

SOFTWARE AND MIND

Andrei Sorin

EXTRACT

Chapter 1: *Mechanism and Mechanistic Delusions*

**This extract includes the book's front matter
and chapter 1.**

Copyright © 2013, 2019 Andrei Sorin

**The free digital book and extracts are licensed under the
Creative Commons Attribution-NoDerivatives
International License 4.0.**

This chapter explains the mechanistic philosophy and its limitations, the mechanistic fallacies, and the difference between mechanistic and non-mechanistic phenomena.

The entire book, each chapter separately, and also selected sections, can be viewed and downloaded free at the book's website.

www.softwareandmind.com

SOFTWARE
AND
MIND

The Mechanistic Myth
and Its Consequences

Andrei Sorin

ANDSOR BOOKS

Copyright © 2013, 2019 Andrei Sorin
Published by Andsor Books, Toronto, Canada (www.andsorbooks.com)
First edition 2013. Revised 2019.

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, without the prior written permission of the publisher. However, excerpts totaling up to 300 words may be used for quotations or similar functions without specific permission.

The free digital book is a complete copy of the print book, and is licensed under the Creative Commons Attribution-NoDerivatives International License 4.0. You may download it and share it, but you may not distribute modified versions.

For disclaimers see pp. vii, xvi.

Designed and typeset by the author with text management software developed by the author and with Adobe FrameMaker 6.0. Printed and bound in the United States of America.

Acknowledgements

Excerpts from the works of Karl Popper: reprinted by permission of the University of Klagenfurt/Karl Popper Library.

Excerpts from *The Origins of Totalitarian Democracy* by J. L. Talmon: published by Secker & Warburg, reprinted by permission of The Random House Group Ltd.

Excerpts from *Nineteen Eighty-Four* by George Orwell: Copyright ©1949 George Orwell, reprinted by permission of Bill Hamilton as the Literary Executor of the Estate of the Late Sonia Brownell Orwell and Secker & Warburg Ltd.; Copyright ©1949 Harcourt, Inc. and renewed 1977 by Sonia Brownell Orwell, reprinted by permission of Houghton Mifflin Harcourt Publishing Company.

Excerpts from *The Collected Essays, Journalism and Letters of George Orwell*: Copyright ©1968 Sonia Brownell Orwell, reprinted by permission of Bill Hamilton as the Literary Executor of the Estate of the Late Sonia Brownell Orwell and Secker & Warburg Ltd.; Copyright ©1968 Sonia Brownell Orwell and renewed 1996 by Mark Hamilton, reprinted by permission of Houghton Mifflin Harcourt Publishing Company.

Excerpts from *Doublespeak* by William Lutz: Copyright ©1989 William Lutz, reprinted by permission of the author in care of the Jean V. Naggar Literary Agency.

Excerpts from *Four Essays on Liberty* by Isaiah Berlin: Copyright ©1969 Isaiah Berlin, reprinted by permission of Curtis Brown Group Ltd., London, on behalf of the Estate of Isaiah Berlin.

Library and Archives Canada Cataloguing in Publication

Sorin, Andrei

Software and mind : the mechanistic myth and its consequences / Andrei Sorin.

Includes index.

ISBN 978-0-9869389-0-0

1. Computers and civilization.
2. Computer software – Social aspects.
3. Computer software – Philosophy. I. Title.

QA76.9.C66S67 2013

303.48'34

C2012-906666-4

Don't you see that the whole aim of Newspeak is to narrow the range of thought?... Has it ever occurred to you ... that by the year 2050, at the very latest, not a single human being will be alive who could understand such a conversation as we are having now?

George Orwell, *Nineteen Eighty-Four*

Disclaimer

This book attacks the mechanistic myth, not persons. Myths, however, manifest themselves through the acts of persons, so it is impossible to discuss the mechanistic myth without also referring to the persons affected by it. Thus, all references to individuals, groups of individuals, corporations, institutions, or other organizations are intended solely as examples of mechanistic beliefs, ideas, claims, or practices. To repeat, they do not constitute an attack on those individuals or organizations, but on the mechanistic myth.

Except where supported with citations, the discussions in this book reflect the author's personal views, and the author does not claim or suggest that anyone else holds these views.

The arguments advanced in this book are founded, ultimately, on the principles of demarcation between science and pseudoscience developed by philosopher Karl Popper (as explained in "Popper's Principles of Demarcation" in chapter 3). In particular, the author maintains that theories which attempt to explain non-mechanistic phenomena mechanistically are pseudoscientific. Consequently, terms like "ignorance," "incompetence," "dishonesty," "fraud," "corruption," "charlatanism," and "irresponsibility," in reference to individuals, groups of individuals, corporations, institutions, or other organizations, are used in a precise, technical sense; namely, to indicate beliefs, ideas, claims, or practices that are mechanistic though applied to non-mechanistic phenomena, and hence pseudoscientific according to Popper's principles of demarcation. In other words, these derogatory terms are used solely in order to contrast our world to a hypothetical, ideal world, where the mechanistic myth and the pseudoscientific notions it engenders would not exist. The meaning of these terms, therefore, must not be confused with their informal meaning in general discourse, nor with their formal meaning in various moral, professional, or legal definitions. Moreover, the use of these terms expresses strictly the personal opinion of the author – an opinion based, as already stated, on the principles of demarcation.

This book aims to expose the corruptive effect of the mechanistic myth. This myth, especially as manifested through our software-related pursuits, is the greatest danger we are facing today. Thus, no criticism can be too strong. However, since we are all affected by it, a criticism of the myth may cast a negative light on many individuals and organizations who are practising it unwittingly. To them, the author wishes to apologize in advance.

Contents

	Preface	xiii
Introduction	Belief and Software	1
	Modern Myths	2
	The Mechanistic Myth	8
	The Software Myth	26
	Anthropology and Software	42
	Software Magic	42
	Software Power	57
Chapter 1	Mechanism and Mechanistic Delusions	68
	The Mechanistic Philosophy	68
	Reductionism and Atomism	73
	Simple Structures	90
	Complex Structures	96
	Abstraction and Reification	111
	Scientism	125
Chapter 2	The Mind	140
	Mind Mechanism	141
	Models of Mind	145

	Tacit Knowledge	155
	Creativity	170
	Replacing Minds with Software	188
Chapter 3	Pseudoscience	200
	The Problem of Pseudoscience	201
	Popper's Principles of Demarcation	206
	The New Pseudosciences	231
	The Mechanistic Roots	231
	Behaviourism	233
	Structuralism	240
	Universal Grammar	249
	Consequences	271
	Academic Corruption	271
	The Traditional Theories	275
	The Software Theories	284
Chapter 4	Language and Software	296
	The Common Fallacies	297
	The Search for the Perfect Language	304
	Wittgenstein and Software	326
	Software Structures	345
Chapter 5	Language as Weapon	366
	Mechanistic Communication	366
	The Practice of Deceit	369
	The Slogan "Technology"	383
	Orwell's Newspeak	396
Chapter 6	Software as Weapon	406
	A New Form of Domination	407
	The Risks of Software Dependence	407
	The Prevention of Expertise	411
	The Lure of Software Expedients	419
	Software Charlatanism	434
	The Delusion of High Levels	434
	The Delusion of Methodologies	456
	The Spread of Software Mechanism	469
Chapter 7	Software Engineering	478
	Introduction	478
	The Fallacy of Software Engineering	480
	Software Engineering as Pseudoscience	494

Structured Programming	501
The Theory	503
The Promise	515
The Contradictions	523
The First Delusion	536
The Second Delusion	538
The Third Delusion	548
The Fourth Delusion	566
The <i>GOTO</i> Delusion	586
The Legacy	611
Object-Oriented Programming	614
The Quest for Higher Levels	614
The Promise	616
The Theory	622
The Contradictions	626
The First Delusion	637
The Second Delusion	639
The Third Delusion	641
The Fourth Delusion	643
The Fifth Delusion	648
The Final Degradation	655
The Relational Database Model	662
The Promise	663
The Basic File Operations	672
The Lost Integration	687
The Theory	693
The Contradictions	707
The First Delusion	714
The Second Delusion	728
The Third Delusion	769
The Verdict	801
Chapter 8 From Mechanism to Totalitarianism	804
The End of Responsibility	804
Software Irresponsibility	804
Determinism versus Responsibility	809
Totalitarian Democracy	829
The Totalitarian Elites	829
Talmon's Model of Totalitarianism	834
Orwell's Model of Totalitarianism	844
Software Totalitarianism	852
Index	863

Preface

This revised version (currently available only in digital format) incorporates many small changes made in the six years since the book was published. It is also an opportunity to expand on an issue that was mentioned only briefly in the original preface.

Software and Mind is, in effect, several books in one, and its size reflects this. Most chapters could form the basis of individual volumes. Their topics, however, are closely related and cannot be properly explained if separated. They support each other and contribute together to the book's main argument.

For example, the use of simple and complex structures to model mechanistic and non-mechanistic phenomena is explained in chapter 1; Popper's principles of demarcation between science and pseudoscience are explained in chapter 3; and these notions are used together throughout the book to show how the attempts to represent non-mechanistic phenomena mechanistically end up as worthless, pseudoscientific theories. Similarly, the non-mechanistic capabilities of the mind are explained in chapter 2; the non-mechanistic nature of software is explained in chapter 4; and these notions are used in chapter 7 to show that software engineering is a futile attempt to replace human programming expertise with mechanistic theories.

A second reason for the book's size is the detailed analysis of the various topics. This is necessary because most topics are new: they involve either

entirely new concepts, or the interpretation of concepts in ways that contradict the accepted views. Thorough and rigorous arguments are essential if the reader is to appreciate the significance of these concepts. Moreover, the book addresses a broad audience, people with different backgrounds and interests; so a safe assumption is that each reader needs detailed explanations in at least some areas.

There is some deliberate repetitiveness in the book, which adds only a little to its size but may be objectionable to some readers. For each important concept introduced somewhere in the book, there are summaries later, in various discussions where that concept is applied. This helps to make the individual chapters, and even the individual sections, reasonably independent: while the book is intended to be read from the beginning, a reader can select almost any portion and still follow the discussion. In addition, the summaries are tailored for each occasion, and this further explains that concept, by presenting it from different perspectives.



The book's subtitle, *The Mechanistic Myth and Its Consequences*, captures its essence. This phrase is deliberately ambiguous: if read in conjunction with the title, it can be interpreted in two ways. In one interpretation, the mechanistic myth is the universal mechanistic belief of the last three centuries, and the consequences are today's software fallacies. In the second interpretation, the mechanistic myth is specifically today's mechanistic *software* myth, and the consequences are the fallacies *it* engenders. Thus, the first interpretation says that the past delusions have caused the current software delusions; and the second one says that the current software delusions are causing further delusions. Taken together, the two interpretations say that the mechanistic myth, with its current manifestation in the software myth, is fostering a process of continuous intellectual degradation – despite the great advances it made possible.

The book's epigraph, about Newspeak, will become clear when we discuss the similarity of language and software (see, for example, pp. 409–411).

Throughout the book, the software-related arguments are also supported with ideas from other disciplines – from the philosophies of science, of mind, and of language, in particular. These discussions are important, because they show that our software-related problems are similar, ultimately, to problems that have been studied for a long time in other domains. And the fact that the software theorists are ignoring this accumulated knowledge demonstrates their incompetence.

Chapter 7, on software engineering, is not just for programmers. Many parts

(the first three sections, and some of the subsections in each theory) discuss the software fallacies in general, and should be read by everyone. But even the more detailed discussions require no previous programming knowledge. The whole chapter, in fact, is not so much about programming as about the delusions that pervade our programming practices, and their long history. So this chapter can be seen as a special introduction to software and programming; namely, comparing their true nature with the pseudoscientific notions promoted by the software elite. This study can help both programmers and laymen to understand why the incompetence that characterizes this profession is an inevitable consequence of the mechanistic software ideology.

The book is divided into chapters, the chapters into sections, and some sections into subsections. These parts have titles, so I will refer to them here as *titled* parts. Since not all sections have subsections, the lowest-level titled part in a given place may be either a section or a subsection. This part is, usually, further divided into *numbered* parts. The table of contents shows the titled parts. The running heads show the current titled parts: on the right page the lowest-level part, on the left page the higher-level one (or the same as the right page if there is no higher level). Since there are more than two hundred numbered parts, it was impractical to include them in the table of contents. Also, contriving a short title for each one would have been more misleading than informative. Instead, the first sentence or two in a numbered part serve also as a hint of its subject, and hence as title.

Figures are numbered within chapters, but footnotes are numbered within the lowest-level titled parts. The reference in a footnote is shown in full only the first time it is mentioned within such a part. If mentioned more than once, in the subsequent footnotes it is abbreviated. For these abbreviations, then, the full reference can be found by searching the previous footnotes no further back than the beginning of the current titled part.

The statement “*italics added*” in a footnote indicates that the emphasis is only in the quotation. Nothing is stated in the footnote when the italics are present in the original text.

In an Internet reference, only the site’s main page is shown, even when the quoted text is from a secondary page. When undated, the quotations reflect the content of these pages in 2010 or later.

When referring to certain individuals (software theorists, for instance), the term “expert” is often used mockingly. This term, though, is also used in its normal sense, to denote the possession of true expertise. The context makes it clear which sense is meant.

The term “elite” is used to describe a body of companies, organizations, and individuals (for example, the software elite). The plural, “elites,” is used when referring to several entities within such a body.

The issues discussed in this book concern all humanity. Thus, terms like “we” and “our society” (used when discussing such topics as programming incompetence, corruption of the elites, and drift toward totalitarianism) do not refer to a particular nation, but to the whole world.

Some discussions in this book may be interpreted as professional advice on programming and software use. While the ideas advanced in these discussions derive from many years of practice and from extensive research, and represent in the author’s view the best way to program and use computers, readers must remember that they assume all responsibility if deciding to follow these ideas. In particular, to apply these ideas they may need the kind of knowledge that, in our mechanistic culture, few programmers and software users possess. Therefore, the author and the publisher disclaim any liability for risks or losses, personal, financial, or other, incurred directly or indirectly in connection with, or as a consequence of, applying the ideas discussed in this book.

The pronouns “he,” “his,” “him,” and “himself,” when referring to a gender-neutral word, are used in this book in their universal, gender-neutral sense. (Example: “If an individual restricts himself to mechanistic knowledge, his performance cannot advance past the level of a novice.”) This usage, then, aims solely to simplify the language. Since their antecedent is gender-neutral (“everyone,” “person,” “programmer,” “scientist,” “manager,” etc.), the neutral sense of the pronouns is established grammatically, and there is no need for awkward phrases like “he or she.” Such phrases are used in this book only when the neutrality or the universality needs to be emphasized.

It is impossible, in a book discussing many new and perhaps difficult concepts, to anticipate all the problems that readers may face when studying these concepts. So the issues that require further discussion will be addressed online, at www.softwareandmind.com. In addition, I plan to publish there material that could not be included in the book, as well as new ideas that may emerge in the future. Finally, in order to complement the arguments about traditional programming found in the book, I have published, in source form, some of the software I developed over the years. The website, then, must be seen as an extension to the book: any idea, claim, or explanation that must be clarified or enhanced will be discussed there.

Mechanism and Mechanistic Delusions

The Mechanistic Philosophy

Three doctrines have dominated Western science and culture since the seventeenth century: reductionism, atomism, and mechanism. Reductionism claims that every phenomenon can be represented, through various transformations (or reductions), as a combination of *simpler* phenomena. Atomism claims that there is an end to reductionism, that every phenomenon can be reduced, ultimately, to some *elementary* entities – some building blocks (or atoms) that cannot be divided into simpler parts. Mechanism adds to reductionism and atomism the claim that every phenomenon can be reduced to *mechanical* phenomena – and hence, ultimately, to the phenomena associated with the motion of bits of matter.

Mechanism was the fundamental doctrine of the Scientific Revolution. It was called the new philosophy, to distinguish it from the scholastic philosophy, which had dominated Europe throughout the Middle Ages. Scholasticism had been worthless as a scientific doctrine, because it prevented its followers from expanding their knowledge of the world. Mechanism, on the other hand, is ideally suited for research in the natural sciences, as it encourages a logical breakdown of complex problems into simpler ones.

With the mechanistic philosophy, scientists working in fields like astronomy, physics, and chemistry quickly found explanations for phenomena that had baffled their predecessors for centuries. These spectacular successes have continued ever since, and have served to establish mechanism as the undisputed method of science: “By the middle of the nineteenth century mechanics was widely acknowledged as the most perfect physical science, embodying the ideal toward which all other branches of inquiry ought to aspire. Indeed, it was the common assumption of outstanding thinkers, physicists as well as philosophers, that mechanics is the basic and ultimate science, in terms of whose fundamental notions the phenomena studied by all other natural science could and should be explained.”¹

The failure of mechanism to explain phenomena in other fields, notably in the human sciences, has been as spectacular as its success in the natural sciences; but this has done nothing to reduce its dominating position in our culture. Despite occasional claims to the contrary, mechanistic theories are still considered, just as they were in the seventeenth century, the only valid scientific theories: science *is* mechanism.

We are interested in mechanism because we want to understand the nature of our mechanistic *delusions*; in particular, the nature of our mechanistic *software* delusions. We suffer from a mechanistic delusion when we attempt to apply the mechanistic principles in a field where they cannot work. And, because of our blind faith in mechanism, much of present-day science is in reality a pursuit of mechanistic fantasies. Software, in particular, gives rise to *non-mechanistic* phenomena; and yet, our software theories are based entirely on mechanistic concepts.

One reason for our mechanistic delusions is that we fail to recognize the mechanistic nature of our theories. Mechanism today is rarely expressed through reductions to mechanical phenomena and particles of matter, or through analogies to machinery, as was the case earlier. So, if we continue to associate mechanism exclusively with these concepts, we are liable to overlook our current fallacies.

Mechanism today manifests itself mostly in the belief that everything can be represented as a hierarchical structure of entities; that is, as a neat structure of things within things. In this chapter, we will see that the hierarchical concept is logically equivalent to the traditional mechanistic concepts. Once we understand this, it will be easier to recognize the mechanistic nature of our theories. The hierarchical model will also help us to understand why so many phenomena *cannot* be represented with mechanistic theories.

¹ Ernest Nagel, *The Structure of Science: Problems in the Logic of Scientific Explanation*, 2nd ed. (Indianapolis: Hackett, 1979), p. 154.



The mechanistic myth has its roots in our desire to understand the world completely, and in our belief that this is possible. For, so long as we don't expect to understand each and every phenomenon, we can simply attribute to some transcendental powers those events that we cannot explain, and then carry on with our affairs.

As soon as we demand complete explanations, we are faced with the task of deciding what answers are good enough to count as explanations. The accepted answers in antiquity and in the Middle Ages were rather circular: the weight of an object is due to a quality that causes heaviness; an object bends or breaks according to a quality that causes elasticity; and so on. "Explanations" like these lead nowhere, obviously, and it is not surprising that Western science advanced so slowly during these periods. Since it is easy to find such circular answers, few scientists felt the need to seek the underlying causes, or to verify hypotheses through experiments.

