SEARCHING FOR SAFETY

Social Theory and Social Policy

Aaron Wildavsky

SEARCHING FOR SAFETY

Series Editor: Ellen Frankel Paul

The Social Philosophy and Policy Center, founded in 1981, is an interdisciplinary research institution whose principal mission is the examination of public policy issues from a philosophical perspective. In pursuit of this objective the Center supports the work of scholars in the fields of political science, philosophy, law, and economics. In addition to the book series, the Center hosts scholarly conferences and edits an interdisciplinary professional journal, *Social Philosophy & Policy*. For further information on the Center, write to the: Social Philosophy and Policy Center, Bowling Green State University, Bowling Green, OH 43403.

Other Books from the Social Philosophy and Policy Center (available from Transaction Publishers)

Laissez Parler: Freedom in the Electronic Media
David Kelley and Roger Donway

Preparing for the Future:
An Essay on the Rights of Future Generations
John Ahrens

Out of the Mouths of Babes: The Infant Formula Controversy
Fred D. Miller, Jr.

Opportunity or Privilege: Labor Legislation in America
Charles W. Baird

The New Crusaders: The Corporate Social Responsibility Debate
Douglas J. Den Uyl

Nuclear War: The Moral Dimension
James W. Child

Plant Closings: Worker Rights, Management Rights and the Law
Francis A. O'Connell, Jr.

The Art of Judging
James E. Bond

The Catholic Bishops and the Economy: A Debate
Douglas Rasmussen and James Sterba

SEARCHING FOR SAFETY

AARON WILDAVSKY

Routledge
Taylor & Francis Group
LONDON AND NEW YORK

First published 1988

First published 2017 by Transaction Publishers

Published 2017 by Routledge
2 Park Square, Milton Park, Abingdon, Oxon OX14 4RN
711 Third Avenue, New York, NY 10017, USA

Routledge is an imprint of the Taylor & Francis Group, an informa business

Library of Congress Catalog Number: 87-37659

ISBN: 978-0-912051-18-5 (pbk)

Library of Congress Cataloging-in-Publication Data

A catalog record for this book has been requested

To the memory of my friend, Bill Havender.

Contents

Acknowledgments

Learning about (and keeping up with) the literature on risk is a monumental task. Even if one excludes the political and social aspects of the subject, arguably its most important, and concentrates, as I do in this book, on the consequences of technology for health and safety, keeping up with risk, where a new hazard is seemingly a daily discovery and views about old ones are continuously revised, has become a full-time occupation. But not, as a political scientist and policy analyst, for me.

In order to broaden greatly the range of phenomena that could be treated, as well as to make up for my lack of knowledge, I sought collaborators. After drafts of the first four chapters were completed, I asked four people—Dennis Coyle, a graduate student in the Political Science Department; William Havender, a geneticist and consultant on risk; Elizabeth Nichols, a graduate student in the Sociology Department; and Dan Polisar, an undergraduate at Princeton—to work with me respectively on how the human body defends itself, on how measures designed to reduce risk go wrong, on whether adding safety measures improves safety, and on why the law of personal injury harms people. I have been extremely fortunate in my collaborators. All of them have made my work better than it otherwise would have been. (Elizabeth Nichols, with whom I continue to work on the inspection of nuclear power plants, also prepared material for me on natural disasters.) Because of the special expertise required, I asked Bob Budnitz, a nuclear physicist and former head of research for the Nuclear Regulatory Commission, to write a section on substituting performance standards for detailed specification in regulation of nuclear power plants. Without them, this book would have been much poorer.

I am also grateful to a number of expert critics who gave valuable advice on parts of the manuscript: Eugene Bardach, Robert Budnitz, Dennis Coyle, Douglas Easterling, Amitai Etzioni, Robert Goodin,

Peter Huber, Robert H. Nelson, Emory Roe, Michael Tartar, and Michael Thompson. The editor of the series in which this volume appears, Ellen Paul, also made helpful suggestions. Responding to their comments has led to considerable improvement.

I have left the special contribution of William Havender for last so it will stand first and foremost. Bill and I met five years ago when I organized a seminar on risk for the Survey Research Center. For most of that time, until illness overtook him, we met every two or three weeks to discuss risk. I learned more about risk from him than from any single person. His detailed critique of the entire manuscript is the toughest and most useful of any I have ever received.

