29: 843–871.
- Walley, Peter. 2000. "Towards a Unified Theory of Imprecise Probability." *International Journal of Approximate Reasoning* 24: 125–148.

- Wilkinson, Richard G., and Pickett, Kate. 2009. *The Spirit Level: Why More Equal Societies Almost Always Do Better*. London: Allen Lane.
- Williams, Bernard. 1976. "Moral Luck." *Aristotelian Society, Supplementary Volume* 50: 115–135.
- Zelizer, Viviana A. 1978. "Human Values and the Market: The Case of Life Insurance and Death in 19th-Century America." *American Journal of Sociology* 84: 591–610.

CHAPTER 4

IN PRAISE OF IMPERFECT COMMITMENT

An Ethic of Power, Professionalism and Risk

SHARON D. WELCH

THERE is one point on which we have political consensus in the United States—there is little, if any, satisfaction with the economic status quo, and there are active calls for fundamental social and political change. Part of the challenge facing us, however, is not just the specific issue of economic policy but the means through which professionals address the making of public policy, and how we as a culture reckon with unpredictability, risk, and intrinsic epistemic fallibility.¹

In this chapter I explore what it takes for economists, other professional policy advocates, and social activists to engage in forms of social analysis and political engagement that acknowledge the limitations of our knowledge of the world around us, and our inability to control complex social and economic processes.

Economists are likely unaccustomed to being addressed alongside social activists, to be sure. Here I will do just that. My reasons for doing so are twofold. First, there is a significant development within the world of social activism. Many of us are recognizing our need for the insights and expertise of economists as we move from social critique to the more difficult work of designing and implementing sustainable and equitable public policies. Second, there are equally significant developments within the world of economic theory and practice, with many economists recognizing a twofold intrinsic ethical challenge within policy implementation. This challenge is one shared with social activists: how to reckon with, responsibly and equitably, two factors: (1) the

¹ Reverend Dr. Lee Barker, president of Meadville Lombard Theological School, coined the term “in praise of imperfect commitment” to describe the vibrant, non-utopian model of social engagement that we are teaching Unitarian Universalist ministers and community activists at Meadville Lombard Theological School.

intrinsic unpredictability of our policy interventions and (2) the differential effects of policies—that which helps some, often hurts others.

The work that we do as activists and professionals can be fundamentally recast because of where we are now culturally and politically. We are immersed in a third wave of political engagement, an era that builds on two prior waves of progressive politics and yet has its own energy and challenges. The first wave was the forceful denunciation of manifold forms of social injustice—slavery, the oppression of workers and the secondary status of women—all forms of oppression defended as divinely ordained or part of the natural order of things. These struggles for social justice have been augmented by a second wave of activism, the work of identity politics, the resolute claim for the complex identities and full humanity of all groups marginalized and exploited by systemic oppression and silenced through cultural imperialism.²

These political tasks are ongoing, yet now occur within a third paradigm. Once we recognize that a situation is unjust, once we grant the imperative of including the voices and experiences of all peoples, how then do we work together to craft and implement just and creative social policies?

For far too long, activists like myself have exercised a form of the prophetic imagination that focused primarily on critiques of what is wrong, and visions of what might be right, and paid little or no attention to the third, most important and most difficult element of the prophetic role—the implementation of social policies that are just, sustainable and resilient. In contrast, professional economists have been very willing to design and implement grand policy solutions to pressing social problems without taking adequate account of the limits of their science and control over the world they seek to improve.

An example for activists of a full prophetic imagination that includes equal attention to critique, vision and implementation can be seen in the current work of the Unitarian Universalist Social Committee. In the 2012 annual report, UUSC's president and CEO William Schulz lays out clearly three elements of constructive social activism:

First, engagement. We're eager to use the people power at our disposal to optimize our effectiveness. Our members, most of whom are associated with Unitarian Universalist congregations, are natural born activists. They're itching to get their hands dirty, be it on their computer keyboards taking online actions or by building an eco-village in Haiti.

Second, innovation. Wherever we go in the world, we ask ourselves, "who's been forgotten and who is doing the most creative, groundbreaking work to transform and empower those forgotten populations." By finding the most innovative, entrepreneurial approaches to problems and crises . . . we encourage new solutions to old quandaries; we engage with communities of women or ethnic minorities, too often marginalized in their societies.

² For a thorough analysis of cultural imperialism and other forms of systemic injustice, see the definitive work of Iris Marion Young, *Justice and the Politics of Difference*. New edition with a Foreword by Danielle S. Allen (Princeton, NJ: Princeton University Press, 2011).