The tendency in early times, thus, was to try to explain a phenomenon by seeing it as the source of other phenomena, rather than by analyzing its origins. Those who recognized the futility of this approach concluded that, to explain the phenomenon, they had to study its causes, not its effects; and then they saw it as consisting of other, simpler phenomena, which are easier to explain. This, in time, led to the mechanistic idea: "When we ask 'why?' concerning an event, we may mean either of two things. We may mean: 'What purpose did this event serve?' or we may mean: 'What earlier circumstances caused this event?' The answer to the former question is a teleological explanation, or an explanation by final causes; the answer to the latter question is a mechanistic explanation. I do not see how it could have been known in advance which of these two questions science ought to ask, or whether it ought to ask both. But experience has shown that the mechanistic question leads to scientific knowledge, while the teleological question does not."²

Once we agree that the only way to understand a phenomenon is by studying its causes, we must decide where to stop. We may well succeed in explaining a particular phenomenon in terms of other, simpler phenomena, and then explaining *those* in terms of even simpler ones, and so on. Which phenomena, though, are the ultimate explanations? That is, at which phenomena can we stop, being certain that we understand them?

The answer, of course, is *never*. But it is the fundamental assumption of the mechanistic doctrine that we all understand, intuitively, simple *mechanical*

² Bertrand Russell, *A History of Western Philosophy* (New York: Simon and Schuster, 1972), p. 67.

operations. So, if we manage to reduce a given phenomenon to mechanical phenomena, we can declare it to be “understood.”

For many phenomena – light, heat, chemical reactions, human feelings – no obvious reduction to mechanical phenomena exists. Nevertheless, mechanism assumes that these phenomena, too, are the result of simpler ones, displayed by some minute particles; and that it is the mechanical properties of those particles that cause, ultimately, the observable phenomena.

We are comfortable with mechanical phenomena because we deal with objects in our everyday activities. We tend, therefore, to perceive everything in the universe as similar to these objects, except for being perhaps much larger, or much smaller, or much faster. We feel that we really understand something only if we can apprehend it with our senses, if we can see or touch it. And, although we admit that certain things cannot be seen or touched as we can common objects, we visualize them nevertheless as similar in nature. Anything else simply doesn't count as understanding.

Thus, the physicists of the seventeenth century attempted to understand complex natural phenomena by assuming that they are ultimately caused by things which, although minute and invisible, resemble familiar devices and processes: “They were always looking for hidden mechanisms, and in so doing supposed, without being concerned about this assumption, that these would be essentially of the same kind as the simple instruments which men had used from time immemorial to relieve their work.”³ Even in the late nineteenth century, the great scientist Lord Kelvin candidly admitted that he understood something only if he could make a mechanical model of it.

In conclusion, mechanism is not an exact, unquestionable concept, but only a *convention*. It is a method that defines “understanding” as the acceptance of some facts with which we are naturally comfortable, and gives us hope that we can explain everything else in terms of these facts.



A belief in mechanism entails a belief in *determinism*: mechanistic theories claim that the future state of a system can be predicted from its current state with any degree of accuracy we want. The greatest success of mechanism, Newton's theory of gravitation, exemplifies this belief: if we know the state of the bodies in the solar system at a particular instant, we can indeed determine their position at any other instant, in the past or in the future. From early successes like this, scientists and philosophers gained the confidence that all

³ E. J. Dijksterhuis, *The Mechanization of the World Picture* (New York: Oxford University Press, 1969), p. 497.

phenomena yet to be explained – not just in physics or astronomy, but also in such fields as psychology, linguistics, and politics – would prove to be, in the end, equally mechanistic.

The most famous expression of this belief is the statement made at the beginning of the nineteenth century by the great mathematician Laplace. The whole universe, he claimed, is a deterministic system consisting of particles of matter acting upon one another according to the law of gravitation; thus, a being who possessed sufficient intelligence to note the state of these particles at a given instant and to solve the relevant equations, could accurately predict every future occurrence in the universe – every entity and every event, every fact and every bit of knowledge. This enthusiasm was shared by most scientists: “When nineteenth-century physicists subscribed to determinism as an article of scientific faith, most of them took for their ideal of a deterministic theory one that defines the state of a physical system in the manner of particle mechanics.”⁴ Strict determinism, however, has been shown since then to be both philosophically and scientifically naive.



Let us pause for a moment and recall the purpose of this discussion. We must study the history of mechanism if we are to understand our mechanistic delusions; and we must understand these delusions in order to recognize our *software* delusions. For, as we will see later, all programming principles, theories, and methodologies invented since the 1970s, despite their great number and variety, can be described with one phrase: *attempts to reduce software to mechanics*.

And it is only because of our long mechanistic tradition that the software charlatans can deceive us with the notion that software phenomena can be reduced mechanistically to simpler ones. While this is true for some isolated problems, which can be represented mathematically, most phenomena related to software and programming are non-mechanistic.

New programming ideas are being introduced every year, all promising a dramatic improvement in programming practices. But once we understand their mechanistic roots, we can easily recognize their similarity. Thus, when analyzed, these ideas reveal a common assumption: software applications consist of neat structures of modules and operations, just as appliances consist of neat structures of subassemblies and parts; so, instead of acquiring programming expertise, we can create applications simply by imitating the methods used to build appliances in a factory. Applications are said to be made of

⁴ Nagel, *Structure of Science*, p. 282.

“components,” which must be assembled using the methods of “software engineering.” Programmers are called “engineers,” and are encouraged to use terminology borrowed from the field of manufacturing. They no longer write, but “build” or “construct” applications. They work in “software factories” and design “data warehouses.”

This mechanistic software ideology is taught in universities, is promoted in books and periodicals, and is supported by businesses and governments with vast amounts of money. But when we search for a logical basis to the belief that the methods of manufacturing can be applied to programming, we find none. All we find is the classic idea that any process can be reduced to simpler ones, and that this reduction can be repeated again and again, down to some processes that are simple enough to implement directly – an idea that was never shown to be true for software.

In reality, then, our software-related pursuits – the so-called information technology revolution, the activities we identify with progress and the future – are steeped in a naive, seventeenth-century mechanistic mentality.

Reductionism and Atomism

1

The earliest documented mechanistic theory, the *atomistic* philosophy, was developed in Greece during the fifth century BC by a school of philosophers we have come to call atomists. This philosophy was founded by Leucippus and Democritus, but was based on the work of earlier Greek philosophers, Empedocles, Parmenides, and Anaxagoras. These thinkers sought to understand the nature of matter and space, and asked questions such as these: What are things ultimately made of? Is space continuous, and thus infinitely divisible into smaller and smaller parts? Or is it discrete, and thus a line of a certain length is composed of a finite number of segments, and segments within segments? They noticed that substances can change into other substances – through heating or mixing, for example; also, animals that eat grass seem to convert the grass into flesh and bones. This means that each substance may have bits of other substances in it, ready to be released through certain processes. Or, more likely, all substances consist of some common elementary entities, and it is the way these entities are combined that determines the observable qualities of each substance.

The atomists concluded that everything is made up of small particles, or atoms, which are infinitely hard and hence indivisible and indestructible. The atoms are surrounded by empty space and are constantly in motion, colliding

and mingling with one another. They come in many shapes and sizes, and the combination of these properties and their motion produces the large-scale phenomena that we observe. The atomists imagined the behaviour of these particles by comparing their interaction to the movement and collision of small objects in a container, or insects in the air, or pebbles in a stream, or grains in a sieve. Thus, while accepting the existence of things that are too small to detect with our senses, the atomists had to visualize them, nonetheless, as similar in nature to the objects encountered in everyday life.

Ancient atomism was a truly mechanistic theory, in that it attempted to explain all phenomena, not just the composition of matter, through reduction to mechanical properties. Sensations and feelings, for example – sight, smell, pain, pleasure – also had to be explained as the result of the collision and interaction of atoms. There is nothing in the universe but atoms and the void, claimed the atomists. Everything is based, ultimately, on matter and motion.

The atomistic philosophy was later embraced by Epicurus, and Epicureanism remained popular in the Roman Empire until the fourth century AD. Then, atomism was forgotten. Centuries later it was remembered, though, owing largely to the book *On the Nature of the Universe*, written in 55 BC by the philosopher and poet Lucretius. Atomism could not be verified, of course, so it was nothing more than speculation. Nevertheless, it greatly influenced the world view of those who accepted it.

After the decline of the ancient civilizations, not much happened in Western science until the Renaissance. Then, during the Scientific Revolution, the ancient atomistic theories were revived and became known as *corpuscular* theories of matter. Their claims had not changed: all phenomena must be explained through the science of mechanics, and hence through reduction to some small and indivisible particles, or corpuscles. Despite the passage of two millennia, scientists could not imagine a world that consisted ultimately of anything but objects and processes which, while invisible, are similar to those we observe in everyday life: “The view became current that all the operations of nature, all the fabric of the created universe, could be reduced to the behaviour of minute particles of matter, and all the variety that presented itself to human experience could be resolved into the question of the size, the configuration, the motion, the position and the juxtaposition of these particles.”¹

It was Pierre Gassendi who introduced ancient atomism in Europe, and the corpuscular notion was quickly embraced by most scientists, including

¹ Herbert Butterfield, *The Origins of Modern Science: 1300–1800*, 2nd ed. (London: Bell and Hyman, 1957), p. 120.

Galileo, Descartes, Boyle, Huygens, Newton, and Leibniz. Innumerable theories were brought forward in an attempt to explain natural phenomena through the motion of particles: solidity and fluidity, rigidity and elasticity, heat and cold, evaporation, condensation, sound, light, colour, taste, etc. These theories explained nothing, of course. They were rather silly conjectures expressing the imagination of the scientists: what they thought particles ought to be like in order to cause those phenomena. There are heat-producing atoms, for instance, and cold-producing atoms, and it is their relative influence that we detect as heat and cold. Solidity of bodies is produced by specially shaped particles that interlock. Melting occurs when heat atoms penetrate the tight spaces between particles. Gravitational attraction is explained as particles emitted from one body and affecting the second body. Vision is explained as the pressure of particles that fill the space between the source of light and the eye.²

Thanks to the vagueness of these conjectures, several theories could easily be formulated by different scientists to explain the same phenomenon. Often, more than one theory was conceived by the same scientist. Thus, although seriously proposed as explanations, the corpuscular theories were mere speculations: “Everything remains in the vaguely qualitative sphere, so that there is no question of an experimental verification of the truth of the theories in question. On the ground of a curious kind of corpuscular imagination, explanatory hypotheses are formulated which may be considered more or less plausible, but which cannot be verified in any way.”³

It is hard to admit today that the greatest scientific minds of that epoch – the epoch we have come to call the century of genius – could engage in such absurd speculations. And it is even harder to admit that the same minds, while engaged in these corpuscular fantasies, produced the brilliant work from which modern science was born. This chapter in the history of science has been told many times, but seldom accurately. Being heirs to the scientific tradition inaugurated by these men, we naturally tend to remember their genius and to forget their mistakes. But we are also heirs to their mechanistic delusions, and if we want to understand our current scientific fallacies we must start by examining theirs. In the following pages, therefore, I want to show the interdependence of scientific thinking and mechanistic delusions, which can help to explain why even today we are liable to confuse one with the other.

² E. J. Dijksterhuis, *The Mechanization of the World Picture* (New York: Oxford University Press, 1969), pp. 416, 427–429.

³ *Ibid.*, p. 430.

2

The early scientists believed that mechanics is the only exact discipline and only the motion of objects has a complete explanation, so we must explain all other phenomena by reducing them to mechanics. This belief gave rise to two types of reductionism, which are still very much with us today. The first one, which may be called *formal* reductionism, is the belief that all phenomena can be reduced ultimately to the simplest mechanical phenomena (the motion of particles of matter), which can then be explained mathematically. The second one, an *informal* reductionism, is the belief that natural phenomena can be explained merely by showing that they are analogous to mechanical devices (since, in principle, these devices can be reduced to simpler and simpler parts, and ultimately to the particles of formal reductionism); clockworks, the most complicated machines in earlier times, were a common model.

The two types of reductionism exist side by side. Scientists prefer *formal* reductions, but, as these theories seldom work, they need some informal analogies as backup. Thus, in the seventeenth century, formal reductionism was represented by the corpuscular theories; but these theories were only speculations, so to support their claims the scientists also described the world informally as a giant machine, with God acting as the master engineer who created it, set it in motion, and is perhaps still supervising its operation.

Today, formal reductionism is represented by theories that attempt to reduce phenomena to simpler ones, in the hope that the latter will be reduced to even simpler ones, and so on, ending eventually with phenomena that can be depicted with mathematical models. But for phenomena that are too complex for mathematical models, we resort to informal models and analogies.

Linguist Noam Chomsky, for example, maintains that it is possible to explain with mathematical precision how a sentence is derived from words, simply by studying its linguistic characteristics. But his theory does not work, so Chomsky is obliged to add an informal model to support it: he postulates an innate language faculty that we possess as part of our genetic structure, so an exact theory of language is no less plausible than an exact theory of the heart or the kidneys; and he defends the idea of an innate language faculty by pointing to the many other innate qualities – our propensity to develop arms rather than wings, for instance.

To take another example, the software theories attempt to reduce programming phenomena – which include in fact many unpredictable business, social, and personal aspects – to precise methods and to formal, mathematical models, on the assumption that software applications can be treated as neat

hierarchical structures of software entities. But these theories do not work, so their proponents support them with informal analogies: they add that these theories must be valid because software development resembles manufacturing and construction activities, and in these activities the concept of hierarchical subassemblies works well.

At first sight, the informal analogy invoked by a reductionist to support his theory may look like additional evidence of its validity. The need for an informal analogy, however, is evidence of its *failure*. If the formal theory worked, the scientist would need no additional, informal evidence to back it up. As it is, the formal theory provides the dignified image that makes the claim appear “scientific,” while the informal analogy is the distraction that masks the failure of the theory and reassures us that there must be something in it after all, so it is worth pursuing. Some of these theories generate “research programs” that continue for years and decades, without anyone realizing that they are in fact worthless, pseudoscientific notions.



A related fallacy today is the idea of *partial* reductionism – the attempt to discover useful mechanistic theories by reducing a given phenomenon to *any* simpler phenomena, not necessarily to the simplest mechanical phenomena. Thus, while all mechanists agree with the *definition* of reductionism and atomism, these principles are so hard to follow that no one actually adheres to them.

Strict reductionism stipulates that the phenomena studied by one discipline are not fully explained unless they are reduced, ultimately, to the simpler phenomena studied by another discipline. Sociology must be reduced to psychology, because sociological phenomena reflect the simpler phenomena involving individuals. Psychology must be reduced to biology, because psychological phenomena reflect the phenomena involving human bodies. Biology must be reduced to chemistry, because bodies are made up of various substances. Chemistry must be reduced to physics, because substances consist of particles like atoms. And physics must be reduced to mechanics – to the motion of those particles.

Just like seventeenth century’s reductionism, which attempted to reduce all phenomena directly to particles of matter, the complete reduction described above is a prerequisite for a valid mechanistic theory. That is, the theory cannot be said to work until all phenomena of lower complexity are reduced to even simpler phenomena. Scientists, however, expect to discover a valid theory by performing only small, partial reductions; for example, they try to explain mental phenomena by reducing them, not to biology, but to simpler mental

phenomena. The reason such theories do not work is that the simplest entities reached are not atomic and independent, so they are not the starting elements in the hierarchical structure required for a mechanistic model. The fallacy, thus, is believing in reductionism and atomism without adhering strictly to these principles. The scientists want the power promised by mechanism, but without the rigours demanded by serious mechanistic theories.

3

Returning to the seventeenth-century scientists, we notice the following, peculiar fact. Even as they claimed strict allegiance to the mechanistic doctrine, even while proposing endless corpuscular theories and clockwork models of the world, their important work – the theories we remember them by – had nothing to do with their mechanical speculations. On the contrary, they were based on deductive reasoning, empirical verification, and solid mathematical principles. They modernized algebra, invented calculus and analytic geometry, and founded the theory of probability; they made discoveries in optics, hydrostatics, pneumatics, physiology, chemistry, and astronomy; all this and more they accomplished within a century. This was the century in which all scientists claimed that mechanics is the only pure science, so everything else must be explained in terms of mechanics; and yet, none of their important theories during this period had to be confirmed through reduction to mechanics.

The greatest irony is that Newton's theory of gravitation, the culmination of seventeenth-century science, was the least mechanistic theory of all. It introduced a mysterious force called gravitational attraction, which horrified everyone, because it was contrary to the accepted mechanistic doctrine. Mechanism asserts that motion can be transmitted from one body to another only by intuitively understandable means; that is, through direct contact and by pushing, like colliding billiard balls. Gravitational force, on the contrary, acts at a distance and by pulling. Newton's contemporaries accused him of relapsing into the medieval belief in occult qualities, of betraying science by abandoning the mechanistic conception of nature. Newton, a mechanist himself, defended gravitational force on the ground that it could be confirmed empirically, and that it was not a primary quality but the result of some hidden mechanical phenomena yet to be discovered. He could not explain it, but he was convinced that one day a mechanistic explanation would be found for it too. He even suggested such a theory himself (a corpuscular theory involving the ether – the substance believed to fill the void between the celestial bodies).

There is an obvious contradiction here. On the one hand, these great

scientists sincerely believed that all phenomena must be explained through formal reductions to mechanical operations, or at least through informal analogies. On the other hand, they did not hesitate to advance theories that stand on their own, based on observation, logic, and mathematics. But this contradiction is resolved once we understand that the corpuscular theories and the informal analogies served in reality only as *metaphors*.

The reductionistic and atomistic concepts inherent in these metaphors gave scientists the confidence to search for new and better explanations. Believing in reductionism and atomism is, in effect, believing that every problem has a solution, that every phenomenon can be explained. These concepts assure us that all things can be divided into smaller things, reaching eventually parts small enough to be considered elementary. At that point, we will know all that can be known: "Physical atomism is more than logical analysis. It is the assumption that there is a quantitative limit to division, that small ultimate units exist... Atomism has rightly been described as a policy for research... But atomism is not merely a policy or a method which has proven brilliantly effective at certain levels of material analysis; it is also a positive assumption regarding ultimate structure. *This assumption has a powerful psychological appeal, for it suggests a limited task with high rewards.* If there really exist ultimate units, we have only to discover their laws and all their possible combinations, and we shall be all-knowing and all-powerful, like gods. So it seems."⁴

Had the seventeenth-century scientists suspected how complex the world really is, they might not have dared look for explanations. Had they known that three centuries later we would still lack a complete understanding of the universe, of matter, of nature, of life, of intelligence, they would have felt unqualified to propose those sweeping theories. But the belief in mechanism inspired them with confidence: assured that the world is well-ordered and fairly simple, they concluded that all natural phenomena can be explained with a few principles – the same principles that govern the behaviour of common objects and devices.

This belief is evident in their corpuscular fantasies. For instance, they postulate the existence of smooth particles and oddly shaped particles to describe phenomena like fluidity and solidity, simply because these phenomena resemble the effects produced by piles of smooth objects and oddly shaped objects;⁵ and they describe magnetism as a vortex of particles shaped like screws, matching appropriately threaded particles present in iron (and magnetic polarity as the existence of particles shaped like right-handed and

⁴ Lancelot L. Whyte, *Essay on Atomism: From Democritus to 1960* (Middletown, CT: Wesleyan University Press, 1961), pp. 14–15 (italics added).

⁵ Dijksterhuis, *Mechanization*, p. 428.

left-handed screws).⁶ Westfall describes this “imaginary construction of invisible mechanisms to account for phenomena” as “the occupational vice of mechanical philosophers.”⁷

Nor did this naivety end when the scientists discovered the *real* theories – which they did, not through their corpuscular fantasies, but through careful observation and experimentation. Although the real theories did not depend on corpuscular ideas, they went back and modified the corpuscular theories to agree with the facts revealed by the real theories.