A geneticist by training, a social scientist by avocation, Bill Havender was far too sophisticated to believe that "the facts" interpreted themselves or that the preponderance of evidence would necessarily prevail in the political system. Nevertheless, now that he is gone (Bill died in April 1987), his passion for truth remains with me. Our first game was to take unfolding events about this or that substance or event (How soon before the latest sweetener was alleged to be worse than saccharin or cyclamates?), and place our bets on what the evidence would show after further studies had been reported. (Would Chernobyl also turn out to be a consequence of trying to improve safety?) He was right, so far as I could see, more often than anyone whose opinions appeared in the public prints. Our last game—trying to figure out why safety measures so often produced the opposite result—appears in Chapter 9. There the argument has been formalized, but I hope some of the fun shines through. Dedicating *Searching for Safety* to Bill is my way of saying that his search will live on.

Financial support for this book has come from three sources: The National Science Foundation, grant No. PRA-8412418; the Survey Research Center; and the Smith Richardson Foundation, through the Institute for Educational Affairs. As always, Joshua Menkes at NSF and Director Percy Tannenbaum and Assistant Director James Wiley at the SRC facilitated my efforts. Doris Patton, my secretary, bore with numerous revisions and names of creatures and substances not found in ordinary parlance, let alone polite company. I try but do not always succeed in doing my job as well as she does hers.

Introduction: The Jogger's Dilemma *or* What Should We Do When the Safe and the Dangerous are Inextricably Intertwined?

...Consider...the publicity over the death of James Fixx this summer. Mr. Fixx, an ardent runner and author of several popular books on running, died of a heart attack while running. His death stirred considerable controversy in both the lay and the medical press on the risks and benefits of exercise.

—*Wall Street Journal*, October 4, 1984, article by Jerry E. Bidrop

Were I asked why anyone would want to read this book, I would answer that it offers a theory accounting for the considerable degree of safety achieved in contemporary society, as well as the danger introduced by technological development. Almost all books on risk treat it as a bad thing to be avoided or diminished, rather than as an inevitably mixed phenomenon from which considerable good, as well as harm, is derived.

Since there can be no safety without risk, as I shall show, the importance of thinking about better and worse ways of searching for safety gives this subject an open-ended quality that is in danger of being foreclosed by thinking in terms of "all good" or "all bad." For if,

as I will further show, both good and bad inhere in the same objects, strategies, and processes—albeit in different degrees and in uncertain combinations—the only way to discover how to achieve more "good" is to search for safety—which, by my way of thinking, involves risk.

To those who see safety as a goal we know how to achieve, the task of decision making is simple: choose the safer path. There is no need to look further, since we already know what to do; no act designed to secure safety could have the opposite effect. Imagine instead, however, safety as largely an unknown for which society has to search. An act might help achieve that objective, or it might not. Often, we would not know which, safety or danger, had in fact been achieved until afterwards. Given this uncertainty, we would always have to be as concerned that the existing level of safety not be lowered as that it might be raised.

The most important part of the safety enterprise is thinking about how to think (and, therefore, how to act) about risk. As things stand, the dangers of new things are viewed in isolation from, and with no attention to, potential advantages. Yet playing it safe, doing nothing, means reducing possible opportunities to benefit from chances taken, and can hurt people. Just as the unknown consequences of technological advance have to become part of our safety calculus, so must "safety risks"—possible damage resulting from measures designed to improve safety—enter into our considerations.

Thinking about risk, I contend, has been one-sided: safety has been over-identified with keeping things from happening. My aim is to redress this imbalance by emphasizing the increases in safety due to entrepreneurial activity. If this essay in persuasion achieves its purpose, the focus of the risk debate will shift from the passive prevention of harm to a more active search for safety.

This book is about how risk and safety are produced, about the fact that they are intertwined, and about what, therefore, should be done to make the search for better combinations both efficient—devoting resources to the worst hazards—and effective—actually improving safety. Because safety must be discovered, and cannot be merely chosen, I shall argue that trial-and-error risk taking, rather than risk aversion, is the preferable strategy for securing safety. Encouraging trial and error promotes resilience—learning from adversity how to do better—while avoiding restrictions that encourage the continuation of existing hazards. Increasing the pool of general resources, such as wealth and knowledge, secures safety for more people than using up resources in

a vain effort to protect against unperceivable, hypothetical dangers. Wealth adds to health.