Third, impact. . . we're experimenting with different approaches to measuring impact because we know that at the end of the day the only thing that really counts is how many lives we've actually changed. (Schulz and Spence, 2012: 2)

A threefold form of prophetic engagement, with far more attention to implementation and impact than to critique or vision, is possible because of nuanced histories of successful social movements and empirical evaluations of the impact of policy interventions. This constructive form of prophetic engagement is leading activists to take seriously what economists have long known, the necessity of attention to implementation and impact. It may also lead economists to pay greater attention to what activists have long known—the imperative of learning from and with those that we professionals purport to understand and serve.

As we take seriously the role that we can play as economists and activists in designing and implementing just social policies, it is important to ground ourselves in two factors—the dynamics of social movements, and the particular role that professionals play in successful social analysis and political engagement.

First, let us look at some of the groundbreaking discoveries in the history of social movements. In his recent book, *The Honor Code: How Moral Revolutions Happen*, Kwame Anthony Appiah charts the transformation of “private sentiment into public norm” in the abolition of dueling in England, footbinding in China, slavery and the slave trade in England, and examines what may well be leading to a similar shift in the honor killing of women. In each case he finds that moral critiques were not enough. In fact, each practice had been discredited by religiously and philosophically based moral critiques decades before the practices came to be seen not just as wrong, but as unseemly.

Changes in behavior occurred only when a form of action that had been accepted for hundreds of years was seen not only as immoral, but as dishonorable, unworthy of respect, and unfitting for all, including elites (2010: 178).

In his massive study of the decline of various forms of physical violence, the social psychologist Steven Pinker explores what has led to significant changes in social norms. Pinker examines the history of the decline of slavery, capital punishment, and torture as accepted and routine forms of political and economic life. He also examines what he calls the ongoing “rights revolution:” “civil rights, women’s rights, children’s rights, gay rights and animal rights” (2011: xvi).

In each case Pinker finds a complex interaction of two phenomena—an expansion of the circle of sympathy, and an escalator of reason. In successful social movements, large numbers of people come to value the lives of people formerly ignored and devalued. Such emotional openness, however, to the dignity and struggles of others, is not enough for fundamental social change. Pinker writes that

. . . the limited reach of empathy, with its affinity for people like us and people close to us, suggests that empathy needs the universalizing boost of reason to bring about changes in policies and norms . . . These changes include not just legal prohibitions against acts of violence but institutions that are engineered to reduce the temptations

of violence. Among these wonkish contraptions are democratic government, the Kantian safeguards against war, reconciliation movements in the developing world, nonviolent resistance movements, international peacekeeping operations, crime prevention reforms . . . (691, 695).

Pinker is clear—empathy can be too easily limited to family, to tribe, to those near and dear. The challenge of the escalator of reason is to build on an expanded sense of sympathy and to find social behaviors, policy changes, rules, and laws that enable greater flourishing for the many, and not just the few (695). As we take up this work of the escalator of reason, finding what actually enables people to flourish in ways that are equitable as well as ecologically and economically sustainable is more a matter of critical experimentation and risk-taking than it is a matter of moral or ideological certainty. Take as a case in point the work of the feminist economists, J. K. Gibson-Graham.

CREATIVE EXPERIMENTATION

J. K. Gibson-Graham (Katherine Gibson, University of Western Sydney, and the late Julie Graham, University of Massachusetts, Amherst, writing as a single persona since 1992) have described a new political imaginary. They analyze, nurture, and celebrate the reality, opportunities, and challenges of community economies. People all over the world are finding ways of shaping their economic lives to recognize the power of interdependence, not a “common being” but a “being in common.” J. K. Gibson-Graham describe different ways of “explicitly recognizing and acting on our interdependence:” “Employee buyouts in the United States, worker takeovers in the wake of economic crisis in Argentina, the anti-sweatshop movement, shareholder movements that “promote ethical investments and police the enforcement of corporate environmental and social responsibly,” the living wage movement, discussions of a universal basic income, social entrepreneurship—all part of a “community economy” that “performs economy in new ways” (2006: 80–81). I would add to this list the growth of B Corporations (the most well known being Patagonia). According to an editorial in the January 2012 issue of *The Economist*, there are now several hundred corporations that have changed corporate laws to “create the legal framework for firms . . . to remain true to their social goals. To qualify as a B Corporation a firm must have an explicit social or environmental mission, and a legally binding fiduciary responsibility to take into account the interests of workers, the community and the environment as well as its shareholders” (2012: 1).

J. K. Gibson-Graham build on the insights of queer theory and political and feminist theory and organizing, emphasizing that shared questions often lead to different answers. Just as there is no one way to be a feminist, there is no single way to perform economic relations justly. There are, however, salient questions, choices to be made in