The mechanistic belief was like the belief in a religious myth: it motivated the scientists to great feats of intellect, and they could not abandon it even when recognizing its fallaciousness. They could continue to accept the corpuscular theories and the machine analogies precisely because they were so vague, and because they had no bearing on the real theories anyway.

It is, again, Newton’s persistent belief in a mechanistic explanation for gravitational attraction even after proving his great theory – which, being based on the concept of force, *contradicted* that belief – that is paradigmatic of the mechanistic obsession. Then, still failing to find a mechanistic explanation but unwilling to abandon mechanism, his followers in the eighteenth century quietly modified the meaning of mechanism to *include* the new concept of force. Eventually, “Newton’s concept of force, which had been rejected as essentially unmechanistic [earlier], came to be regarded as precisely the most characteristic feature of a mechanistic conception of nature.”⁸

It was their naivety as much as their acumen that led the early scientists to recognize mechanism – the combination of reductionism and atomism – as an effective method of discovery. It was their belief that all the mysteries of the world are ultimately accessible to the human mind, that they are merely a collection of puzzles waiting to be solved. With this belief comes the confidence that the phenomena we observe on a large scale, and which we do not understand, must be caused by some phenomena that exist on a smaller scale. So it is those phenomena that we must study. Then, even if we are wrong and those phenomena are not the elementary causes we thought they were, we only need to treat *them* as the large scale phenomena, and repeat the process. Eventually, we are bound to reach the lowest level: the elementary entities and processes that cause everything else.

Consciously or not, therefore, anyone who believes in determinism and in the existence of complete explanations will become a believer in reductionism and atomism. Indeed, if we define *understanding* as a process similar to knowing how a complicated machine works, there is no other way to under-

⁶ Richard S. Westfall, *The Construction of Modern Science: Mechanisms and Mechanics* (New York: Cambridge University Press, 1977), pp. 36–37.

⁷ *Ibid.*, p. 41.

⁸ Dijksterhuis, *Mechanization*, p. 497.

stand than by perceiving all phenomena as systems of things within things, waiting for us – just like the subassemblies of a machine – to take them apart and study, one level at a time.

But this method, while useful in disciplines like physics and chemistry, is only partially effective in disciplines like biology and physiology, and is totally useless in those disciplines that deal with mental capabilities or human relations. This is true because reductionism and atomism are valid only for *independent* phenomena; that is, when the lower-level phenomena can be isolated and studied separately. When the lower-level phenomena share some of their elements, they interact, and isolating them alters the nature of the high-level phenomenon by severing the interactions.

The early scientists were unaware of this limitation, and were convinced that *all* phenomena would eventually be explained mechanistically. It is this naivety that must be stressed – not to denigrate their work, but to recognize the same mistake in our own beliefs. It would be historically inaccurate to see the adoption of the mechanistic philosophy as the result of a careful consideration of research methods. It was the result of a naive view of the world: the belief that the world is simpler than it really is.

When scientists today attempt to describe with neat models the human mind, or human communication and relations, or economic phenomena, or the process of software development, they are guilty of the same naivety – the childish belief in a simple and well-ordered world. Although scientists have been attempting for three centuries to reduce everything to mechanics, this has succeeded only for phenomena that can indeed be isolated, and hence approximated with mechanistic theories. All other attempts have resulted in mechanistic delusions.

We are continuing in this tradition, but with one important difference. Whereas the earlier delusions rarely went beyond academic discourse, the current mechanistic theories immediately become fashionable and influential. No matter how worthless, they are accepted with enthusiasm by experts and laymen, by universities, corporations, and governments – while the non-mechanistic theories, even when successful, are dismissed as “unscientific.” We have redefined science to mean simply the pursuit of mechanism, and we no longer care whether a theory works or not, or whether an idea is useful or not, as long as it is mechanistic. As a result, we are wasting more and more of our resources on mechanistic delusions. Moreover, because the pursuit of a mechanistic delusion is just like the pursuit of a pseudoscientific theory, our society is increasingly dominated by crackpots and charlatans. Finally, with the mechanistic *software* theories and their repercussions – the destruction of knowledge and minds, the promotion of totalitarianism – our mechanistic delusions have reached the point where they are threatening civilization itself.



To summarize, the function of mechanism is mainly as myth. In both their informal theories (the analogies of phenomena with machines) and their formal theories (the reduction to particles), the early scientists never insisted on a perfect mechanistic representation, and were ready to overlook deficiencies and contradictions. All that mattered was to discover a theory that worked.

What made mechanism such a powerful myth were the twin concepts of reductionism and atomism. These concepts gave the scientists both the confidence to advance revolutionary theories and the methods to verify them: “The dream of final atomic knowledge seized many Western scientific minds, consciously or unconsciously guiding them to unexpected discoveries which in a sense justified the dream.”⁹ “In our not fully successful attempts at reduction, especially of chemistry to physics, we have learned an incredible amount. . . . Thus from the point of view of method, our reduction programmes have led to great successes, even though it may be said that the attempted reductions have, as such, usually failed.”¹⁰

Final atomic knowledge was never attained, but the *belief* that such knowledge is possible has helped us reach levels of knowledge that otherwise we might not even have envisaged. If we doubt the value of mechanism, we need only compare the spectacular advances in natural science during the century of the Scientific Revolution with the stagnation during the preceding one thousand years. Mechanism acted merely as a motivating force, as a psychological aid. But the fact that it is mainly a myth does not lessen its value in those fields where reductionism and atomism lead to useful theories. The danger is to conclude that mechanism is a universal method of science, and to apply it with the same confidence in fields where it is worthless. Its powerful psychological appeal works then *against* science, by preventing us from trying other, possibly better, research methods.

4

Few of us realize how much of our academic research, and how many of the ideas that guide our society, are but mechanistic delusions. The reason we accept absurd theories time and again, and thus repeat the mistakes of the past, is that we don't know how to identify a mechanistic theory.

⁹ Whyte, *Essay on Atomism*, p. 15.

¹⁰ Karl R. Popper, *The Open Universe: An Argument for Indeterminism* (London: Routledge, 1988), p. 146.

Psychologists, linguists, sociologists, and software theorists may no longer attempt to reduce their problems to particles of matter or to clockwork models, and they may even criticize the mechanistic doctrine. When we study their theories, though, we find nothing but mechanistic beliefs. One mechanistic theory after another are being invented, and hailed as scientific breakthroughs. We clearly see that they do not work, but we remain convinced that they are valid. And we do not recognize their similarity to the *past* theories, which also did not work, and which no one is taking seriously any longer. It is our failure to recognize their common mechanistic grounding that tempts us to accept each one as novel and important.

Our task now is to find that common concept – the concept shared by all mechanistic theories. We have already determined that the corpuscular theories of the seventeenth century, the atomistic theories of antiquity or of modern physics, and the analogies to machines and manufacturing, are *metaphors*. Reductionism and atomism provided the actual method, while the metaphors provided only psychological support. The metaphors may change, but the underlying methods do not. So, if we want to find that common concept, we must search for the simplest model that provides the functions of reductionism and atomism. It is not hard to see what that model is: the hierarchical structure (see figure 1-1). Clearly, this structure embodies both functions: if we move from the top element toward the terminal elements – that is, from the highest level of complexity to the lowest – the elements at each level are reduced to the simpler elements of the lower level; and the terminal elements are the structure’s atoms – its basic building blocks, or alphabet.




Figure 1-1

We can use the hierarchical structure to model anything that can be depicted as a system of things within things. Thus, to represent the atomistic and corpuscular theories, the terminal elements of the hierarchy are the particles, and the top element is the object or phenomenon described by the theory. When the mechanistic metaphor is a machine, the terminal elements

are the simplest operations, the top element is the whole machine, and the intermediate elements are the various subprocesses.¹¹

So the hierarchical structure can represent any mechanistic theory and any mechanistic phenomenon. This is true because the elements and relations can stand for anything – physical entities, processes, bits of knowledge, pieces of software, and so forth. The hierarchical concept is what matters, what provides the mechanistic qualities: describing the elements at one level of complexity as combinations of elements from the lower level provides the reductionistic quality, and basing the entire structure on a relatively small set of terminal elements provides the atomistic quality.

Thus, the hierarchical concept can be equated with mechanism because it is the simplest concept that provides the two fundamental properties of the mechanistic method, reductionism and atomism. Let us call *immediate* metaphors the traditional mechanistic metaphors of particles and machines; the hierarchical concept is then the *hidden* metaphor. The hierarchical metaphor is inherent in every mechanistic metaphor; that is, we can always reduce an immediate metaphor to a hierarchical model. The immediate metaphors are expendable, but the hidden, hierarchical metaphor is always present. So, if we want to determine whether a theory is mechanistic, all we have to do is verify whether it can be represented with a hierarchical model. We must not be distracted by the immediate metaphors that accompany the theory; if based ultimately on the hierarchical concept, it is mechanistic.

In the case of software theories, the most popular immediate metaphors are from engineering, but we must not take them too seriously. It is the hidden metaphor that we must look for. And indeed, whether or not employing engineering metaphors, the software theorists resort in the end to hierarchical structures. In fact, they misunderstand and misrepresent the engineering practices altogether. They claim that software development must resemble engineering projects, but even for an informal metaphor their notion of engineering is highly distorted. If we remember the hidden hierarchical metaphor, however, we can easily understand the source of their engineering delusions. What the software theorists like in engineering – what they wish to emulate in software development – is the high degree of success, the precision, and the predictability. They superficially study the engineering practices, and, lacking engineering knowledge, all they see is the use of hierarchical structures: the design of complicated objects as modules within modules. They conclude that this is all there is to engineering, failing to realize that the hierarchical concept was their own metaphor, their own delusion.

¹¹ We will study hierarchical structures, which I also call simple structures, in the next section. They are introduced here in order to complete the present argument.

We like the hierarchical concept because it is reassuring: its reductionism and atomism create the illusion that it can explain any phenomenon and solve any problem. Guided by this belief, we see hierarchical structures everywhere we look. These hierarchies, as well as their ability to create large structures from simple elements, are real enough; but it is a mistake to conclude that this is the most complex concept possible. What creates richness and complexity in the world is not a large number of levels or elements, nor a large number of hierarchies, but the *interaction* of hierarchies.

There are indeed hierarchies in engineering projects, as there are in software projects. But on their own, without the expertise and creativity contributed by human minds, they would remain simple, mechanistic concepts. The software theorists see the hierarchies of engineering, but they underrate the knowledge added by individuals, and which cannot be reduced to hierarchical structures. Thus, they fail to see the *non-mechanistic* aspects of the engineering practices. In programming, the non-mechanistic aspects are even more important, and there is very little that can be accomplished with independent hierarchical structures.

5

To understand the influence of mechanism on science we must also consider the role of mathematics. Science became practically synonymous with mathematics in the seventeenth century, at precisely the time when it became mechanistic. This is hardly a coincidence.

Mathematical systems are hierarchical structures: in a given system, we start with a few basic concepts and create more complex ones by combining, one level at a time, concepts from the previous level; or, conversely, we analyze a complex concept by breaking it down into simpler ones, one level at a time, until we reach the basic concepts. Thus, mathematics and mechanism employ the same hidden metaphor. The concepts of a mathematical system are the entities that make it up, and the functions and theorems that use these entities.

The first mathematical system designed as a hierarchical structure was Euclid's geometry. Euclid, who worked in Alexandria from about 300 BC, collected all the geometrical knowledge accumulated by his predecessors and organized it as one book, the *Elements*. As is well known, he started with a small number of elementary assertions – assertions that appear to state self-evident truths, and can therefore act as premises. Then, by logically combining them, he showed how to demonstrate the truth of more and more complex assertions. The premises, the theorems built from them, those built from the former ones, and so on, form a perfect hierarchical structure, and

the progression from one level to the next is based on logical deduction. This guarantees the validity of the theorems at the higher levels, which are too complex to prove directly. And, even though Euclid dealt mainly with geometry, this principle became the foundation of all of mathematics. The simplicity of the idea, coupled with the seemingly unlimited complexity of the entities, functions, and theorems that can be created, explains perhaps our infatuation with hierarchical structures.

Even in ancient times mathematics was appreciated for its ability to solve practical problems. But it was only at the time of the Scientific Revolution, when mechanism became a universal research method, that its ability to represent natural phenomena was recognized. Since mathematics and mechanism are both based on the hierarchical metaphor, it was at this time that science, mathematics, and mechanism became practically indistinguishable. Kepler and Galileo were convinced that “the structure of the external world was essentially mathematical in character and a natural harmony existed between the universe and the mathematical thought of the human mind.”¹² They likened the universe to an open book, but a book written in the language of mathematics: we can read it, so we can discover all the secrets of nature, but only through mathematics.

Descartes went further and *identified* science and knowledge with mathematics. He was impressed by “the long chains of simple and easy reasonings by means of which geometers are accustomed to reach the conclusions of their most difficult demonstrations,”¹³ and decided to adopt the hierarchical method for all fields of knowledge. Moreover, he says, he “had little difficulty in determining the objects with which it was necessary to commence, for [he] was already persuaded that it must be with the simplest and easiest to know.”¹⁴ He believed, in other words, that all we have to do is represent everything as entities within entities, and ensure that the starting elements of these structures are simple enough to understand directly. Then, if we rigorously follow the hierarchical links, if we “always preserve in our thoughts the order necessary for the deduction of one truth from another,”¹⁵ we will be able to create all the knowledge that human minds can attain.

It matters little to us whether the popularity of mechanism was enhanced by the successful use of mathematics, or whether, conversely, scientists adopted mathematics as tool because it was based on the same metaphor as mechanism. What is significant is the transition to the modern world view that took place at this time and involved both mechanism and mathematics: “The

¹² Dijksterhuis, *Mechanization*, p. 404.

¹³ René Descartes, *A Discourse on Method* (London: Dent, 1912), p. 16.

¹⁴ *Ibid.*

¹⁵ *Ibid.*

mechanization of the world-picture during the transition from ancient to classical science meant the introduction of a description of nature with the aid of the mathematical concepts of classical mechanics; it marks the beginning of the mathematization of science, which continues at an ever-increasing pace in the twentieth century.”¹⁶ All scientists in the seventeenth century shared in the new belief that “nature has to be described in mathematical language and that it can only be understood by man to the extent that he can describe its workings in that language.”¹⁷

Since mathematics is based on the hierarchical concept, the mathematical belief expressed in the last sentence – a belief that continues to guide us even today – can be summarized as follows: to understand a phenomenon means to be able to represent it with a hierarchical structure. But this, as we saw earlier, is also the idea of mechanism. The mathematical belief, therefore, is identical to the mechanistic belief. Mathematical models are mechanistic models.



If mathematics is identical to mechanism, it too can represent accurately only *deterministic* phenomena; it too is useless, therefore, for phenomena involving minds and societies, which are indeterministic. All mechanistic delusions in the human sciences, we will see later, stem from the belief that complex human phenomena can be represented with exact, mathematical models.

Similarly, despite many exact aspects, software-related phenomena involve minds and societies; so they are, in the end, indeterministic. The idea of defining and developing software applications with the formal tools of mathematics is, therefore, absurd. Accordingly, all programming theories based on this idea – the relational database model, structured programming, and the like – are mechanistic delusions.

The theorists look at software and see that, just like Euclid’s geometry, applications are ultimately made up of some basic elements (the individual operations), which are simple enough to verify directly. Why can’t we, then, simply combine these elements hierarchically into more and more complex ones (blocks of operations, modules, and so on), each one guaranteed to work perfectly (because built from previously proven elements), until we reach the complete application? Why can’t we, in other words, apply Descartes’s method, said to work for all human knowledge, to the knowledge embodied in a software application? Programs are nothing but parts within parts, so we should be able to build applications of any size and complexity with geometrical precision, simply by designing them as strict hierarchical structures.

¹⁶ Dijksterhuis, *Mechanization*, p. 501.

¹⁷ *Ibid.*

This is what, in one form or another, is promised by all theories and methodologies. But, as we will see, this promise cannot be fulfilled. The hierarchical concept is useful in geometry because the hierarchical structures represented by geometrical theorems are a *direct mapping* of the world of lines, angles, and objects that geometry is concerned with; so a geometrical structure provides a one-to-one correspondence to the real structure. The phenomena we wish to represent with software, on the other hand, are systems of interacting structures. We may perhaps succeed in mapping *each aspect* of a given phenomenon into a software structure, but we cannot develop the application by implementing these structures separately. Since the actual structures interact, the corresponding software structures must interact too, if the application is to represent the phenomenon accurately. To treat this phenomenon as we do geometrical problems, we must first separate it into its constituent structures. But if we ignore the interactions, then even if we successfully program the individual structures, the application will not reflect reality. (We will discuss the concept of interacting structures later in this chapter, and software structures in chapter 4.)

6

Having equated the hierarchical structure with mechanism, we can define it as the model that represents all mechanistic phenomena, and all mechanistic theories. Our next task is to study this model. Then, we will extend it so as to represent *non-mechanistic* phenomena; this will be a system of *interacting* hierarchies.

I call the mechanistic model a *simple structure*, and the non-mechanistic one a *complex structure*. We will make good use of these models later, when we study various phenomena and the theories that attempt to explain them, and especially when we study *software* phenomena. These models will help us to determine whether a given phenomenon can be represented with a simple structure (in which case it can be explained mechanistically), or whether it can only be represented with a complex structure (in which case no mechanistic theory can explain it).

These models will also help us to understand why it is so easy to fall prey to mechanistic delusions. Because a complex structure appears to be just a collection of simple structures, it is tempting to try to explain the complex phenomenon by studying several simple phenomena in isolation. We will see that, in the final analysis, what all mechanistic theories do is attempt to represent a certain phenomenon with simple structures. This idea works when the phenomenon is reducible to simpler phenomena, but fails when it is not.

Mechanistic theories provide complete and precise explanations. Those who believe in mechanism claim that *all* phenomena can be explained, so their particular phenomenon *must* have a mechanistic theory. If the current theory does not work, they say, it will be improved, or a better one will be discovered in the future. With our models, we can see that this optimism is unwarranted when the phenomenon can only be represented with a complex structure, because no theory can provide a complete and precise explanation for complex phenomena.

I want to emphasize again that it is not mechanism in itself that is the target of this criticism, but the mechanistic *dogma*. Mechanism is an important concept, and we must always start by determining whether a given phenomenon can be usefully represented with a mechanistic model. When successful, such models are invaluable. Our concern here is with mechanistic *delusions*, which stem from the belief that *every* phenomenon can be represented with a mechanistic model.

History abounds with mechanistic delusions, especially following the Scientific Revolution. In our own time, however, they have reached epidemic proportions. One reason may be that we have pushed our knowledge beyond the capability of deterministic models. As we will see presently, even the *successful* mechanistic theories are only *approximations* of reality, because nothing in the world can be purely mechanistic. So, as we expand the range of phenomena that we wish to explain, only complex structures, which are indeterministic, can act as models. But we have yet to reach the level of scientific maturity where we can admit that some phenomena lie beyond the explanatory power of mechanistic principles. So we continue to invent mechanistic theories, whose determinism is comforting, even as we note that they fail to explain these phenomena.