Objective Risk

My subject is the objective aspects of safety and of danger (popularly known as "risk"). By "objective" I do not mean to assert my own particular respect for facts. Nor would I suggest that all people see eye to eye about what risks under which circumstances should or should not be taken. Nevertheless, almost all participants in the risk debate do claim a respect for facts and attempt to legitimate their policy preferences by reference to canons of scientific inquiry. Citing ghosts or gremlins would not do. Assuming that the common objective is to secure safety, let us all try to see what sort of actions would best achieve that goal.

By objective aspects of risk, therefore, I mean both observable dangers as well as the observable consequences of actions undertaken for the ostensible purpose of increasing safety. Do, or do not, actions undertaken in the name of risk reduction achieve safety? In the end, to be sure, evaluation of strategies is done by human beings, like myself, who bring biases to bear on their perceptions. Though I shall try to be persuasive, I make no claim to some transcendental objectivity. It is the certainty of the question—How best to secure safety?—not the uncertainty of the answers that leads me to talk about objective risk. Except in passing, therefore, I shall not deal with the subjective aspects of safety: What sorts of people are risk-taking or risk-averse for which reasons in regard to different matters (say nuclear energy or acquired immune deficiency syndrome). Nor shall I emphasize the overtly political aspects of the controversies over risk, though mobilizing support is indeed a vital and often decisive activity. In this book I shall concentrate instead on what ought to be done to improve safety regardless of its immediate political feasibility.[1]

That this book is not overtly concerned with political action does not signify that it is devoid of political ideas. On the contrary, I have deliberately set out to redress existing risk-averse biases by countering them with my own risk-taking bias. Nor does the undoubted fact that all positions in the risk debate reflect views of the good life imply that all are equally correct, i.e., that all would result in the same consequences for safety. Not so. What sort of strategies work best in securing safety remains a problem subject to investigation. A good place to start is with the multiple meanings of risk.

"Safety," E. Siddall writes, "is the degree to which (a) temporary ill health or injury, (b) chronic or permanent ill health or injury, or (c) death are controlled, avoided, prevented, made less frequent or less probable in a group of people."[2] The reference to a particular group rather than to an entire society makes us aware that risk may have distributive consequences. "Safety for whom?" is always a good question.[3]

Following current usage, William Rowe defines risk as "the potential for harm."[4] According to this conception, risk is the probability of encountering negatively-valued events. This definition has the virtue of being clear, concise, and general. Because it views risks as never carrying associated benefits, however, it also has the disadvantage of prejudging the conclusion of inquiries. To take risks thus stands condemned as being bad, the only question being, how bad. Yet, in ordinary usage, we often speak of taking risks in order (hopefully) to secure gains. If and when we succeed, risk taking can help improve safety. Altering the definition to refer to "the potential for harm and/or for safety" complicates matters. I favor this definition, however, because it compares gains as well as losses, accepting the possibility that a single act or thing can cause both positive and negative safety consequences.[5]

Consider the case of the "rational" potato. Like all growing things, it could not survive the process of evolution unless it was able to ward off predators. Unable to run away or fight directly, the homely potato has evolved chemical defenses. When mother told us that the potato's vitamins were concentrated in the jacket, she was right. What she did not know, however, was that the poisons the potato uses to ward off predators also were in the jacket. Under unusual conditions, these poisons accumulate and can be dangerous to human life. Nevertheless, when we are urged to eat potato jackets because it is healthy, I concur, providing only we understand that the potato's poisons and its nutrients are largely in the same place.

The Principle of Uncertainty, the Axiom of Connectedness, and the Rule of Sacrifice

Once we stop thinking of health and safety as qualities we already know how to achieve—and government's task being simply to choose what it confidently knows is life-enhancing, while rejecting what it already knows is life-denying—the question of how to search for safety becomes paramount. Which set of principles, axioms, and rules, I ask,

helps us discover how to reduce risk overall so that society as a whole becomes safer? Though no one can say for sure, my candidates in opening up the inquiry on risk and safety are: the principle of uncertainty, the axiom of connectedness, and the rule of sacrifice.

Dangers will arise; everyone agrees on that. Uncertainty about the consequences of present acts and about others as yet unforeseen cannot be reduced to zero. The principle of irreducible uncertainty is based not only on the self-evident premise that no one knows it all, but also on the slightly less obvious consideration that even as human beings act on their environment they are creating consequences of which they are as yet not fully aware. Although some uncertainties may be reduced under some circumstances, a modicum of uncertainty is a universal condition. Hence Kenneth Boulding writes about "irreducible uncertainties."[6]

To the principle of uncertainty I wish to add another—the axiom of connectedness. This states that the good and the bad (safety and harm) are intertwined in the same acts and objects.[7] There may be unalloyed goods (though I doubt it) but, if so, they are few and far between. Take the two principles together—uncertainty cannot be eliminated and damage cannot be avoided. This combination stipulates the conditions (old risks cannot be reduced without incurring new ones) under which the question of how to increase safety should be considered.

Suppose the things of this world, and the practices people follow, were all good or all bad in terms of securing safety. Then nothing would matter except to discover and choose the safe, and reject the dangerous. The end—securing safety—is given, and a decision rule exists for choosing among available means.

But life is not so straightforward. For the most part, safety and danger coexist in the same objects and practices. Under the right (or wrong) conditions, everything we need for life can also maim or kill: water can drown, food can poison, air can choke. Babies cannot be born without risk to the mother, nor can they grow to adulthood without facing innumerable dangers.

The trick is to discover not how to avoid risk, for this is impossible, but how to use risk to get more of the good and less of the bad. The search for safety is a balancing act. For if the axiom of connectedness holds, there is no choice that results in no harm. Merely minimizing danger to some people would not meet the safety criterion if it resulted in less safety for a larger number of other people than there otherwise would have been. If every act and every thing has harm for someone

somewhere, indeed, if the safety we seek is ineluctably bound up with the danger that accompanies it, more harm than good may be done by efforts to avoid risk.

The principle of uncertainty and the axiom of connectedness can be generalized to all (including human) systems: each part of every system cannot be stable at the same time. Economist Burton Klein expresses the thought well:

> It is true, of course, that in an imaginary, unchanging world economies can be predictable both in the small and in the large: that is, they can survive simply by taking the classical law of supply and demand as a given. However, if an economic system is to make smooth adjustments in dealing with new circumstances—if it is to remain predictable in the large, so to speak—it must be able to adapt itself to new circumstances. In fact, what I mean by "dynamic" *is the ability of a person, a firm or an economy to adapt itself to new circumstances by generating new alternatives.* But, it should be apparent that if predictability in the small is defined as "microstability" and predictability in the large as "macrostability" neither an individual firm nor an entire economy can simultaneously conserve its micro and macrostability. Only in heaven can microstability be equated to macrostability. Here on earth, the greater the insistence on microstability—the greater the insistence on preserving a way of life—the lower will be its macrostability. Conversely, if a system is to enjoy a high degree of macrostability it must enjoy the ability to generate new alternatives, when confronted by necessity.[8]

I will call this—the safety or macrostability of the whole being dependent upon the risk taking or instability of the parts—the "rule of sacrifice."

Does risk in the human context mean that specific individuals must give up their lives or limbs for a collective entity? No. It does mean that if the parts of a system are prevented from facing risks, the whole will become unable to adapt to new dangers. The concern with harming individuals for the benefit of the collective is warranted (see Chapter 10 for discussion) but misapplied. For harm is being done anyway; existing hazards are already taking their toll. The only real alternatives involve understanding the balance among those helped and hurt. Since there is no way of avoiding harm for everyone, the search for safety has to proceed on the understanding that the rule of sacrifice is inexorable: there can be no stable whole without some unstable parts.[9]

Rival Strategies for Securing Safety

Assertions about risk should be treated as hypotheses. That is why I have decided to focus on the agreed objective—securing safety—rather than the disputed means. Is "taking no chances," for instance, conducive to safety? In the first chapter, "Trial and Error Versus Trial Without Error," I argue that it is not. Many predicted dangers have turned out to be exaggerated or nonexistent. Others are real, but the policies invoked merely shifted dangers from one place or people to another. Still other, worse dangers resulted because the products or practices protested against were replaced by more harmful substitutes. There are, to be sure, unanticipated consequences of technology that do more harm than good. But there are also, as Michael Thompson calls them, "unconsequent anticipations," i.e., predictions of unsafe events that never come to pass. When government acts to prevent all imaginable dangers, the costs associated with such predictions are no longer cheap. The cost is counted not only in money but also in damage done by those very preventive measures, and in harm that might have been mitigated had resources not been used up on unrealized risks.