Philosopher Karl Popper reminds us that our sciences are founded on conventions: the simplicity, determinism, and universality we seek to attain with our scientific theories are criteria we have invented ourselves, because this is the only way that we, human beings, can practise science. But theories are only *models* of the world, so no matter how successful some of them are, it doesn't mean that the world itself is simple, deterministic, and regular. We prefer simple theories, but "the world, as we know it, is highly complex; and although it may possess structural aspects which are simple in some sense or other, the simplicity of some of our theories – which is of our own making – does not entail the intrinsic simplicity of the world."¹⁸

Similarly, we prefer deterministic theories, because they are relatively easy to test; but "it seems no more justifiable to infer from their success that the

¹⁸ Popper, *Open Universe*, p. 43.

world has an intrinsically deterministic character than to infer that the world is intrinsically simple.... The method of science depends upon our attempts to describe the world with simple theories: theories that are complex may become untestable, even if they happen to be true.... We have much reason to believe that the world is unique: a unique and highly complex – perhaps even infinitely complex – combination of occurrences of interacting processes. Yet we try to describe this unique world with the help of *universal* theories. Do these theories describe universal features of the world, regularities? Or is universality, like simplicity, a characteristic merely of our theories – perhaps of our theoretical language – but not of the world?”¹⁹

The conclusion we must draw is this: We defined science long ago as a body of narrow, mechanistic principles. But if we are ambitious enough to attempt to understand those aspects of the world that lie beyond the explanatory power of mechanistic principles – aspects like our mental capabilities, our social affairs, or our software pursuits – then we must also be wise enough to supplement these principles with new ones, adequate for complex phenomena. For, if we continue to believe that complex phenomena can be explained mechanistically, we will be exploited forever by the cranks and charlatans who, under cover of science, promise to solve our complex problems with mechanistic methods.

Simple Structures

1

The model I call *simple structure* is the hierarchical structure introduced in the previous section (see figure 1-1, p. 83). Thus, I also call this model a simple hierarchical structure.

A hierarchical structure is a system of things within things. Hierarchies are usually represented as inverted trees, with the branches spreading downward; but they can also be drawn with the branches spreading upward or sideways. Hierarchical structures have *levels*: two or more *elements* at one level are combined to form an element at the next higher level. These combinations signify the *relations* between elements: the value of an element at a particular level is a function of the values of the elements at the lower level and the particular *operation* that combines them. At the highest level there is only one element, whereas at the lowest level there are usually many elements. Different branches may have a different number of levels, so the *terminal*

¹⁹ Ibid., pp. 44–45.

elements (the elements at the lowest level) may actually be at different levels relative to one another.

When studying the hierarchy by moving from low to high levels, it is more appropriate to call the terminal elements *starting* elements. And, since it is common to view hierarchical structures both ways, the two terms, “terminal” and “starting,” are used interchangeably for these elements.

The *definition* of a hierarchical structure must include its starting elements, its operations, and the various rules for using them; that is, everything we need in order to derive all valid alternatives for that hierarchy – all possible ways to combine elements, and all the values that the top element can display.

The designation of levels as low and high, although matching the inverted tree representation, has a different reason. Even if we reverse the tree and show the terminal elements at the top, these are still the lowest levels; and the element at the bottom is still called the *top* element. The reason for designating the terminal elements as the lowest level is the way hierarchical structures are used: the terminal elements are the simplest entities of the structure, and the elements at any given level are more complex than those at the next lower level (because one element is the result of an operation between several lower-level elements). The hierarchical levels indicate, therefore, *levels of complexity*.¹

The hierarchical structure is an extremely versatile model. The elements in the structure can stand for almost anything – physical objects, persons, processes, events, situations, categories, ideas, linguistic entities, pieces of software – and the connections between them can represent any kind of operations. The structure is completely general, therefore, and can model any system involving neatly related “things within things.” There is no limit to the number of levels, elements, or operations, so the structure can be extended indefinitely. Let us briefly examine some common uses of the hierarchical model.

Most classification systems are hierarchical. If we imagine a structure starting with the element *life* at the top, a simple classification is as follows: the top element branches into the elements *animals* and *plants*, one level down; the element *animals* branches into *domestic* and *wild*; *domestic* includes the elements *dogs*, *horses*, *chickens*, and so on; *dogs* includes various *breeds*, and each breed finally branches into the terminal elements – the individual animals we call dogs. In a classification, the operations are the criteria whereby the elements at one level are combined to form the next higher level. The only

¹ Do not confuse this complexity with the complexity of complex structures. Although the same word is used, the two types of complexity are so different (as will become evident later) that it is always obvious from the context which one is meant. In simple structures, the complexity is caused by the shift from low to high levels within one structure; in complex structures, it is caused by the interaction of several simple structures.

elements that are real things are the terminal elements; the elements at higher levels are abstract concepts that reflect the way we choose to group these things.

We find another example of hierarchies in human organizations like corporations, governments, and armies. The elements in these structures can be the units (departments, divisions, sections) or the people who make up or head these units (managers, commanders, workers); the terminal elements are the smallest units or the actual individuals; and the operations indicate how the elements at one level are grouped to form the elements of the next higher level.

Natural systems like organisms can be seen as hierarchical structures. The human body consists of subsystems (nervous, digestive, respiratory, etc.), which in their turn consist of various organs (heart, kidneys, etc.); organs are made up of specialized parts, and these parts are made up of cells, the terminal elements.

Similarly, artificial systems like cars and appliances are designed as hierarchical structures. A car consists of a number of subassemblies, which are composed of simpler subassemblies, and so on, down to the thousands of individual components that are the terminal elements.

We will encounter many other hierarchies later. The examples we have examined, however, already allow us to discuss the most important characteristics of the hierarchical structure. We note, first, that hierarchies are easy to understand when their elements are physical objects. Since physical structures are common and obvious, we have no difficulty visualizing their elements, levels, and operations. But, while simple structures make indeed excellent models for physical structures, this is their least interesting application. We are going to use simple structures to model such phenomena as knowledge, language, and software. Unlike physical objects, the elements and levels in these structures may not be obvious, and the relations may not resemble assembly operations, so the hierarchy may be more difficult to visualize.

We note, next, that the hierarchical model is useful for both analysis and synthesis. When used for analysis, we move from high to low levels: given a complex problem, for instance, we divide it into simpler problems, which become the elements at the next lower level; and we repeat this process of division until we reach problems that are simple enough to solve directly. When used for synthesis, we move from low to high levels: we start with simple concepts, for instance, and combine them into more and more complex ones. A hierarchical tree diagram can be viewed, therefore, as a process that moves in both directions. The elements at the lowest level remain the simplest in the structure; but they can be either terminal elements or starting elements. We can study a physical system, for example, either starting from the completed structure and moving toward the individual parts (in a disassembly operation),

or starting from the parts and moving toward the higher levels (in an assembly operation).

Some confusion often arises between the *definition* of a hierarchy and the tree diagram used to *represent* it. In the tree diagram, many elements at the lowest level combine into fewer and fewer elements, until we reach just one element at the top. The concept of a hierarchy, however, appears to imply the opposite: we start with just a few values for the elements at the lowest level (a limited “alphabet”), and by combining these elements, more and more values are possible for the elements at the higher levels. But there is no contradiction here. A tree diagram depicts *one instance* of the hierarchy, not its definition; it shows one particular set of values for the starting elements, and the combinations leading to the corresponding value of the top element. To *define* the hierarchy we usually employ other methods – rules, descriptions, formulas, etc. It is impractical to depict the definition itself with tree diagrams, but we can visualize it like this: we would start by drawing all possible trees (perhaps an infinite number of them); then, if we looked at *all* these trees, we would see *the same few* starting values repeated in all the trees at the lowest level, a greater variety of values at the intermediate levels, and the greatest variety of values for the top element (as many values perhaps as there are trees).

2

Let us discuss next the concept of abstraction. To abstract means to leave something out, to extract one aspect out of a whole. And hierarchical structures, clearly, function as systems of abstraction. When we combine the elements of one level through a particular operation to form the next higher level, what we do is abstract, from all the attributes possessed by these elements, those attributes that are important in the relation between the two levels. Thus, while hierarchical operations can take many forms, it is ultimately the attributes of elements, and the abstraction of attributes, that determine how one level gives rise to the next.

This is easy to understand in a classification. When we recognize a great number of dogs as members of a particular breed, what we do is extract, from all the attributes that distinguish each dog, those attributes that are relevant in identifying the breed. So, we may consider the dog’s colour and the shape of its ears, but ignore its internal organs and its age. Similarly, when we move to the next level (dogs), we abstract the attributes that distinguish dogs, regardless of their breed, from other kinds of domestic animals; and on the following level (domestic animals), we abstract the attributes that distinguish domestic from wild animals.

Let us look at another hierarchy, a physical structure this time. The transmission of a car consists of subassemblies within subassemblies, on several levels, down to the smallest parts. Here we abstract from the attributes of an individual part, those attributes relevant to the operation of the subassembly to which it belongs, and ignore the others. So for a gear we may note attributes like dimensions and number of teeth, but ignore its date of manufacture. The resulting subassembly has its own set of attributes, but, when combining it with other subassemblies, we abstract only those attributes relevant to the operation of the subassembly at the *next* level. Finally, when the transmission is installed in a car, we ignore its internal details altogether and abstract only attributes like size, weight, and function.

We saw earlier that the levels in a hierarchy indicate levels of complexity: the higher the level, the higher the complexity of the elements, because the value of those elements is affected by more and more elements from lower levels. Now we see that levels have a second and related significance: as we move to higher levels, we increase the degree of abstraction. For this reason, the hierarchical levels are also called *levels of abstraction*. Through abstraction, we lose at each level some of the details that were important at the lower level. What we gain is the ability to deal with just one element. This is beneficial when we can ignore the lower levels, when only the attributes at a particular level are important. Thus, we can discuss the concept of animals without having to think of individual creatures; and we can make use of the car's transmission while knowing nothing about its internal operation. We couldn't even think of animals in general if we had to recall every single one; and it would be difficult to design or build a car if we had to deal with the smallest parts, without subassemblies.



The concept of abstraction leads us to one of the two fallacies born from the popularity of the hierarchical model. We will discuss this fallacy in detail later, but a brief introduction is in order.

We saw that moving to higher levels increases the complexity of the elements, as each element is the result of many operations and elements from lower levels. This may tempt us to associate higher levels with “power”: just by moving to higher levels we seem to be getting something for nothing. If the top element of the structure is our goal – as in a manufacturing project, or when creating a structure of knowledge in the mind – then, we may think, it is foolish to start from a low level. For, the higher our starting level, the more complex the starting elements, and the faster we will reach the top. This belief leads in practice to the principle that we should always attempt to start with the largest

subassemblies available – whether the subassemblies are physical objects or pieces of knowledge. In programming, the subassemblies are pieces of software, and this principle has engendered an endless series of theories that promise (often using these very words) higher levels of abstraction, as if this were such an obvious benefit that no discussion is needed.

But, while levels impart to hierarchical structures their “power,” it is not always beneficial to start from higher levels. A high-level element does indeed replace a whole combination of elements and operations, but it can only replace *one* combination. Thus, as we move to higher levels and higher complexity, we suffer an *impoverishment* in the values possible for our starting elements. As a result, all levels above them, including the top one, will be impoverished. We do indeed shorten our route to the top element, but we lose many alternatives in its value. In an extreme case, when our starting level is the top element itself, there can be only one alternative.

To start from higher levels, then, we must accept the consequences of abstraction. For example, when we start with the concepts of *white*, *blue*, *red*, and so on, there are many ways to perceive an object; but if we replace them with one concept, *colour*, we must accept any colour, and we may even lose the idea of colour differences altogether. And if we go even further and replace *colour*, *shape*, *size*, and so on, with one higher-level concept, *quality*, we will no longer be able to differentiate objects at all. What we will have done, in effect, is replace an infinity of objects of different sizes, shapes, and colours, with only one object. Such an object is an abstract concept, useless as a starting element in any project, except perhaps for creating even more abstract concepts.

The higher the level we start from, the less detailed are our starting elements. This principle can be useful, as we saw: we design cars in such a way that the worker who installs the transmission doesn’t need to understand how it works. But we should be wary of applying this principle whenever we see a structure of things within things. Computers and software, for instance, are important tools, and specific applications can be of great benefit. If, however, instead of studying each application we are content with the high-level, abstract concept “technology,” we will accept any software novelty, useful or useless.

The consequences of abstraction can be even more serious, of course. Millions of citizens who enthusiastically vote for “freedom,” “democracy,” “rights,” “socialism,” or “equality,” but without understanding clearly the meaning of these abstract concepts, vote in effect for words; so they can easily bring to power a dictator who uses these words but whose intentions are the opposite of what the words stand for. It happened more than once, as we well know. Similarly, millions of software practitioners who enthusiastically embrace concepts like “software engineering,” “object-oriented programming,” or “relational databases,” mistakenly perceiving them as progress, embrace in

effect words, and bring to power evil organizations whose intentions are the opposite of progress: the prevention of programming skills and software freedom, and, ultimately, the destruction of all human knowledge and freedom.

Complex Structures

1

Given the versatility of the simple hierarchical structure, one may wonder: how can there be anything else? It is precisely this versatility that has led to our mechanistic delusions, to the belief that nothing lies beyond the modeling power of the simple structure. Everything around us, it seems, fits into one hierarchy or another, so there is no need for another model.

But all these hierarchies are not, in fact, simple structures. Let us consider an ordinary object. We immediately notice that it has a shape. And, since all objects have shapes, through this attribute it is related to other objects. The simplest way to represent this relation is as a three-level hierarchy: the top element, which represents all objects, branches into two elements – the category of objects with this particular shape, and the category of objects with any other shape; and these two elements branch into the terminal elements – the individual objects. (See figure 1-2. More elaborate hierarchies can be created, with categories within categories, if we consider also the details of the objects' shape; but here we only need the simplest one.)




Figure 1-2

The object also has other attributes – colour, size, texture, weight, etc.; and each one relates it to other objects, since all objects have these attributes. These relations too can be represented as three-level hierarchies, and each hierarchy would likely be different (two objects can have, for instance, the same shape but different colours).

The object has, in fact, many attributes, and hence many relations with other objects. No matter how we study this object, we can find a hierarchy where it belongs quite naturally: through its colour, it is part of a hierarchy that relates objects according to their colour; through its location, it is part of a hierarchy

that relates objects according to their location; and so on. We can find relations based on its uses, or its manufacturer, or the ratio between its weight and its length, or the age of the person who owns it. Clearly, we can continue the list of attributes indefinitely. Nor are these relations and hierarchies a figment of our imagination; they are as real as the objects and attributes themselves. And, while the starting elements in the example were physical objects, we can follow the same logic starting with any other entities. Thus, we can think of any number of attributes relating persons, events, situations, or concepts.¹

Recall the idea behind simple structures: at each level, a certain operation combines several elements to form an element for the next higher level. The critical condition is that these operations be precisely defined, so that, given the value of the lower-level elements, we can determine the value of the resulting higher-level element. This condition ensures the determinism of simple structures, but restricts the kinds of relations possible between a particular structure and the external entities – those entities that do not concern it. Specifically, whatever relations exist between its elements and the rest of the world must not alter the nature of its own, internal relations.

Ideally, then, a simple structure is completely isolated from the rest of the world. But we just saw that this is impossible, that there are no entities which are isolated from all other entities. Since each entity has attributes – and has, in fact, a large number of *significant* attributes – it is necessarily related to all the other entities that have the same attributes. Each attribute causes the entity to be an element in a different hierarchy, so the entity is an element in several hierarchies at the same time. As a result, no structure can be isolated from all the other structures.

If a structure shares some of its elements with other structures, the value of its top element depends also on other factors, besides the lower-level elements and their operations: it depends on the elements and operations of the other structures. Thus, when sharing elements, structures interact. And consequently, the behaviour of each structure will be different from its behaviour when isolated. Its definition – in particular, the operations specifying how elements are combined from one level to the next – is no longer sufficient. To describe its behaviour, we must also take into account the structures it interacts with. But the behaviour of each one of those structures is itself affected by the interactions, so it too deviates from the structure's definition. Clearly, then, if we place no restrictions on the interactions, the only

¹ A note on terminology: In this book, “entity” and “element” are used to refer to the same things, but from different perspectives. Objects, categories, processes, pieces of software, etc., are *entities*, and at the same time they are *elements* of structures. So, while the two terms refer to the same types of things and are often interchangeable, “entity” is generally used for the thing in itself, and “element” to stress its role in a structure.

way to study the behaviour of one structure is by studying the system of structures as a whole.

Let us call *complex structure*, thus, a system of two or more interacting simple structures.

2

The most important difference between simple and complex structures, and probably the one most difficult to understand, is this: a simple structure can represent the relations created by *only one* attribute; we need a complex structure if we want to represent the relations created by *multiple* attributes. Let us examine this problem.

Recall, first, the three-level structure that represents the relations created by each attribute (figure 1-2). In this structure, one element stands for the category of entities for which the attribute has a particular value, and the other for the category with any other value. We are going to use this type of structure to show how entities are related through their attributes, but bear in mind that this is only a simplified representation.

Attributes can have a whole range of values, and even categories of ranges, so a complete structure would require many levels and many elements on each level. If the attribute is colour, for example, we can have light, medium, and dark on the first level, various colours within these categories on the next level, a division into shades on additional levels, and finally the terminal elements – the groups of entities that have a precise hue. And if the attribute is the time when events occur, we can have years on one level, the twelve months on the next level, then days of the month, hours of the day, etc., down perhaps to fractions of a second. Note that, no matter how many levels and how many elements per level we create, these structures remain correct hierarchies as long as the elements are depicted strictly as categories within categories.

It is for this reason that we can ignore here the additional levels, and limit ourselves to the simplified hierarchy, as if the attribute had only two significant values. The actual structure is much larger, but the additional details are unimportant; all we want is to confirm that elements which share more than one attribute cannot be represented with a simple structure. To put this differently, if it is impossible to combine several attributes in a *simplified* hierarchy, there is no point in demonstrating this problem with the larger, actual hierarchies.



Here is why one structure cannot represent more than one attribute, when attributes may be shared by all entities. If we want to represent several attributes with one structure, we must use one of them for the top element and the rest for the lower elements, within one another, as if each attribute were a finer detail of the previous ones. This problem is illustrated in figure 1-3. If we start with the totality of entities, then *E100*, *E110*, etc., are the categories of entities formed when separating them, repeatedly, based on their attributes (*A1*, *A2*, and *A3*). Each attribute has only two values, and the *E* digits designate the combinations of values that define the various categories. But this is an incorrect hierarchy, because the attributes are repeated: each one must be included for each branch created by the previous attributes. In reality, entities that possess several attributes possess them all in the same way, not as one within another. Thus, the diagram is incorrect because it shows some attributes as subordinate to others while the attributes are, in fact, independent of one another.




Figure 1-3

In addition to the repetition of attributes, we note the absurdity of this representation in that we can depict the attributes in any order while the terminal elements remain the same: *A3* within *A2* within *A1* (as in figure 1-3), or *A2* within *A1* within *A3* (as in figure 1-4), etc. This freedom means that the structure is illogical. Real entities cannot possess the same attributes in several ways.

A simple structure, therefore, does not represent real entities and attributes correctly: this is neither the way entities actually exist, nor the way we perceive them. We can develop a replica of the complex phenomenon *in our mind*, but

we cannot represent the phenomenon with a precise diagram. All we can do is depict the three attributes as *separate* structures, while knowing that the terminal elements are in fact the same in all three structures (see figure 1-5). We will discuss this problem in greater detail in chapter 4.




Figure 1-4

When each attribute is shared by only *some* of the entities, it may be possible to include several attributes in one structure, as one within another. In figure 1-6, for instance, while *A1* is shared by all the entities, *A2* and *A3* are shared by fewer entities, and *A4*, *A5*, *A6*, and *A7* by fewer still. None of the entities that possess *A5*, for example, possess *A6*. The final structure *is* a correct hierarchy, since attributes are no longer repeated. Moreover, there is only one




Figure 1-5

way to depict them: we cannot modify their order here as we could in the structure of figure 1-3, where all entities possess all three attributes. A6, for example, *must* be within A3, since only *some* entities possess A3 and A6, and the others possess A3 and A7. Also, it is impossible to redraw the structure so as to place A3 within A6, rather than A6 within A3.

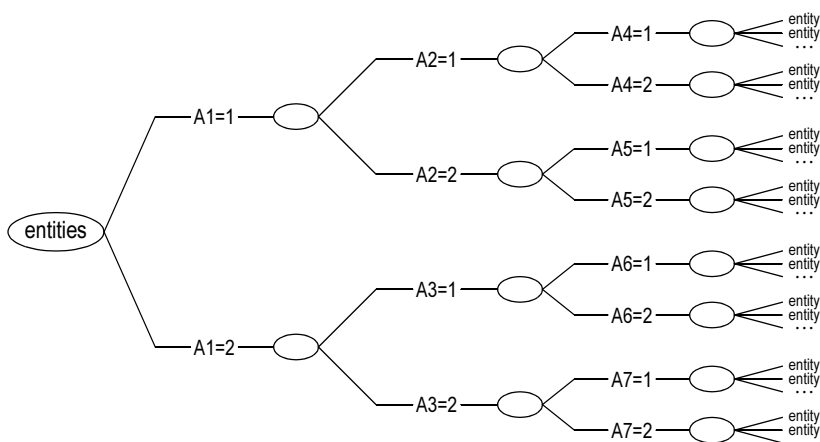


Figure 1-6

One can argue, however, that in fact we can *never* represent a combination of attributes in one hierarchy, as in figure 1-6. This is true because, even when an element does *not* possess an attribute, it still belongs, logically, to the structure established by that attribute: instead of being part of a branch with a particular value, the element is part of a branch called “does not possess the attribute” (as in figure 4-8, p. 359). In other words, because not possessing an attribute may be as significant as possessing it, we can combine several attributes in one hierarchy only when we can indeed ignore the cases where an attribute is not possessed by an element. (Thus, we cannot ignore these cases in *software* structures; for, attributes affect the performance of an application just as significantly when *possessed* by a software element as they do when *not possessed*. For example, an application will malfunction either because a database field is used when it shouldn’t be, or because it is not used when it should be. We will discuss software structures in chapter 4.)

Very few phenomena, in the end, consist of entities that possess a perfect combination of attributes – a combination that permits all attributes to be represented together in one hierarchy, as in figure 1-6. The classification of animals provides a good illustration of this problem. An attribute describing their teeth, to pick just one example, is shared by animals that belong to

different branches of the formal classification (the neat hierarchy of classes, orders, families, genera, and species). In other words, if we classified animals through a hierarchical structure based on their teeth, the hierarchy would be different from the formal one. Thus, since they possess many attributes of this type, animals are related through many hierarchies at the same time; and the only way to represent these relations in one diagram is by repeating attributes, as in figure 1-3. The formal classification avoids this repetition and remains a correct hierarchy simply because the scientists who designed it restricted themselves to attributes that *can* be represented within one another, as in figure 1-6. It is an artificial classification – correct, but capable of depicting only *some* of the relations between animals. It is impossible to represent *all* their relations with one hierarchical structure.

If you still don't see why the repetition of attributes in a structure indicates that the structure is complex, look at it this way: If entities belonging to different categories (say, *E100* and *E200* in figure 1-3) possess the same attribute (*A2*, in this case), and some of them (those belonging to *E110* and *E210*) even possess the same *value* of that attribute, it means that these entities are logically related; and, through this relation, they form a hierarchical structure that is different from the structure we see in the diagram. This second structure reflects the way the entities possess the attribute *A2*; but it is not manifest – because it is dispersed throughout the structure shown in the diagram – so to detect it, we must identify and combine *in our mind* the categories concerned. (*A3*, thus, gives rise to a third structure, and in its case we must identify and combine in our mind other categories.) Redrawing the diagram as *A1* within *A2* would not solve the problem, of course: the structure representing *A2* would then be manifest, but the one representing *A1* would not.

3

Hierarchical structures are models of reality, and if we can have two kinds of structures, simple and complex, it is because there exist two kinds of phenomena: those that can be represented as simple structures, and those that can only be represented as complex structures. We have equated simple structures with mechanism, so we can call the complex structures *non-mechanistic*. It is the central tenet of my argument that complex structures cannot be reduced to simple ones, and therefore certain aspects of reality cannot be represented with mechanistic models.

The immediate consequence of this claim is that the methods we have developed for studying mechanistic phenomena are inadequate for complex ones. These are the *deterministic* methods – those that attempt to represent a

phenomenon with precision. They include, specifically, mathematics and software. The fact that many phenomena cannot be represented mathematically is well known. What is more interesting for us is that the same phenomena – namely, indeterministic phenomena – cannot be represented with software either. Thus, complex phenomena cannot be exactly simulated or explained with software (although they can be *approximated* with software). Note, however, that software can *generate* complex phenomena: a running application is a system of interacting structures, and its behaviour cannot be represented with a mechanistic model.

At first sight, it may not be obvious why complex structures cannot be reduced to simple ones. If complex structures are nothing but systems of interacting simple ones, why can't we study the simple ones individually with mechanistic methods, and then combine the results so as to predict the behaviour of the system as a whole? The reason is the mutual influence of interacting structures: one structure affects the others at the same time *it* is affected by *them*, so they cannot be studied separately. While a solution may be found in trivial cases, no solution is possible for real-world phenomena. In general, the structures can share any number of elements, at any number of levels; they can overlap and intersect in any conceivable fashion.

This becomes evident if we recall the long list of attributes that an entity can possess, each attribute causing it to be an element in a different structure. Since the same entity can be an element at different levels in different structures (when we have several levels of detail for each attribute), an interaction can arise through a feedback loop, as follows: element *e* in structure *A* is also an element in structure *B*, so it affects the elements at higher levels in *B*; but if one of these elements is shared with structure *A* and is at a lower level in *A* than *e* is, then *e* will be affected by *B*, thus closing the loop. With this feedback, the value of the top elements in *A* and *B* may well be unpredictable. Also, a loop can link more than two structures (*A* to *B* to *C* to *A*), and several loops can exist at the same time. In practice, such interaction is not only possible, but probably very common.



The conclusion we must draw from the foregoing analysis is that there can exist no simple structures in the world, but only complex ones. This is true because any element of a simple structure, if one existed, would also be part of other structures, through its other attributes, thus making the simple structure immediately part of a complex one.

But we saw that many phenomena *can* be represented as simple structures, so how can we explain the contradiction? The answer is that those structures

which appear to be isolated simple structures are *approximations* of the phenomena. Although their elements take part in the operations of several structures, the operations of one structure are much more important than the others. Or, to put it differently, the attributes that determine their relations in one structure are much more important than the attributes that make them part of other structures. If this is true for all the elements in the structure, the structure behaves in practice approximately like an ideal, isolated simple structure. It is entirely up to us to decide what level of approximation is acceptable, and that depends on how we plan to use the structure.

Classifications of things, for instance, are possible only if we take into account *some* of their attributes, and ignore the others. The formal biological classification of animals is a perfect hierarchy of classes, orders, families, etc., as noted earlier, only because we agreed to consider just *a few* of their attributes. If we wanted to include *all* their attributes, no simple hierarchical structure would be possible. The totality of animals, with the relations generated by all their attributes, is a complex phenomenon, and can be represented only with a complex structure. But in our studies of animals we needed a strict hierarchy, so we settled for an *approximation* of the phenomenon.

The most famous example of mechanistic approximation is Newtonian mechanics – the theory that, for two centuries, served as a model for all phenomena involving matter and force. Then, the growth in knowledge pointed to some anomalies in its predictions, and it was replaced with a more accurate mechanistic model: Einstein's theory of relativity. But we continue to use the Newtonian model where its approximations are useful, because it is simpler. There are difficulties with the relativistic model too, and scientists, including Einstein, have been searching for a better theory. Einstein predicted that one day his theory would be replaced by an even more accurate one, and would then be seen as merely a better approximation of reality than Newton's. Despite repeated failures, however, he continued to believe in complete determinism: he was convinced that the next theory would also be mechanistic. Many scientists today doubt that a mechanistic model can provide a complete and exact explanation for all phenomena, from the smallest particles to the whole universe.

4

Once we recognize that mechanistic theories can provide only degrees of approximation and not a complete explanation, we begin to understand why it is so easy to fall prey to mechanistic delusions. The problem is not that mechanistic theories provide only approximations. This is not a weakness, but

their strength; for, if we insisted on a perfect representation of reality, we would never find a working theory and the world would appear incomprehensible. Approximate theories about the world are the only explanations we can hope to find in a world where no phenomenon can be totally isolated, where none is truly mechanistic. And we can often expect improvements in these models as our knowledge, methods, and instruments progress.

The problem, rather, is to determine whether the approximations are practical, whether they can model the world accurately enough to be useful. It is childishly easy to find a mechanistic theory for a particular phenomenon if we don't care how poorly it approximates reality. As we saw earlier, there are patterns and regularities, and hence hierarchies and mechanistic models, everywhere we look. What this means is that we can always find some formulas or diagrams that match reality fairly well in certain situations.

And this is how a mechanistic delusion is born. Formulas and diagrams always look neat. So, even though the model works only in a few situations, we are tempted to conclude that, being neat, it must be important. After all, it is acceptable for a model to provide only an approximation. One day, we think, it will be accurate enough to be useful.

The anomalies displayed by a mechanistic model when approximating a complex phenomenon spring from the fact that it is a simple structure struggling to represent a complex one. Since it can represent *only one* of the structures that make up the phenomenon, we must ignore the interactions between this structure and the others. We must ignore, in other words, those attributes that relate its elements to the other structures. The more important the ignored attributes, the stronger the interactions with the other structures, and the poorer the approximation when these interactions are ignored.

It is this simple explanation that the mechanists fail to grasp. They notice the abnormal relations that exist between the elements of one level and the next – the deviations from the structure's definition. But their assumption of isolation prevents them from recognizing these anomalies as interactions with other structures. Since they only see one structure, they conclude that the anomalies are caused by some deficiencies in the structure's definition; so they try to correct these deficiencies. They may work with a tree diagram or with a mathematical representation of the model. They may modify or add elements, operations, branches, or levels, and they may even diverge from a hierarchical appearance. But they never give up the requirement for complete determinism. They fail to see that, no matter what changes they make, if the behaviour of the resulting structure can be described completely and precisely, it is still a deterministic model, still a simple structure. Even if it no longer looks like a tree diagram, we know that there exists a hierarchical structure (although perhaps a very complicated one) that precisely describes its behaviour.

Some of this work may improve the model, but the need for enhancements never ends: the model never reaches a state where it can provide a practical approximation of the complex phenomenon. All the time, however, the mechanists are convinced that they are making progress, that what they are doing counts as research. They believe that, if the process of improving a mechanistic model often leads to a working theory in the exact sciences, it must also be effective in *their* discipline. But this process succeeds in the exact sciences because *there* we can often isolate a simple structure from the complex whole and use it as an approximation. In disciplines that deal with human minds or human societies, the links between the interacting structures are too strong to be ignored. There seldom exists one dominating structure, and any research project that assumes so becomes a mechanistic delusion: a futile attempt to approximate a complex structure with simple ones.



Note the informal nature of this discussion. Terms like “interaction” and “feedback” are quite vague, but I cannot improve on this. I will make no attempt, for instance, to draw several simple structures and somehow connect their elements to depict a complex structure; or to find some equations that explain how the behaviour of a simple structure is modified by its interaction with the others; or to develop some software that simulates these interactions. Many researchers fall into this trap, failing to see that, no matter how they approach this problem, their work always amounts to the same thing: an attempt to reduce a complex structure to simple ones. This is an impossible task, because what it seeks, in effect, is to represent an indeterministic phenomenon with a deterministic model.

Thus, the structures I described earlier (where entities are related through each one of their attributes) are merely a way to demonstrate the interactions caused by shared attributes. It is impossible to analyze the actual interactions, or the results. If we could perform such an analysis, if we could describe the interactions with precision, if we could explain and predict them with diagrams or equations or software, then we would have no complex structures to begin with. For, if the behaviour of a particular complex structure could be exactly described, there would perforce exist a certain simple structure that displayed exactly the same behaviour. This is true because simple structures are logically equivalent to deterministic theories and methods: any phenomenon that can be described precisely and completely can be represented with a simple hierarchical structure.

I don't need better arguments, because what I claim, ultimately, is only this: all phenomena are complex structures, and are irreducible to simple ones

(i.e., deterministic models); while some can be approximated with simple structures, for others no useful approximation exists. If I hold this thesis, I need to prove nothing. What I have done is shift the burden of proof onto those who claim otherwise – onto the reductionists – which is the only logical way to address this problem. I hold that the normal state of affairs in this world is complexity, so it is up to those who insist that the world is simple, or those who promise to reduce complexity to simplicity, to prove their claims. A problem is deemed to have no solution until one is found: we wouldn't take someone seriously if he maintained that all problems have solutions, so there is no need to *actually* solve a given problem. Yet this is exactly the position taken by those who confidently set about to reduce a complex phenomenon to simpler ones, without trying first to solve what is in fact the primary problem: whether mechanistic methods can work at all for that phenomenon.

We must not forget that mechanism is only a doctrine, not unlike the religious doctrines of the past, and true scientific thought does not accept methods on the strength of their popularity or past reputation. We must not let the early successes of mechanism tempt us to see it as the only valid method of science. It must earn our respect like any other method, by solving *today's* problems.

5

Let us examine a few complex phenomena and interpret them as complex structures. The simplest complex structure is probably the phenomenon arising when three bodies attract one another according to the law of gravitation, in what is known as a *three-body* system. The attraction of the bodies in a *two-body* system, like the earth and the moon, and consequently their motion, can be described mathematically with fairly simple equations. If we add just one more body, however, as in the system comprising the sun, the earth, and the moon, we can no longer have an exact mathematical representation. The phenomenon of combined attractions in a three-body system is so complex that it is impossible to express the motion of the bodies analytically for a general case.² Intuitively, we can see the difficulty of describing the interaction between *A* and *B* if at the same time *B* interacts with *C*, and *C* interacts back with *A*. When more than three bodies are involved (*n-body* systems), the phenomenon is, of course, even more complex.

² It is possible for special cases, as when one of the masses is small enough (relative to the other two) to be negligible; otherwise, the motion can be calculated with various methods of successive approximation.

What we witness here is more than just an increase in complexity from two to three bodies, which would merely necessitate more complicated equations. The two kinds of systems appear to create entirely different phenomena. But the bodies in a three-body system are not different from those in a two-body system; and it is the same gravitational force that governs their motion. We must conclude, therefore, that the two phenomena are very similar in reality, and only appear different to *us*, when we try to represent them with our mathematical, and hence mechanistic, models.

The two-body system is a simple structure. We can view it as a trivial hierarchy of two levels: the two bodies, *A* and *B*, are the terminal elements, gravitational attraction is the operation performed on them, and the system as a whole is the resulting higher-level element. The three-body system is a complex structure. One way to view it is as a system of three simple structures (three two-body systems: *A* and *B*, *A* and *C*, *B* and *C*), which share elements and therefore interact with one another. It is not surprising that we can find a mathematical representation for two-body systems but not for three-body systems, since we already know that only phenomena which are simple structures can be reduced to mathematical models. Any attempt to find an exact model for a three-body system amounts to an attempt to reduce a complex structure to simple ones, which is impossible.



Another complex phenomenon is the recognition of a human face. Even though we can easily *recognize* familiar faces, we cannot as easily *describe* them. That is, we cannot depict in words a familiar face with sufficient precision so that someone unfamiliar with it could recognize it as easily as *we* do. And this is not due to a deficiency in our languages; for, if we were permitted to create new words or even to invent a new language to describe faces (using perhaps a special notation, as we do in mathematics), we wouldn't know what words or symbols we needed. This is what we do, in essence, when programming a computer to recognize complex patterns, and the performance of this type of software is always inferior to our own capabilities. We can recognize a face because the mind can process complex structures. Describing the same face with a system of symbols reduces the recognition process to simple structures, and this can only *approximate* the actual phenomenon.

When we recognize a face, our mind processes it as a whole, not by reducing it to simpler patterns first. We don't know how the mind does it, but we can try to understand this process by representing it as a complex structure. It is obvious that we do not remember faces by storing picture-like images of them, but by storing certain information about them. For example, we can recognize

a face – even a barely familiar one – from any angle, not just as we previously saw it. There can be little doubt, therefore, that we recognize faces thanks to their attributes: the size, shape, and proportion of the various parts of the face; the relations between sizes, shapes, and proportions; the combinations of relations and proportions; and possibly even the combinations of these combinations.

If I wanted to describe to you the face of a person I know well, so that you could recognize that person in a crowd, I would have to be aware of all these attributes, relations, and combinations. But these facts do not exist in my mind in a form that I can use consciously. I never had them separated and classified in my mind, not even when I first learned to recognize that face; for, I didn't learn to recognize it by matching it with a list of attributes. So the best I could do is watch or visualize that face, and convey to you its attributes – most of which I would probably notice for the first time. Then, if we are both skilled in this task, you may even manage to recognize that person using my instructions.

At that point, you are acting the way a computer would; namely, matching faces with a list of attributes. But as you get to know that person, you will begin to recognize his face instantly, intuitively, as I do. You will not simply run that “program” – that list of instructions – in your mind faster and faster. You will have acquired some knowledge which lets you recognize that face *without* matching it with a list of attributes; but you will have no idea what that knowledge is. If asked to describe the face, you would have to draw up a list of attributes, as I had done earlier.

We can think of a number of individual processes that make up the complex process of face recognition: we may perceive the basic arrangement of the principal parts (eyes, nose, mouth, chin, hair) as a generic template that we interpret as a human face, the actual faces being then particular deviations from this template; we may perceive the sizes of these parts as big or small, their shapes as round or long, their proportions as average or not, and so on. Each one of these processes could be represented with a simple structure, because the elements and relations that make it up can be precisely described. But these structures share their elements – the various parts of a face – so the whole process of face recognition must be a complex structure. When we describe a face with words, we create in effect an approximation of the complex structure by means of several simple structures; and, depending on the face and the description, this may provide a useful substitute for the process of recognition.



Our linguistic capability provides another example of complex phenomena. We will study the important subject of language and its relation to software later, but we can already see why our capacity for language is not a mechanistic phenomenon.

When we think of language, we immediately think, in the case of written text, of a hierarchical structure of letters, words, sentences, and paragraphs; and in the case of speech, we think of a hierarchical structure of phonemes, words, sentences, and ideas. We can also view each sentence, though, as a hierarchical structure of grammatical units – words, phrases, clauses. Moreover, in addition to being an element in the structures just mentioned, each word has a meaning: it names a thing like an object, an action, or a concept. This makes it an element in yet another structure, one including related objects, actions, or concepts. But we saw that entities have in fact *many* attributes, and are therefore related in many different ways. So, through the entities they represent – through their meanings – words are elements in *many* additional structures. Language is possible because the mind can process all these structures simultaneously.