Acting safely, it is said, requires rejection of risky things—chemicals, processes, apparatus—that could cause grave damage. Instead of making progress by trial and error, some proponents of safety propose a rule of allowing no trials without prior guarantees against harm. This across-the-board risk aversion actually increases danger, I shall argue, because in thus achieving "safety" the harm avoided is part and parcel of even larger, as yet unknown, benefits that will be forgone. By taking a narrow view—a certain thing does harm—instead of a broad perspective—what is the balance of harm and help here compared to alternatives—measures designed to increase safety can lead to an overall decrease. When a strategy of "no trials without prior guarantees against error" replaces "trial and error," the opportunity to take risks in order to achieve beneficial consequences is lost.

This is the theme of the second chapter, "Opportunity Benefits Versus Opportunity Risks."

The risks stemming from trial and error should also be discounted by the "safety risks" that flow from trying to prevent damage. For if all things are potentially dangerous, merely giving a measure a safety label is no guarantee that it will not do harm. Rem Khokhlov, an eminent mountain climber and scientist, for instance, was also part of the Soviet political elite. When he suffered a pulmonary embolism during

a training mission at 20,000 feet, a helicopter was sent for him, after which he was taken to Moscow for treatment. "It has been suggested to me by several Soviet friends," Peter Franken wrote in *Science*, "that Khokhlov's seniority in the Soviet system was partly responsible for his death in that it led to his being treated by physicians who were less familiar with this particular illness than was the medical community nearer the site of the accident."[10]

Overwhelming evidence shows that the economic growth and technological advance arising from market competition have in the past two centuries been accompanied by dramatic improvements in health—large increases in longevity and decreases in sickness. One might expect a focus on profit to have led to neglect or even disregard of safety. Why, then, does society nevertheless usually end up healthier? Economists, so far as I know, have not turned their attention to the relationship between markets and safety: Does economic competition increase or decrease safety?

The positive association between market processes and economic wealth is well documented. "Richer Is Sicker Versus Richer Is Safer," the third chapter, begins by postulating another truth, just as well-established, though less well-publicized: human health is a function of economic wealth. Whether comparisons are made between nations or among people within them, wealthier is healthier. Why this is so, however, is not self-evident.

Disputes over the moral implications of economic costs as a consideration in regulation of risk (Should life be sacrificed to material gain?) do not make sense if health and wealth are positively related. In a wealth-health analysis, I ask whether a regulation or device or practice adds more to human safety than its cost subtracts from that safety.

The common wisdom is captured in the proverb "an ounce of prevention is worth a pound of cure." As is usually true with proverbs, however (he who hesitates is lost, but, look before you leap), they do not state the conditions of applicability. Cure may well be better than prevention if the former is feasible and the latter is not or if cure increases flexibility in dealing with future dangers while prevention induces rigidity.

Is it better, I ask in Chapter 4, to attempt to anticipate dangers before they occur or to inculcate a capacity to respond resiliently, i.e., to learn from experience to cope with untoward events? If seeking to anticipate dangers saps a system's energies without enabling it to guess right, for

instance, then the system might end up without the benefits of either anticipation or resilience. And, further, if anticipation works well enough to suggest that stability is permanent, unexpected challenges may yet overwhelm the system. How, then, can society protect itself against unknown dangers?

Human beings can engage in prevention. They can plan. But how well? Planning depends on prediction, which presumes the elimination (or, at least, reduction) of surprise. I am not talking about minor (a small change in probabilities) but about major surprise, a change in kind, a change like acquired immune deficiency syndrome that central decision makers could not imagine might occur. Confusing quantitative surprise (where we surmise what might happen but not its probability) with qualitative surprise (where we have no idea of the kind of event) trivializes the problem of unexpected danger. If only "expected surprise" existed, there would be a lot less to worry about. Fire drills may protect against expected surprises, but not against the qualitatively unexpected.

How, then, might members of society develop defenses without knowing what might be in store for them? How might a society sample the unknown so as to get hints about future dangers before they become massive? One good question deserves another. Who has both the capacity and the incentive to undertake this strange sort of task? Who might so benefit from outguessing the future as to be willing to absorb the losses from what might be a high rate of failure? The larger and more centralized the organization that seeks to predict the future, the longer it will take to get agreement, the fewer hypotheses it can try, and the more costly each probe is likely to be. Instead of assuming that anticipation must be centralized, therefore, I would like to open up the possibility that safety might be improved by spreading the anticipatory function more widely throughout society.