The phenomenon of language, thus, is the interaction of many structures. It can have no precise representation, no mechanistic model, and this ought to be obvious to anyone who thinks about it. This has not stopped scientists and philosophers throughout history, though, from attempting to find a mechanistic model. Nor has this stopped them from searching for an artificial language – a language they hope would have the same potency as the natural ones while being simpler, and hence reducible to a mechanistic model. (We will study these delusions in chapter 4.)

The reason it is tempting to seek a mechanistic model for language is the ease with which we can discover and isolate linguistic structures. All we have to do is extract one of these structures from the whole phenomenon of language. Then, if we base a theory on this structure, the theory is guaranteed to work for *some* sentences. In general, it is easy to find a simple structure that represents *approximately* a complex one, and hence a mechanistic theory that *approximates* a complex phenomenon. These approximations typically yield theories that work in some situations but not in others, and this explains their occasional success. Thus, if the approximation is useful, we may be doing valuable research. In the case of language, however, the interactions are too strong. Consequently, no mechanistic theory can usefully approximate the whole phenomenon of language.

Abstraction and Reification

1

Two great fallacies arise from mechanistic thinking. Let us discuss first the fallacy I mentioned earlier, when we studied simple hierarchical structures: the belief that starting from higher levels of abstraction confers certain benefits. We saw that, as we raise the level of our starting elements, we reduce the number of values possible for the elements at the higher levels. Ultimately, we reduce the number of values that the top element, our final goal, can take.

Starting from higher levels of abstraction, thus, causes an *impoverishment* – a reduction in alternatives. The versatility of the hierarchical structure derives, not from a large number of possible values for the starting elements, but from the large number of combinations of values generated at the higher levels. For a hierarchy to be practical, the number of values we start with must be small. To attain a large number of values at the top level, we increase, instead, the number of levels and the types of operations that relate them.

When starting from a higher level, we could still have, in principle, as many values as we had at that level previously. We could, in other words, define as a set of starting values the *combinations* of values occurring at that level when our starting level was the lower one. In practice, though, this set would be so large that we would use only a fraction of it. Thus, when we lose the lower levels of a hierarchy we are bound to lose also many combinations of values.

Starting with only twenty-six letters in the English alphabet, for instance, we can create thousands of words one level up, and an infinite number of sentences on the next level. If we were to start with *words*, each word having its own symbol (that is, an elementary symbol, independent of the symbols of the other words, just as the symbols for letters are now independent of one another), our language would be impoverished. We would have to limit our vocabulary to a practical size (say, a few hundred symbols), and communication would be very limited. And if we were to skip one more level and start with *sentences* (assigning a symbol to each sentence and restricting ourselves to, say, a few hundred sentences), communication would break down completely.

Note that this fallacy can be committed even with simple structures, since it involves only one hierarchy; so it can be committed even when a mechanistic model is otherwise adequate. I will use the term *abstraction* to describe this fallacy, but bear in mind that “abstraction” also refers to the normal transition from one level to the next. Abstraction, therefore, means both the process of generalization that is part of any hierarchical structure, and the mistaken view that we can skip the lower levels of the structure. (Although the sense in which

the term is used will be obvious from the context, I will sometimes use the whole phrase, “fallacy of abstraction,” for the second sense.)

It is especially easy to commit this fallacy when the elements of the hierarchy are abstract concepts. We may think of the “average” man or woman, for example, or the “typical” salesman or accountant, and this may be a useful concept. But we can never *meet* an average or typical person; there is no such being. So when we treat an actual person, or think of ourselves, according to this idea, we are committing the fallacy of abstraction. In a hierarchy, if the average person is an element at a certain level of abstraction, it subsumes millions of real individuals who are elements at a lower level. The elements at the lower level are real things, whereas those at the higher level are abstract concepts. This situation is found in many hierarchies. (Classifications, for instance, have actual things at the lowest level, and categories of things – i.e., abstract concepts – at the higher levels.)

But the fallacy of abstraction is not limited to a transition from concrete things to abstract concepts. It can be committed even when both levels have concrete things (as with the words and sentences we saw previously), or when both levels have abstract concepts. The fallacy occurs whenever we start illegitimately from a higher level of abstraction. When we accept the elements at the higher level as *starting* elements, what we do is perceive them mistakenly as similar in nature to those at the lower level, as providing the same versatility. So the consequence of abstraction is a reduction in alternatives: *the structure is impoverished through the destruction of levels.*

It is easy to see how this fallacy can be exploited, if an elite gains control of the levels of abstraction in an important structure. Our alternatives may be restricted even as we think that we are gaining something. The elite can tempt us to start from higher levels by promising us expedience: why start from low-level elements when high-level elements are available? Since each high-level element includes a combination of many low-level ones, the elite tells us, we will reach the top element, our goal, much faster.

Now, when the high-level elements are provided *in addition* to the low-level ones, as an *option*, we may well find them more effective in certain situations, and we are committing no fallacy in using them. But they are usually provided as a *substitute* for the low-level elements, not as an option. Once we lose the lower levels, we lose countless alternatives for the top element – alternatives that may be important. If we forgo those alternatives, our life will be impoverished. And if we do want them, we will depend on the elite, which alone can access the low levels. Having lost the capability to create those alternatives on our own, we are at the mercy of the elite every time we need a new alternative.

Abstraction can be subtle. A common way of losing alternatives without

realizing it is by confusing freedom of choice with a large number of alternatives. But a large number of possible values for the top element may be only an *illusion* of freedom: we may have a large number of values, even an infinite number, and still be severely restricted.

This problem is related to the unusual properties of large numbers. To understand this, imagine a system represented by a structure where the elements are numeric values, and the top element – the result of various operations, performed on several levels – can be any integer. The top element, thus, can have any one of an infinite number of values. Now imagine that, through the elimination of one level, we restrict it to integers divisible by 10; we are left with only one tenth of the original values, but we still have an infinite number of them. And if we eliminate further levels and thereby restrict the top element to integers divisible by 100, and then to integers divisible by 1,000, and so on, we are left each time with only one tenth of the values previously possible. There will be fewer and fewer values, but we will continue to have, nevertheless, an *infinite* number of values.

So, even though the alternatives for the top element are being reduced to the point where the system may become useless, if we judge it by the sheer *number* of alternatives we may feel that we haven't lost much; after all, there are still an infinite number of them. This paradox is obvious in a trivial structure like the one just described, especially if we are already familiar with all the alternatives. But it may be hard to detect in a structure where the values possible for the top element consist of *novel* alternatives – alternatives which we never encountered before, and which we cannot even imagine in advance. In this case, we may not even realize that we are missing alternatives.



The more ignorant we are, the easier it is for an elite to gain control of the levels of abstraction in structures on which we depend, and to exploit us by eliminating alternatives. Thus, widespread programming incompetence has permitted the software elites to establish a sort of business best described as software charlatanism. The chief goal in this business is to destroy levels: under the pretext of efficiency, the elites are constantly raising the level of abstraction in development systems. When starting from higher levels, they tell us, we reach the top level – the complete application – much sooner. Certain features may no longer be possible, it is true, but we still have an *infinite* number of alternatives. So we can implement about the same applications as before.

We commit the fallacy of abstraction, however, if we interpret the infinity of alternatives as evidence that we have lost only a few. This will become clearer when we discuss the second mechanistic fallacy, reification, because the

two fallacies are usually committed together. Applications comprise *many* structures, not one; so when starting from higher levels we lose also the low-level *links* between structures, and with that further alternatives.

We are aware of lost alternatives only if we once had them. If we had to abandon words, for example, and restrict ourselves to ready-made sentences and ideas, we would immediately recognize the dramatic impoverishment in language-related processes. But programming is a new human endeavour, and the software charlatans gained control of our software-related affairs, and restricted them, before we could discover all possible alternatives – all the ways that human minds can find to create and use software. Software controlled by an elite is the only software we have ever had, so we cannot know what we have lost.

To enable us to create useful applications, the elites often restore some of the low levels. But they do it through some complicated means, which they control. So, instead of being free to create any application, we now depend on the elites for their high-level systems, and also for the low-level elements that had previously been available to us directly. (We will study this charlatanism in “The Delusion of High Levels” in chapter 6.)

2

If the fallacy of abstraction can be committed with simple structures alone, the second mechanistic fallacy is committed with complex structures; specifically, when we extract the simple structures from the whole that is a complex structure.

We already know that a mechanistic theory employs a simple structure to approximate a complex phenomenon. The approximation is achieved by isolating one structure – one aspect of the phenomenon. This is a legitimate procedure when the isolated structure approximates the complex one well enough to be useful; in other words, when its interactions with the other structures are much weaker than its internal relations. If, however, we attempt to extract a structure when the interactions are too strong to be ignored, we are committing a fallacy: the resulting simple structure is not a useful approximation of the actual phenomenon; if employed as model, it will not represent correctly the behaviour of the complex structure. I call this fallacy *reification*, borrowing a term that is already used to describe similar fallacies.

In philosophy, reification is often used to describe the fallacy of perceiving an abstract or hypothetical concept as a real thing. We do this, for example, when we treat transcendental entities as similar in nature to the concrete things of everyday life. Outside philosophy, however, the idea of reification has gained

a broader meaning, and is used in any situation where this type of fallacy is committed. In psychiatry, for example, reification describes a common thought pattern displayed by schizophrenic patients. Many schizophrenics are incapable of comprehending a particular topic while viewing it as part of a context. They extract it from the complex reality and treat it as a separate entity, or they attach it to a wrong context. Most topics, however, have different interpretations in different contexts, so schizophrenic conversation and behaviour is often incoherent. To take another example, in certain social theories, reification (along with notions like alienation and false consciousness) is used to describe our tendency to perceive human lives and societies as made up of separable parts. We isolate a person's knowledge or skills, for example, from the whole that is his existence, and thereby distort our social relations. It seems logical then to rate and to purchase a person's skills or time as if they were objects the person owned.

What is common to these examples is a mechanistic form of thinking: an illegitimate attempt to reduce a complex phenomenon to a simple one by taking something that is part of a whole, something that cannot possibly exist in isolation, and treating it as a separate thing. Since that aspect can exist as a separate thing only in our imagination, what this type of thinking does is objectify an abstract concept. It is quite appropriate, therefore, to use the term "reification" for the fallacy we are considering here – separating a complex structure into simple ones. If the simple structure we attempt to isolate has strong links to the others, it can exist as a separate structure only in our imagination. While the complex structure represents reality, the individual simple structures that constitute it are imaginary: they cannot exist as separate things. Or, more accurately, when viewed as independent structures they represent a *different* reality – an approximation of the complex structure. The fallacy is in the belief that the approximation is close enough to be useful.

A round and red object is always both round and red. It cannot have only shape or only colour, so it cannot occur separately in a structure of shapes or in a structure of colours; it *must* exist in both structures at the same time. In our imagination, we may be able to extract the object's roundness or redness. We may be able, that is, to construct an isolated hierarchical structure of shapes with no colours, or a structure of colours with no shapes – structures that would function as classifications of objects (see figure 1-2, p. 96). But these structures cannot exist in reality. If shapes and colours are defined as attributes of objects, the only structures of shapes and colours that can exist in reality are those that interact and make up a complex structure, because they always share their elements (the objects).

We can represent shapes or colours with hierarchical diagrams of categories, describe them in words, or define them with codes. What these devices do is

convert our imaginary structures into real ones, and they may be useful when we can study each structure on its own; for instance, when only the shape or only the colour of objects is important. But if *both* are important, then the two structures interact, and the diagrams and codes are a form of reification: we convert into real things, concepts that can exist only in our imagination. At this point, it may seem quite natural to treat objects on the basis of diagrams or codes, as if these devices embodied the same knowledge as the complex structure that represents the actual objects.

We can view, for example, all the objects in the world as a simple hierarchy with two branches leading to two categories, round objects and other objects. Our round and red object will then be one of the terminal elements branching out of the round-objects category. We can also classify all the objects in the world into red objects and other objects, and our object will be one of the terminal elements in this hierarchy too: an element branching out of the red-objects category. But clearly, the terminal elements in these two hierarchies represent the very same objects. Since an object cannot be in two places at the same time, these hierarchies do not represent reality, although they are correct from the perspective of the *individual* attributes. The only way to take into account *both* attributes, roundness and redness, at the same time – which is the way they actually exist – is by treating the whole phenomenon of objects and their attributes as a complex structure. In our mind we can readily do so (because the mind can process complex structures), and this is why we can *perceive* objects as being round and red at the same time. But we cannot represent this phenomenon accurately with mechanistic models (because it is impossible to combine several attributes in one hierarchical structure when the attributes are possessed by all the elements, see pp. 98–102).

To put this differently, if objects and their attributes could be represented with mechanistic models we would be able to represent the world with software, and our computers would perceive reality just as *we* do. But, in fact, we cannot program a computer to perceive even that single round and red object as it really is, as humans perceive it. If we represent the two attributes as two code systems, for example, the computer will “see” first one code and then the other; but it cannot see both codes at the same time. The attributes, however, do exist together in reality, so we must conclude that the computer does not represent the phenomenon exactly as it is.

Thanks to its ability to process data quickly, the computer may provide an *approximation* of a complex structure – if we manage to reduce the phenomenon to its most important structures and interactions. For example, we could use *one* system of codes instead of two, and represent *combinations* of shapes and colours; the computer will then “see” the roundness and redness simultaneously. But this method, while adequate in the trivial case of one

object and two attributes, breaks down in practical applications, when we must represent *many* entities and *all* their attributes. And the reason is not only that there is no computer fast enough to simulate so many interactions, but that we don't even know how to identify for the computer all the structures involved, and all their interactions. (It is for this reason that the attempts to make computers perceive the world as humans do cannot advance beyond trivial situations. We will return to this issue in chapter 2.)



Three kinds of structures exist in the process of reification: the complex structure, which alone reflects reality, and which is the only true representation of the complex phenomenon; the simple structures that exist only in our imagination, when we view the complex structure as several interacting structures; and the real, reified simple structures that we create from the imaginary ones. The fallacy consists in thinking that the real structures we created ourselves represent the same phenomenon as the imaginary ones.

It is easy to demonstrate the fallacy: when we reconnect the real structures the way we think the imaginary ones were connected, we will not re-create the same phenomenon as the one represented by the original, complex structure. In practice, if our intent was to reduce the complex structure to simple ones, what we see is that our model, theory, or method does not work as we expected: it does not explain the complex phenomenon adequately, and the reason is that it cannot account for all the alternatives that the phenomenon can display. *For a phenomenon that is a complex structure, the only exact representation is the phenomenon itself.*

Recall the complex structures we studied earlier. We may well think of a three-body system as made up of three two-body systems, but these two-body systems can exist only in our imagination. We may be able to separate, for example, any two of the bodies from the third one (or perhaps the two did exist once as a separate system); but this isolated two-body system behaves differently from the two bodies that are part of the three-body system. We have two kinds of two-body systems: the real ones, and those we visualize as making up the three-body system; and the latter *must* be imaginary. If we separated bodies *A* and *B* as one system, how could we separate at the same time *A* and *C* as a second system? *A* cannot be in two places at the same time. The phenomena peculiar to complex structures are caused precisely by the fact that its elements belong to several structures at the same time.

Similarly, when recognizing a face we may well *imagine* several separate processes, but these processes cannot be separated. There may exist, for example, a process involving the size of the nose. But it doesn't follow that we

can view this process as a separable structure. To separate it from the other processes that make up this phenomenon, we would have to separate the nose from the face; and this would alter the other processes that involve the nose. So the simple structures that exist as separate things *in reality* (sizes, proportions, combinations – big eyes, long face, etc.) are not the same kinds of things as the equivalent simple structures that are part of the face recognition phenomenon, and which can exist only in our imagination.

We also note the fallacy of reification in the study of language. Linguists start by extracting language from its complex human context; then, they go even further and extract various aspects of language – syntax or semantics, for instance. Language, however, could not possibly exist separately from the human beings who use it; and human societies would not be what they are, were the form of language different. The two evolved together and are inseparable. We may view the linguistic and human structures – syntax, word meaning, social relations, individual behaviour, cultural traditions – as separable phenomena, if we want. But, clearly, they can exist as isolated phenomena only in our imagination: separating any one of them would separate many elements that also belong to the others, so they cannot actually exist as independent structures. The *real* phenomena studied under these labels are different, therefore, from the *imaginary* phenomena we perceive as parts of the phenomenon of language. The real phenomena constitute a reification, and this is why studying them cannot help us to understand the phenomenon of language.

3

We saw previously that abstraction – treating the higher-level elements of a structure as starting elements – impoverishes the structure by reducing the number of alternatives for the top element. Now we see that the effect of reification, too, can be described as an impoverishment. Similarly to abstraction, the impoverishment manifests itself as a reduction in the number of alternatives for the top element; but this reduction derives from a loss of interactions rather than a loss of levels.

The impoverishment caused by reification is probably even more severe than the one caused by abstraction, because even more alternatives are lost now. It is more difficult to recognize this loss, though, because we cannot describe the interactions between structures (which is where the loss occurs in reification) as precisely as we can the relations between the elements and levels of each structure (which is where the loss occurs in abstraction).

We are tempted to separate structures when we want to understand a

complex phenomenon. Separating structures, however, severs the links between them – links that create in fact the complexity. The complexity is due to the unspecifiable alternatives: those values or states which are caused by interactions, and which cannot be predicted from the properties of the individual structures. Reification eliminates the unspecifiable alternatives, and hence the indeterminism, but it is precisely this indeterminism that imparts to complex phenomena their richness and potency. It is silly to try to understand a complex phenomenon by eliminating the interactions, since this destroys its complexity: what is left is only the simpler, deterministic phenomena represented by the individual structures.

It is obvious, then, why reification is such a tempting fallacy. The only way to understand a complex phenomenon is by studying it as a whole, by relying on the mind's ability to process complex structures. (Understanding is then taken as an informal concept, not in the mechanistic sense of breaking down the phenomenon into simpler ones.) Thus, if there is a way to understand a complex phenomenon, it entails a great deal of knowledge and experience, intuition and creativity. Reification simplifies the task by separating it into isolated structures, which can be handled by researchers who lack these qualities. With isolated structures, they can apply mechanistic methods; in particular, separating these structures into even simpler ones. But mechanistic methods are futile in the case of complex phenomena. The isolated (real) structures studied by the mechanists are not the same as the (imaginary) structures that make up the phenomenon. So, even if they make good progress in their study of the isolated structures, the mechanists contribute nothing to the understanding of the phenomenon.

Isolated phenomena, thus, are academic inventions, reifications of complex phenomena, and their study seldom has a practical value. In the domain of programming, the fallacy of reification has given rise to the theories of software engineering, to concepts like structured programming and object-oriented programming. These theories have generated an enormous amount of academic research (not to mention the effort wasted by those who try to use them), while contributing nothing to the real issue – the advance of programming knowledge. Expert programming entails a capacity for complex software structures, because those phenomena we want to represent with software are complex; and the only way to develop this capacity is through lengthy practice. Reification creates isolated software phenomena – simple structures that reflect isolated aspects of reality. These structures are then within the capabilities of inexperienced programmers, but the problems associated with the original, complex phenomena remain unsolved.