Decentralized anticipation (numerous independent probes of an uncertain future) can achieve a greater degree of safety. Since innovations are introduced piecemeal as a result of these independent probes, a larger proportion of emergent dangers can be perceived early while each hazard is still localized and small. The ability of market competition to interrogate the unknown at a low cost to society, while simultaneously encouraging individuals to overcome adversity, has been underappreciated.

Anticipation and resilience, the broadest strategic alternatives for attempting to secure safety, subsume other strategies. Decentralized,

rapidly moving trial and error contributes to a strategy of resilience. Centralized, slow-moving regulation of trials to prevent errors is essential to a strategy of anticipation. Are these strategies mutually exclusive, or can they be combined?

Why, the reader may well ask, are these strategies dichotomized into extremes instead of joined together? Obviously, the intelligent person would want to combine anticipation with resilience. Even though, in the end, a mixed strategy should be adopted, I think these large questions, with their polar-opposite alternatives, are good to ask.

Each extreme alternative has real advocates. The strategy of "no trials without prior guarantees against error" has substantial support; as the second chapter shows, it is written into law. The debate about whether the "no threshold effect" (a single molecule of a carcinogenic substance, some say, is enough to kill) should be used as the major criterion of regulatory choice is instructive in this regard.

But if we suppose, following Ronald Heiner's seminal essay on "The Origin of Predictable Behavior," that human uncertainty is deeper and more pervasive than has heretofore been thought,[11] reliance on simple rules of thumb makes more sense. Has the individual on the verge of decision reached the right conclusion about both his objectives (or preferences) and/or the means for obtaining them? Perhaps he has the right preferences, but at the wrong time or for different conditions. Since the consequences of actions are so entangled—undoubtedly the product of diverse causes, some recent, the others long past—our decision maker cannot be at all certain about whether past solutions have worked or to what degree. Evolution proves survival, to be sure, but survival alone is hardly convincing evidence that the strategy itself was beneficial. The doctrine of "adverse selection"[12] reminds us that organisms or practices may be selected because they inhere in desirable objects despite the fact that the item selected, were the truth known, is actually harmful.

Torn by uncertainty, decision makers place a high premium on reliability, i.e., on increasing the likelihood that a choice will work reasonably well under most conditions, though not necessarily exceptionally well under any. The chance that a given objective can be perfectly realized, Heiner argues, is exceedingly low compared to the probability of missing by so wide a margin that nothing is attained. Blame is reduced by reasonable success, not by failing to achieve perfection. The need for reassurance in an uncertain world, therefore, may lead to adopting

a polar strategy—resilience or anticipation—even though neither is expected to work supremely well in every instance.

Also, mixed strategies are not always feasible; when a decision maker must choose a single strategy, therefore, it is desirable to estimate which one would yield the best overall results. It may be advantageous to start with one possible strategy, modifying it as circumstances warrant. The danger is less in beginning than in ending with a single strategy. Polarizing proposed solutions is largely a device that ultimately enables us to devise mixed strategies.

Diverse Arenas for Studying Risk

The first four chapters in Section I contain examples drawn from current controversies over risk—regulation of chemical carcinogens, efforts to reduce pollution of different kinds, safety on the job, and so on—that appear to have reached an impasse in terms of productive thinking. Wishing as I do to alter how we view risk and safety, I have sought new areas in which to study rival efforts to reduce risk. Hence, the rather unusual subject matter in Section II—nonhuman life forms, the inspection of nuclear power plants, the human body, and the law of personal injuries—deserves a word of explanation.

I want to show that resilience and anticipation are universal strategies. One way to do this is to pick an area of life remote from current concerns, an area which contains both sets of strategies, fairly evenly matched, operating under different conditions. The safety (ecologists speak of systemic or species stability) of plants, animals, and insects is well-suited for this purpose. Also, I want to see what happens when a strategy of anticipation is used far more frequently than a strategy of resilience. The inspection of nuclear power plants, where one safety measure after another is used in an effort to ward off danger, meets the criterion of predominant anticipation. The other side of the risk coin—more resilience than anticipation—is exemplified by the human body. Though the body tries to anticipate by building some barriers against certain dangers, most of its rather economical defenses work toward mitigating harms as, and after, they occur. A fourth arena for studying risk illustrates the worst of the two worlds, in which there is neither much resilience nor much anticipation. The law of torts is avowedly a form of resilience; people cannot sue for damage arising from personal injuries unless they first show that there has been injury. Nevertheless, the tort law has been applied so as to cause the

worst consequences of anticipatory regulation with almost none of the benefits of resilience.