I want to stress again how easy it is to confuse the structures we interpret as the components of a complex phenomenon, and which are only imaginary, with some independent, real structures. Even *describing* complex structures as “interacting simple structures,” as I do here, is a form of reification (because it suggests that the simple structures were once, or could also exist as, independent structures). We use such phrases because we have no other way to discuss this subject, but we must not forget that those simple structures exist only as a combination, only as the complex phenomenon. They cannot exist on their own. The *real* structures that we think are their identical counterpart are indeed independent, but they represent different phenomena: if we combine them, we will not reconstruct the complex phenomenon. *Thinking* of interacting simple structures may help us to understand the complex structure, but we must go no further.

It is with language, again, that the differences between the imaginary and the real structures are easiest to observe. We may think of language as one of the many phenomena that, together, make up a human society; but this is where we should stop. The very idea of “language” as a distinct subject of study is an absurdity. The (real) phenomenon studied by the academic discipline known as linguistics is not the same as the (imaginary) phenomenon of language that exists only as part of a human society. Were it the same, we would be able to implement our linguistic theories with software and then converse with our computers just as we do with people.

It is generally accepted that the difficulty of programming computers to understand language is not due to the intricacy of the language structures, but to the fact that every word, every expression, every idea, has rich meanings – meanings which depend on such knowledge structures as the current context, related ideas, and previous experiences. It is impossible to store this type of knowledge in a computer, because the only way to acquire it is by “being in the world”: by being born human, by having a body, by growing up in a human society. This knowledge cannot be reduced to a number of independent structures – facts, methods, processes – with language structures among them. It is a complex structure, the result of many interacting phenomena, and hence possible only in a human mind.

4

The two mechanistic fallacies, abstraction and reification, are usually committed together. It is easy to see why: we are tempted to abstract because we want to start from higher levels; but to abstract we need simple structures, so we first reify the complex structure. Reification impoverishes the complex

structure by eliminating the interactions between its component structures; and abstraction impoverishes each structure by reducing the number of levels. Each fallacy contributes its own kind of impoverishment, but the result is the same: a reduction in the number of alternatives for the values of the high-level elements.

Note that when we commit both fallacies we seldom keep the reified structures separate. We still need a complex structure, usually; so we combine the impoverished structures and allow them to interact again. But the number of interactions and alternatives possible now, with the low levels missing, is only a fraction of those present in the original phenomenon.

Note also that, while reification can be committed even without abstraction, abstraction always entails reification. Thus, we may reify a complex structure but retain the low levels of the individual structures; and the only alternatives lost would then be those resulting from the interactions. But we cannot start from higher levels in the complex structure without also causing some reification. The interactions between structures are due to the shared elements, especially at the lower levels, so when we lose these elements we lose levels *and* interactions. The higher the level we start from, the greater the effect of abstraction, but also of reification, because fewer elements are left that can be shared.

So, when we commit both fallacies it doesn't matter whether we view this process as abstraction, or as reification followed by abstraction, or as committing the two fallacies at the same time. All three interpretations are equivalent to committing abstraction on the individual structures while still part of the complex structure (because this abstraction would also reify them). Destroying the lower levels of the complex structure directly is, therefore, the same as committing reification first, then abstraction on the individual structures, and then allowing the impoverished structures to interact again. The result in both cases is a complex structure where the interactions and alternatives previously generated by the low-level elements are missing.



It should now be obvious why the model of simple and complex structures can help us in the two subjects that concern us in this book: the failure of mechanistic theories to explain complex phenomena, and the methods employed by charlatans to deceive us.

Regarding the first subject, both fallacies consist in the belief that we can simplify the complex structure that represents a complex phenomenon and still represent the same phenomenon. But the resulting structure can represent only a fraction of the alternatives that constitute the complex phenomenon.

If explanation means accounting for all the manifestations of the phenomenon, the mechanistic theories fail because they can only account for a small portion of the alternatives.

Regarding the methods used to deceive us, the same model can explain how they work. What the charlatans do, in essence, is tempt us to commit the two fallacies; they tempt us, in other words, to hold in our minds an impoverished picture of reality, of the facts and events that constitute our existence. Just like the scientists who construct invalid mechanistic theories about the world, we end up developing invalid mechanistic notions about our personal, social, or business affairs. It is not difficult then for the charlatans to exploit us.

This process is well understood for language. Practically everything in a human society is related to language: culture, traditions, social relations, knowledge, communication. Reasoning, in particular, would be impossible without language: no theory ever offered a satisfactory explanation of reasoning without also involving language, or vice versa. So it is not surprising that those who want to deceive us – leaders, gurus, visionaries, advertisers – always attempt to control language. And they achieve this by impoverishing it: instead of a rich vocabulary and precise ideas, they use slogans, jargon, and standard phrases, which subsume many meanings and thereby obliterate the differences between them. They also invent new terms, or use old terms in a new sense, but without properly defining them; so these terms too subsume, in effect, many meanings.

The charlatans, thus, employ forms of communication that start with high-level elements: their messages appear to convey information, while expressing in fact only vague, misleading notions. This practice makes it hard for us to connect these messages to our prior knowledge, and without the benefit of prior knowledge we are more easily deceived. Recalling the mechanistic fallacies, the deception is achieved by raising the level of abstraction, but also through reification; for, when losing the low-level elements, we also lose the low-level links between the language structures and various knowledge structures we hold in the mind.

If the complex structure is our existence, and the simple structures that make it up are the various aspects of life and of language, then a reduction in the ideas conveyed by language will impoverish not just communication but also other aspects of our life. By destroying alternatives in language, the charlatans can also destroy alternatives in our knowledge, in our beliefs, and in our expectations.

An extreme form of language manipulation was described by George Orwell in his account of a hypothetical society where English was replaced with Newspeak – an artificial language specially devised to control minds. By restricting the vocabulary to carefully selected words, Newspeak limits mental

processes to high levels of abstraction and prevents people from linking knowledge structures. This makes it impossible to develop any thoughts that contradict the official ideology. (We will study linguistic manipulation, and Newspeak, in chapter 5.)



Since software is acquiring in our society the same role as language, if we reify and abstract software we should expect to see similar effects. Software, in fact, exemplifies even better how the combination of the two mechanistic fallacies can be used to exploit us, because this has already happened – in the world of programming, at least. (Reification and abstraction in language and in software are discussed in chapters 4, 5, and 6.)

Software is one structure among the many structures that make up society: our affairs are increasingly dependent on computers, and hence on software; and our business and social relations increasingly reflect our dependence on software practitioners. For institutions as much as for individuals, knowledge and thinking are increasingly linked with software, in the same way they are linked with language. But, like language, software lends itself to abstraction and reification: we are just as easily tempted to start from higher levels of abstraction, and to treat the various aspects of software as separable processes.

Thus, it is common to think of a software application as an entity that interacts with the world only through its specifications, or parameters, or input and output; in other words, through means that can be defined precisely and completely. Once we commit this fallacy, we are likely to think of programming – the act of creating that application – as a project that can be separated from the other aspects of our affairs. Moreover, it is easy to commit further reifications *within* the act of programming: we extract certain software processes (user interface, reporting, database operations, etc.) and treat them as independent structures, linked to the rest of the application through precise specifications. But these processes share their elements (the software entities that constitute the application), so they form structures that interact with one another. Finally, these structures interact with the structures created by the business, social, and personal issues reflected in the application. All structures are different aspects of the same phenomenon – the complex phenomenon of software development and use – and we sever the links between them when addressing them separately.

In addition to *reifying* the application, we are tempted to raise the level of abstraction of the starting elements within the isolated structures. And this further reduces the number of alternatives that can be implemented. The interactions between the software processes that make up an application, and

between these processes and our affairs, must take place at low levels of abstraction (as low as the individual operations), because this is the only way to represent with software the *details* of our affairs. The experts, nevertheless, manage to convince us that starting from higher levels affords us efficiency, or “power.” The higher levels come in the form of programming aids – methodologies, development tools and environments, database systems.

Another way to start from higher levels of abstraction is to use ready-made applications. Since these applications cannot include all possible alternatives, they function in effect as high-level expedients. When relying on ready-made software, therefore, we lose low-level elements *and* low-level links – the links between software structures, and between software structures and the structures that make up our affairs. So we commit both fallacies, abstraction and reification, and the result is the same as when relying on programming aids: impoverished applications. As we saw earlier, committing both fallacies can be interpreted either as abstraction on the complex structure directly, or as reification followed by abstraction on the separated structures (see p. 121). With software, we see now, using ready-made applications is equivalent to the first interpretation, while using programming aids is equivalent to the second one.

As a consequence of these fallacies, more than 90 percent (and perhaps as much as 99 percent) of the cost of business computing is due to adopting new applications and programming aids over and over; in other words, to solving software-related problems instead of business problems. The reason for this incredible inefficiency is the incompetence of programmers. The complex phenomena of software can only be understood as a whole, and this requires great expertise. Programmers need to isolate software processes and to start from higher levels because they cannot deal with the complex software structures directly. But higher starting levels allow fewer alternatives. The resulting structures are impoverished, and so are the structures they interact with – those structures that make up the business environment. This impoverishment manifests itself in the inadequacy of applications: some of the alternatives – that is, some of the business needs – cannot be met. And this perpetual inadequacy compels businesses to seek new applications and new programming aids all the time.



Now that we have developed our mechanistic and non-mechanistic models, we can study complex phenomena and the failure of mechanistic explanations by representing them with these models. For each phenomenon, we must do two things: first, determine that it is in fact the result of several interacting

phenomena, so it can only be represented with a system of interacting structures; then, show that the proposed explanations are mechanistic, that they attempt to represent the phenomenon with isolated simple structures. We will use these models frequently in the following chapters.

The only way to study a complex phenomenon mechanistically is by extracting one of its structures. This is the fallacy of reification, which is always present; we may or may not also find the second fallacy, abstraction. Scientists underestimate the links which that particular structure has to the other structures, and conclude that on its own it can provide a practical approximation of the complex phenomenon. Once they isolate a simple structure, any number of mechanistic theories present themselves to explain that one reified phenomenon; but the original, complex phenomenon remains unexplained. Sometimes, the scientists extract *several* structures, hoping to explain the complex phenomenon by somehow combining their separate, mechanistic explanations.

We recognize a theory as mechanistic when it is deterministic, when it promises a complete and exact explanation, when it is based on reductionistic and atomistic concepts. It may use simple hierarchical structures directly, but it may also use a machine model, or a mathematical method, or diagrams, or software. Regardless of the concept or model employed, we already know that if it is deterministic it is also mechanistic, so there exists a simple hierarchical structure that could represent it. That structure may be large and unwieldy, but we don't have to create it, or even to visualize it. We can continue to study the theory in its given form (mathematics, diagrams, software, etc.), and keep in mind that it is logically equivalent to a simple hierarchical structure.

Scientism

1

Most attempts to explain complex phenomena mechanistically are found in the human sciences: in psychology, sociology, linguistics, economics, anthropology, history, politics, etc. And we must add programming and software use to this category, because, apart from their precise and deterministic aspects, these activities are performed in a social context; they are, therefore, subject to the same processes that make all human activities complex, and hence non-mechanistic, phenomena.

It is not hard to understand why these fields have engendered so many mechanistic theories. They deal with our most important concerns, and at the same time they are the most complex phenomena we know. So we are

anxious to discover models that can explain and predict, just like the models we use in physics or astronomy; and we forget that we are now dealing with indeterministic phenomena, which cannot be represented with exact models.

Human phenomena always involve interacting structures, starting with those that occur in one mind and ending with those that involve many minds in a society. Each one of us has many characteristics, roles, and preoccupations: we are workers, parents, friends, neighbours, citizens, spouses, investors, consumers, drivers, clients, patients; we have fears, wishes, beliefs, hopes, memories, plans; we use and own objects; we have bodies. The list of things that make us what we are is practically endless. But the important thing is that we are all of these things at the same time, all the time.

Each one of these attributes places us in a different structure; for example, a structure of people, events, activities, locations, or beliefs. We are the elements of many structures at the same time, so a human existence is a complex structure. We cannot even *identify* all these structures, much less separate them from one another. Even as we are engaged in one activity, or think of one subject, we are influenced by everything else that makes us what we are. At any instant, we are the repository of countless experiences, traditions, emotions, propensities, and pieces of knowledge. Our minds provide the interactions that make our individual lives complex structures, and each one of us provides the interactions that make society a complex structure. We are, each one of us, both the result and the cause of the many interacting structures that make up society.

It ought to be obvious, then, that it is impossible to study *any* aspect of individual or social life in isolation. Yet this is precisely what we do in the academic disciplines called psychology, sociology, linguistics, and so on. The disciplines are separated as if our behaviour, our social relations, or our use of language were separate phenomena. Various theories assume that I am either a man, or a programmer, or a speaker, or a consumer, or a hundred other things. But I am *all* these things *at the same time*. How could I be only one thing at a time? I am not running several unrelated programs simultaneously, like a multitasking computer. It is my mind and body, which are always there, that cause the endless interactions between these roles.

If we seek a theory that explains one aspect of human life, we must extract that aspect from the complex whole. We must treat it as a simple structure, and assume that its interactions with the other structures are weak enough to be ignored. But this is rarely true, so the approximations created by these theories are very poor. We already saw that language can be explained only through non-mechanistic theories – theories that take into account *all* the knowledge present in a mind. And the same is true of our other capabilities and acts, as individuals and as society.

2

Scientism is a derogatory term used to describe mechanistic delusions: the application of mechanistic concepts in fields where they cannot work, and especially in the human sciences. Scientism is the opposite of science; it is a dogmatic approach to research, embraced by those who are incapable of discovering useful theories. Mediocre scientists prefer to stress *methods* rather than *results* as the criterion by which the value of research is measured. Only exceptional people can make a real contribution, but almost anyone can follow methods. Through their dominating influence in academia and in society, these bureaucrats have shifted the definition of science to match their incompetence: the practice of science in many disciplines today means an interminable “research program” – activities that must obey the mechanistic principles, but need not produce any useful results.

Here are three views: “Scientism is the profoundly unscientific attempt to transfer uncritically the methodology of the physical sciences to the study of human action.”¹ Scientism describes “an attitude which is decidedly unscientific in the true sense of the word, since it involves a mechanical and uncritical application of habits of thought to fields different from those in which they have been formed. The scientific as distinguished from the scientific view is not an unprejudiced but a very prejudiced approach which, before it has considered its subject, claims to know what is the most appropriate way of investigating it.”² “The progress of modern science has been due to its rigorous confinement to the measurable aspects of elements of experience which are contained in a causal system. But science does not encompass nor does it profess to encompass all human experience. Science seeks to approximate the truth about the world only within the limitations of specific and rigorously defined contexts. No true scientist will claim more; no educated layman should expect more. Yet the vulgarization of science – scientism – has led many people, including not a few scientists who have lost sight of the philosophical foundations of their craft, to assert that science holds the key to *all* problems of human experience.”³

These views, expressed many years ago, show that we have been aware for a

¹ Murray N. Rothbard, “The Mantle of Science,” in *Scientism and Values*, eds. Helmut Schoeck and James W. Wiggins (Princeton, NJ: D. Van Nostrand, 1960), p. 159.

² F. A. Hayek, *The Counter-Revolution of Science: Studies on the Abuse of Reason*, 2nd ed. (Indianapolis: Liberty Fund, 1979), p. 24.

³ Robert Strausz-Hupé, “Social Science Versus the Obsession of ‘Scientism,’” in *Scientism and Values*, eds. Schoeck and Wiggins, p. 223.

long time of the limitations of mechanism. This has not stopped it from spreading, however. Although many thinkers have shown why mechanism cannot work in the human sciences, the mechanistic temptation is too great to resist. This is true because there are no easy alternatives, if we seek an exact theory.

For complex phenomena, the approximations provided by mechanistic theories are not close enough to be practical. It is highly improbable that we can discover a theory which explains and predicts with great accuracy individual or group behaviour. And if we ever do, it would be a *non-mechanistic* theory, because it would have to explain many of our capabilities and actions at the same time. Its model would be a complex structure, and hence impossible to represent with diagrams, mathematics, or software. Such a theory might evolve in the mind of one person (because minds can process complex structures), but that person would be unable to describe it to others with precision, as we explain mechanistic theories. This superior knowledge would manifest itself in the form of correct decisions and predictions, and would be interpreted as the result of intuition or personal experience.

In our mechanistic culture, however, such a theory would be rejected as “unscientific.” Researchers, therefore, do not even try to discover non-mechanistic explanations. They are trained to think only in mechanistic terms, to see the world as nothing but isolated hierarchical structures, so they can discover nothing better than mechanistic theories even when non-mechanistic explanations exist. We are caught in the mechanistic trap: we only permit ourselves to discover theories that cannot work.



If mechanism leads to scientism, scientism leads to utopianism – the belief that society can be greatly improved by implementing some rational or scientific plan. Thus, modern utopian ideas tend to parallel contemporary scientific knowledge. Their enthusiasm reflects the latest successes in the natural sciences, and the belief that similar successes are possible in the human sciences. Not surprisingly, utopian theories have been multiplying at an ever increasing rate since the Scientific Revolution. Utopian ideas never work, of course, and few were attempted on a large scale. Nevertheless, they have been an important factor in the evolution of social thinking.

The most striking characteristic of utopianism is that it always leads to totalitarianism; that is, a society where individual freedom is sacrificed in the name of an ideal, where an elite controls the political, social, and even personal life of every citizen. Central control is deemed necessary in order to maximize efficiency in industry, commerce, education, and all other public domains.

This efficiency, it is believed, will ultimately benefit all people by permitting the elite to create a perfect society.

The history of scientific utopianism in Western culture starts with Plato's ideal society in the *Republic*. Karl Popper,⁴ in his study of totalitarian philosophy, traces its roots to Plato's theory of forms or ideas. According to this theory, all objects, attributes, and processes in the real world derive from a set of pure and abstract ideas. The ideas depict the original, perfect forms, while the real things derived from them are deviations from these forms. The idea of a circle, for example, is a perfect circle. There can be only one perfect circle, and that is an abstract concept; the actual circles found in the world are only approximations.

Plato disliked the political turmoils and the drift toward democracy that were taking place in his time, favouring instead a stable and unchanging society – a society founded upon sound political principles. According to his theory of ideas, there must exist an idea of the perfect social system, and the reason why societies in the real world are so changeable and corrupt must be that they are deviations from that system. It should be possible, therefore, to deduce rationally what the perfect society ought to be like. This he does, and the result is the depiction of a totalitarian state.

Like all utopias, Plato's ideal society can only be realized if we disregard the interests of the individual. Society is to be ruled by an elite – a body of highly trained experts who, Plato feels, are alone qualified to make political decisions. Thus, using what seems to be irrefutable logic but is in fact scientism, Plato promoted totalitarianism as the best form of government. This escaped the notice of most interpreters until quite recently, when the similarities between Plato's imaginary state and the totalitarian ideas of our own time became evident.

While in Plato's utopia the rulers had to be highly trained in the sciences and in mathematics, Francis Bacon took this idea to its extreme: in his utopia, *The New Atlantis*, the rulers *are* scientists. There is no politics, as the chief preoccupation of the rulers is scientific research and technological development. All the inhabitants of his imaginary island are supremely happy, thus illustrating Bacon's belief that science and technology will solve all personal and social problems.