Anticipation and resilience are generic strategies, capable of being employed by diverse life forms under many different circumstances. These strategies are useful in classifying the ways in which "Nonhuman Life Forms Cope With Danger" (Chapter 5). The efforts of plants, insects, and animals to protect themselves have much to tell us about diverse approaches to safety.

Consider, for instance, a principle enunciated by Gerald Rosenthal: "No defense is inviolate."[13] The emerging discipline of chemical ecology studies the interaction between plants and their enemies. In certain desert areas the hairs or trichomes of plants store natural toxins, some of which cause allergic reactions that deter herbivores from feeding on them. Other defensive mechanisms are more indirect. Some plants produce analogues to juvenile hormones that keep the insect in perpetual youth, so that it cannot metamorphose and become a pupa. But juvenoids, which resemble juvenile hormones, must also kill their predators, for otherwise they merely prolong its larval stage, which happens to be the most destructive stage of the insect.[14]

Defenses may be counteracted. Even safety alarm systems can be used against an organism. When attacked by a predator, the aphid *Myzus persicae* secretes an alarm pheromone that informs other aphids of the coming danger. This aphid also preys on wild, tuberous potatoes. In apparent response, the wild potato releases a substance that mimics the aphid's alarm signal.[15] Organisms, we may conclude, have to live with incomplete defenses.

Are human beings so different? Can *Homo sapiens* devise foolproof defenses? Is safety always secured by multiplying the number of safety measures? The problem is not only whether measures purporting to increase safety always accomplish that purpose but, less obviously, whether measures that, taken singly, do increase safety have the same positive effect if deployed one on top of the other. A study of inspection and inspectors in nuclear power plants, the sixth chapter, reveals a curvilinear relationship: introducing a few devices tends to increase safety, but multiplying them decreases safety—they get in each other's way; the devices themselves become the causes of new failures.

If relying on anticipation has its drawbacks, what can we say about a preponderant strategy of resilience? "How the Human Body Defends Itself," the seventh chapter, is largely a story of resilience. Rather than evolving systems that never or rarely break down, biological processes

are accompanied by a variety of repair mechanisms. The most important body strategies, so to speak, appear to be "search and destroy" and "redundant repair"; they are based on learning how to bounce back from insults. Some dangers are anticipated, but most are left to learning through trial and error. Sometimes, however, a safety mechanism such as the immune system turns against the body it is designed to protect; in the body as in society, the axiom of connectedness holds, for the sources of life and death are intertwined.

The eighth chapter looks at the evolution of the law of torts, concerned with personal injury, from a mode of resilience to one of anticipation. Consequently, tort law now operates (read your local newspaper) as if defendants were subject to regulation—i.e., to enforceable prohibitions in advance of action.

The concluding chapters, comprising Section III, seek to set out principles that reduce harm and increase safety. The ninth chapter—"Why Less is More: A Taxonomy of Error"—accentuates the negative by developing principles specifying what not to do to improve safety. The normal revulsion to risk is misleading. Paradoxically, overconcentration on danger has led to neglect of safety: measures to increase safety often end up decreasing it, while courting danger may reduce it. In the same spirit, Michael Novak suggests that we pay less attention to the causes of poverty, which we do not wish to produce, and more to the causes of wealth.[16] An overwhelming concern with large consequence, low probability events screens out strategies that have probability of accumulating small, health enhancing benefits. To the degree that securing safety is our ultimate concern, principles guiding thought and action—the objective aspects to which I referred earlier—should be reconsidered and redirected, from risk aversion to risk taking.

By seeking to eliminate all but infinitesimal sources of risk, conformity to the single best safety strategy is sought. Thus, organizations in society become fewer and larger, and, in order to prove that products do no harm, much more capital is required. Consequently, small organizations are driven out. Diversity among organizations declines as they are subject to the same regulations. Hence the available responses to unexpected adversity diminish. The growth of knowledge slows down because fewer hypotheses are tested. Increases in wealth are held to a minimum, both because resources are devoted to anticipation and because of regulatory restrictions on market transactions. Because earlier efforts at prevention lead to less experience in coping with

unanticipated risk, resilience declines. Eventually there is also less ability to prevent damage through such anticipatory measures because past efforts to ward off innumerable dangers, most of which do not materialize, have led to internal exhaustion. Just as the effort to mobilize too many resources against threats can lead the body's defenses to turn on the body itself, destroying healthy cells and tissue, so, too, society, while preoccupied with eliminating harm, can inadvertently destroy its sources of safety. What understanding of risk and safety, I ask, would lead to better strategies for securing safety?

If all things are potentially risky, losses here may be made up by gains there. Advantages in one place may be given up in another. There is no problem in protecting a part (of an apparatus, a person, a group, a society). The difficulty lies in advancing the whole so that more people are gaining than losing at any one time, most people are safer over a period of time, and almost all people are better off than they would have been in decades past. Steady improvements in safety, I hypothesize, depend upon enlarging society's overall resource capacity so as to increase its resilience.

General and Global Resources

On the global level, safety is a function of general resources in a given society at a certain time. By general resources I mean knowledge, education, wealth, energy, communication, and any other resource that can be shifted around. General resources can be converted into other things. These are not just food crops, for example, but the capacity to grow food and to alter what one grows according to the conditions of the time. The global level stands for overall resource capacity—the ability to mobilize and to redeploy general resources.

Conceivably there could be a limit on global resources, a limit not only at an instant in time but for all time thereafter. While the materials out of which energy is transformed should last for eons, for instance, it is possible that at some time their cost will become prohibitive. While physical resources may be limited (a topic covered under the elusive subject of entropy), however, the true currency of human ingenuity is information, and I am not aware that it faces exhaustion.

Since my subject here is objective risk, assertions about limits have to be treated as subject to empirical inquiry. No limits are in sight, I think, even though one day they may reveal themselves. In this book, I shall assume, in accord with the evidence of the last several hundred

years, that global resources can and do grow. If this assumption is unwarranted, if mankind is on a declining resource curve, then drastic, perhaps desperate, measures might be indicated; for allowing things to go on as they are might be worse than using up resources in trying preventive measures that have very low probability of success. But I doubt it.

Hypotheses about global resource potential do not dispose of the question of how safety might best be achieved at any stipulated level of global resources. Whether an act or a program designed to increase safety actually works as intended cannot be determined in advance. Every act must be treated as a potential source of harm as well as help. Just as abuses are committed in the name of liberty, so measures taken in the name of health can still make us sick. Purported safety measures, in addition to local consequences—their effectiveness against the specific hazard at which they are aimed—may also affect the global level of resources. Since I hypothesize that this global level determines the safety of society overall, it follows that safety measures that cause a decline in general resources will decrease the net safety of society.

While acts designed to secure safety may have the opposite effect, so, too, acts that do not intentionally consider safety at all may actually increase it. When the global level of resources grows, so my hypothesis suggests, safety grows with it whether intended or not. "Richer is Safer" not necessarily because of good safety intentions, but rather because the global resource level has gone up, thereby increasing the capacity to undertake both anticipatory and resilient strategies. Conversely, measures intended to secure safety actually may increase harm—not only because they may be misguided, but also (and, for my thesis, more important, since mistakes are inherent in life) because they reduce the amount of global resources. In the tenth and final chapter ("The Secret of Safety Lies in Danger") I shall try to explain how and why citizens of nations in which there has been technological progress through trial and error have become healthier and safer than they or their forebearers used to be.

The jogger's dilemma brings us full circle to the essence of the relationship between courting danger and securing safety, for the two are different sides of the same coin. Too much or too strenuous exercise too soon is unsafe. Too little, too infrequently is also bad. So far, so simple. The complication is that during the limited time devoted to

the most strenuous exercise, the risk of heart attack rises. That's the bad news. The good news is that for the rest of the day, as well as the days in between regular exercise, the body is safer. "Although the risk of primary cardiac arrest is transiently increased during vigorous exercise," a study in *The New England Journal of Medicine* reports, "habitual vigorous exercise is associated with an overall decreased risk of primary cardiac arrest."[17] You cannot have the one—a safer organism—without the other—expanding its resilience by allowing it to face risks. As the experience of joggers shows, safety is the other side of risk.