Bacon is known as the founder of modern science, but not because of any discoveries of his own. This honour is due to the fact that he was the first to specify and promote the new methods of empirical research – observation, experimentation, inductive logic – while also being an influential philosopher

⁴ Karl R. Popper, *The Open Society and Its Enemies*, vol. 1, *The Spell of Plato*, 5th ed. (Princeton, NJ: Princeton University Press, 1966).

and public figure. So, although most contemporary scientists were already using these methods in their work, it was Bacon's reputation that made them respectable. It is interesting, however, that for someone who stressed the importance of logical methods of inquiry, Bacon found it unnecessary to apply such methods to his vision of the ideal society. He takes it for granted that scientific and technological advances can be used to improve human lives and human societies just as we use them to improve things in the material world. He assumes – without trying to verify this assumption – that if each need and desire can be addressed separately by specialists, this is bound to create happy individuals and hence an ideal society. He gives no thought to the complex interactions that take place in a society, and how these interactions might affect the implementation of his vision.

In the modern era, the first attempt to implement a utopian idea took place in France during the revolution of 1789. This attempt resulted immediately in a totalitarian society, followed after a few years of failures by an even more oppressive regime – the Napoleonic dictatorship. A variety of mechanistic social theories were developed in Europe during the next one hundred years, eventually giving rise to the two greatest utopian concepts, and hence totalitarian systems, of our time: Communism and Nazism.

Albert Salomon and F. A. Hayek, among others, have studied the growth of social scientism in the nineteenth century, and its totalitarian tendencies.⁵ They trace its origins to Henri de Saint-Simon and Auguste Comte, the French thinkers credited with founding sociology. These thinkers were quite outspoken about their vision. The only way to have a perfect society, they said, is by designing it scientifically, from scratch. Then, an elite should control it, scientifically. Individual freedom is an unscientific concept, and its abolition must be the first step in any project whose goal is social progress. Thus, while seeking a system that benefits all citizens, these thinkers ended up advocating totalitarianism as the answer. This is the type of absurdity that scientific beliefs lead to: “The early sociologists made one fatal mistake: they placed their faith in the methods of natural science. In their fervent hopes for an intellectual revolution, they believed that knowledge about human beings was of the same sort and could attain the same precision as that of physics or biology. . . . But the application of scientific progress to rational transformation of society in the name of a humanitarian ideal ended in a clearly articulated vision of a totalitarian society.”⁶ We will return to this subject in “Totalitarian Democracy” in chapter 8.

⁵ Albert Salomon, *The Tyranny of Progress: Reflections on the Origins of Sociology* (New York: Noonday Press, 1955); Hayek, *Counter-Revolution of Science*.

⁶ Salomon, *Tyranny of Progress*, pp. 103–104.



We are interested in utopian thought because, if we understand how utopian social concepts are born from mechanistic social theories, we will be in a better position to understand how our mechanistic *software* theories are engendering today a new sort of utopia: the Age of Information. By exposing the similarities – the ignorance of the elites, the pseudoscientific nature of their theories, the use of propaganda instead of reasoned argument, the need for deception to cover up failures – we may be able to recognize the totalitarian tendencies of our software culture, and to avoid perhaps the mistakes of the past.

Scientific social theories, we recall, offer solutions by separating the complex phenomenon that is a human society into isolated aspects; in other words, by separating a complex structure into its constituent simple ones. Human beings are seen as members of political systems, or production systems, or educational systems, or communities, or families, or organizations. But human beings are elements in these structures, and in many others, *at the same time*. They are, in fact, the shared elements that cause these structures to interact. In our imagination, it is not difficult to find precise, mechanistic methods to improve any *one* aspect of society. It usually seems possible to improve that aspect simply by modifying or replacing the structure that represents it. When we try to implement these changes, however, we find that it is impossible to isolate that aspect from the others. What we may see in practice is that people are reluctant to adopt the changes because these changes affect negatively other aspects of their life, or that the other aspects affect the one modified so that the promised benefits do not materialize.

For example, we may have nothing in principle against a centrally planned economy. But in order to implement such an idea, the elite must decide which social institutions are important and which ones are not, must determine what goods are to be available and their prices, must control everyone's education and training, must restrict individual choices in selecting a career, and so on. Even if it were sincere when claiming that it wanted to control only one thing, the elite would end up controlling *all* aspects of life. The utopian system would become totalitarian.

Similarly, our software elites may be sincere when claiming that all they want is to improve the way we create and use software. But if everything we do depends directly or indirectly on computers, the only way to control software is by controlling *all* aspects of our existence. Thus, a society can become totalitarian through a utopian software ideology as easily as it can through a utopian economic, political, or religious ideology.

Their ideology notwithstanding, all totalitarian systems are alike. The elites are so confident in their plans that they cannot understand our doubts. How

important are individual rights compared to the vision of a perfect society? They attribute our resistance to ignorance, to our failure to appreciate the ultimate benefits of totalitarianism. They alone are knowledgeable enough in these matters, so it is their duty to persuade us to accept the new system, for our own good. And if this necessitates deception, propaganda, indoctrination, or even force, then they must carry out these unpleasant tasks, because achieving that dream is more important than the rights of the individual. In any case, these are only temporary measures, which will become unnecessary once we are all enlightened enough to accept the new system on our own accord.

3

The influence of the natural sciences on the social sciences took place in two stages.⁷ In the period encompassing most of the seventeenth century, social thinkers, impressed by the successes of the new mechanical and mathematical concepts, looked for ways to apply the same concepts in their own fields. After Newton's publication of the *Principia*, however, their expectations grew accordingly: they became convinced that a *universal* theory can be found, similar to Newton's theory of universal gravitation, which would explain all social and individual behaviour.

No useful theories based on mechanical models were ever discovered in the human sciences, but a brief survey of these attempts may help us later to recognize our *software* delusions. We immediately notice the use of mechanical terms to describe mental and social phenomena: "mass," "attraction," "motion," etc. Similarly, the software experts describe *programming* concepts with mechanical terms: applications are "assembled" from software "components," are "built" by software "engineers," and so forth. We recognize in both cases the same circular, fallacious logic: The experts *assume* that a given phenomenon can be reduced to mechanics, so they adopt mechanical notions *before* discovering a useful mechanistic theory. As a result, models and terms that have a precise meaning in mechanics are only metaphors in the other fields, where they are employed illegitimately. But now, these very metaphors are taken as evidence that the theory is valid.

Thus, John Locke and the other empiricist philosophers applied the model of contemporary mechanics to the working of the mind: "The mind was treated as if it were a box containing mental equivalents of the Newtonian

⁷ I. Bernard Cohen, *Interactions: Some Contacts between the Natural Sciences and the Social Sciences* (Cambridge, MA: MIT Press, 1994), pp. 101–102.

particles.”⁸ Everything we know, the empiricists claimed, originates with the simple sensations perceived by our bodies; these sensations result in simple ideas, or atoms of thought; and the complex ideas that constitute knowledge and reasoning are combinations of these simple ideas, in the same way that physical objects are combinations of some elementary particles.⁹ David Hartley, who introduced the concept of associations of ideas, explained them as purely mechanical processes: each type of sensation is transmitted as a different type of vibration through the nerves and causes a corresponding vibration in the cerebral material, which then becomes predisposed to vibrate in that way. This results in a simple idea, and when different vibrations occur simultaneously, these ideas become associated. Ultimately, “Hartley’s associationalist psychology led to a mechanical theory of creativity.... The creative function of the mind consisted of breaking down complex ideas into their component parts and rearranging them into new ideas. That is, the imagination functioned in a purely mechanical way – simply rearranging parts into new wholes.”¹⁰

George Berkeley tried to explain social relations by viewing them as mutual attraction between people; and David Hume claimed that the phenomenon of associations of ideas can be explained as a form of attraction. Both indicated Newtonian mechanics as the source of their inspiration, comparing their concepts of attraction to the gravitational attraction between physical masses.¹¹

Charles Fourier, one of the founders of socialism, created a utopian social theory based on Newtonian mechanics. (A few communities in France and in America actually attempted to implement this theory.) He “claimed to have discovered an equivalent of the gravitational law, one that applied to human nature and social behavior.”¹² He had a “calculus of attraction” and a “calculus of harmony,” and claimed that his “laws of social motion” were superior to Newton’s laws of physical motion.¹³

Many economists were impressed by the science of mechanics. Leon Walras, for example, published an article titled “Economics and Mechanics,” in which he “argued that identical differential equations appear in his analysis of economics and in two examples from mathematical physics: the equilibrium of a lever and the motion of planets according to gravitational celestial mechanics.”¹⁴ Vilfredo Pareto claimed that “pure economics is a sort of mechanics or akin to mechanics.”¹⁵ Pareto, as well as Fisher, drew up detailed tables showing many concepts from mechanics and their counterparts in

⁸ Isaiah Berlin, *The Age of Enlightenment: The 18th Century Philosophers* (New York: Mentor, 1956), p. 18.

⁹ Ibid.

¹⁰ David F. Channell, *The Vital Machine: A Study of Technology and Organic Life* (New York: Oxford University Press, 1991), p. 44.

¹¹ Cohen, *Interactions*, p. 19.

¹² Ibid., p. 20.

¹³ Ibid.

¹⁴ Ibid., p. 41.

¹⁵ Vilfredo Pareto, quoted *ibid.*

economics.¹⁶ J. E. Cairnes claimed that the principles of economics are identical in character to the physical principles of the laws of gravitation and motion.¹⁷ In the end, “with a sense of security coming from the use of equations homologous to those in physics, the new economics assumed the metaphor of rational mechanics.”¹⁸

Harvey’s discovery of the principles of blood circulation, occurring as it did when the mechanical philosophy was beginning to dominate science, contributed greatly to the view that the body works just like a machine. Few doubted that, like blood circulation, all physiological functions would ultimately prove to be nothing but mechanical operations: “The ‘new philosophy’ no longer accepted organic development as self-explanatory. Rather, it insisted on analysing all natural processes into fixed patterns of mechanical action and interaction. The bodies of animals, quite as much as inanimate objects, were to be regarded as configurations of material parts, moving and interacting like the pieces of a machine.”¹⁹ “At first the machine served only as an analogue for biological processes.... But as mechanical philosophy became successful as a method of explanation, people no longer saw the machine as simply an analogue for life – life became literally mechanical.”²⁰

Although Harvey himself recognized that most biological phenomena cannot be reduced to mechanics, his work was vulgarized by Descartes and the other mechanists. What emerged as a result was the school of mechanical biology called iatromechanics.²¹ Some scientists applied Newtonian mechanics directly, and attempted to explain physiological functions and diseases as the effect of forces they believed to act between organic particles.²² Most theories, however, were directed toward finding parallels between organisms and machines, on the assumption that live systems work on the same mechanical, hydraulic, or pneumatic principles as the machines of that period.

It was easy enough to find parallels to mechanical operations when superficially studying, say, the movement of limbs. But the iatromechanists attempted to explain *all* physiological functions as mechanical operations: digestion, metabolism, respiration, reproduction, sensation – no function lay beyond the power of their imagination, and they invented machine-like operations to explain them all. Although they believed that they saw evidence to confirm their theories, these theories were only speculations: “For the most part, iatromechanics was simply irrelevant to biology.... Beside the subtlety of

¹⁶ *Ibid.*, pp. 44–47.

¹⁷ *Ibid.*, p. 42.

¹⁸ *Ibid.*

¹⁹ Stephen Toulmin and June Goodfield, *The Architecture of Matter* (Chicago: University of Chicago Press, 1982), p. 168.

²⁰ Channell, *Vital Machine*, p. 30.

²¹ Richard S. Westfall, *The Construction of Modern Science: Mechanisms and Mechanics* (New York: Cambridge University Press, 1977), p. 94; Channell, *Vital Machine*, p. 36.

²² Channell, *Vital Machine*, p. 39.

biological processes, the 17th century mechanical philosophy was crudity itself.... In fact iatromechanics made no significant discovery whatever.”²³

The movement culminated with the work of Julien Offray de la Mettrie, *Man a Machine*, published in 1748. This philosophical treatise went beyond Descartes’s mechanistic doctrine by asserting that not just the human body but also the phenomena of mind and life are mechanical in character. Everything concerning human beings, therefore, can be explained by applying the principles of motion to organic matter.

The political theory of Thomas Hobbes was very advanced for his time, and influenced several generations of thinkers. Hobbes, however, was also interested in the natural sciences, and his political theory reflects this: “He was fully convinced that a science of politics or of human society must be similar to a natural science, based on two primary concepts: movement and matter or substance, in accordance with what was known as the ‘mechanical philosophy.’”²⁴ In *Leviathan*, Hobbes likened the state to a great animal, combining the mechanical theory of Galileo and the physiological theory of Harvey into a political theory. The state had been likened to a live body by earlier political thinkers, but Hobbes modified this concept into that of “a great animal machine, acting like an animal but composed of mechanical parts.”²⁵ He “used the new discoveries in physiology to transform the organismic concept of the body politic by giving it a mechanical basis in conformity with Descartes’s reductionistic philosophy. The political and social world of Hobbes is a hybrid kind of organic structure operating mechanically and conceived under the sign of Galileo, Descartes, and Harvey. His system of society was a collection of human beings acting as ‘mechanical systems of matter in motion.’”²⁶

Hobbes had great admiration for the deductive methods of science, as illustrated by the theorems of geometry, and believed that a perfect society could be attained if we found a way to apply these methods to social relations: “Were the nature of human actions as distinctly known, as the nature of quantity in geometrical figures,” all irrational and dishonest motives would disappear, and mankind would enjoy “an immortal peace.”²⁷ The method of science “was said by Hobbes to lead to predictive rules for a human science and so to produce a guide for obtaining predictable results in the domains of ethics or morals and of political action. In short, Hobbes envisioned a social science that would have some of the same qualities of exactness and of predictability as the physical sciences.”²⁸

²³ Westfall, *Modern Science*, p. 104.

²⁵ *Ibid.*

²⁶ *Ibid.*, p. 123.

²⁴ Cohen, *Interactions*, p. 120.

²⁷ Thomas Hobbes, quoted *ibid.*, p. 121.

²⁸ *Ibid.*

In France, too, in the decades preceding the Revolution, political thinkers held that human phenomena are similar to the phenomena studied by the exact sciences. So, they concluded, it should be possible to design a perfect society simply by emulating the methods employed in those sciences. The Revolution was, in effect, an attempt to implement such a society. Thus, Morelly claimed that a science of morality can be developed, “as simple and as self-evident in its axioms and consequences”²⁹ as a mathematical system. He assumed that there existed in nature an ideal, objective pattern of human affairs, which he perceived as “a social mechanism, a ‘marvellous automatic machine.’”³⁰ Similarly, Mably “strove for scientific certainty in social and human affairs. He believed that politics could develop from the most conjectural into a most exact science, once the recesses of the human heart and passions had been explored, and a scientific system of ethics defined.”³¹ And Condorcet was convinced that the events of the Revolution proved that a general method of investigation had been found, equally applicable to the exact sciences and the human sciences: “Once this instrument had been applied to morals and politics, a degree of certainty was given to those sciences little inferior to that which obtained in the natural sciences.”³²



The most flagrant manifestation of scientism is to be found, however, not in previous centuries but in our own time. The mechanistic fallacies that dominated the human sciences in the mid-twentieth century were denounced, in an exhaustive and scathing study, by the eminent sociologist Pitirim Sorokin.³³ Still, even though considered a classic, this study did not put an end to mechanistic theories.

In its simplest form, this scientism manifests itself in the use of jargon, neologisms, pompous terms, and platitudes. This practice, mockingly called by Sorokin “speech disorders,” includes the “blind transference of terms and formulas from the natural sciences into sociology and the related disciplines. The net result of the transference is a distortion of the precise meaning the terms have in the natural sciences and a contamination of the social sciences by terms that now become either meaningless or vague. Being incomprehensible, such terms impress the uninitiated as exact and ‘scientific.’”³⁴ With innumerable examples, Sorokin shows that any definitions and explanations based on such terms are meaningless, because the terms themselves are not

²⁹ J. L. Talmon, *The Origins of Totalitarian Democracy* (New York: Praeger, 1960), p. 17.

³⁰ *Ibid.*

³¹ *Ibid.*, p. 18.

³² *Ibid.*

³³ Pitirim A. Sorokin, *Fads and Foibles in Modern Sociology and Related Sciences* (Chicago: Henry Regnery, 1956).

³⁴ *Ibid.*, p. 22.

properly defined. To take just one case of “speech disorders” cited by him: “Every psychological activity may be ordered to a *two-dimensional plane (surface)* where organism and goal represent certain spatial regions within *the surface...*’ Psychological activity of all sorts will be ordered to a *path*, and may be said to represent *locomotion* in the psychological field.’ This psychological field is a ‘topological medium,’ with ‘fluidity,’ ‘cohesiveness,’ ‘permeability,’ ‘hodological space,’ etc.”³⁵

Another imitation of the exact sciences is the concept of operationalism. Operationalist methods (which consist of precise definitions, experiments, and measurements), while important in disciplines like physics and chemistry, are worthless in the human sciences: “As usual, without the necessary study of the real nature of operationalism, of its role in the progress of the natural sciences, of its limitations and doubtful elements, and forgetting the important role of pure intuition, deduction, and nonoperational induction in the progress of science and wisdom, our sociologists, psychologists, and anthropologists were converted into ardent operationalists and began *en masse* to apply operational method on their study of social, cultural, and mental phenomena. A sort of operational orgy rapidly spread throughout these disciplines.”³⁶

G. K. Zipf, for example, using operational preciseness, arrives at this fantastic definition of an organism: “A movable mathematical point in time-space, in reference to which matter-energy moves in such a way that a physical situation exists in which work is expended in order to preserve a physical system from a final gravitational and electromagnetic equilibrium with the rest of universe.”³⁷ S. C. Dodd, in his so-called system of operationally defined concepts for sociology, attempts to explain social change with terms borrowed from physics, but he does not actually define any operational procedures. “Instead,” Sorokin notes, “he simply takes the terms with their symbols (like T for time) from the physical sciences and concludes that by such a transference he has satisfactorily solved the problem of operationally defined concepts for sociology.”³⁸ So the resulting theory, despite the presence of equations similar to those we find in physics, is nonsensical: “However impressive this simplified transcription of physical concepts and their symbols looks, in application to ‘societal’ time, duration, change, acceleration, and force, these definitions are empty and useless. For *they do not give any real unit* for the measurement of social change or of its acceleration, velocity or force.”³⁹

T. Parsons and R. F. Bales claim that social action can be explained with such notions as the principle of “inertia,” the principle of “action and reaction,” and

³⁵ Ibid., p. 25, citing J. F. Brown and K. Lewin.

³⁷ G. K. Zipf, quoted *ibid.*, p. 30.

³⁶ Ibid., p. 32.

³⁹ Ibid.

³⁸ Ibid., p. 41.

the principle of “virtual displacements.” But, while valid in mechanics, in a social theory these principles are merely “distorted transcriptions” and “logical and empirical nonsense.”⁴⁰

Some of the authors of mechanistic social theories are not sociologists at all. For instance, P. W. Bridgman is a physicist, and J. Q. Stewart is an astrophysicist. Their avowed goal is to find “uniformities in social behaviour which can be expressed in mathematical forms more or less corresponding to the known patterns in physical science.”⁴¹ So they attempt to express such psychosocial concepts as people, activities, interactions, and desires in physicalist terms like social mass, social temperature, and social distance. Stewart “views the social universe as six-dimensional or made up of six ‘social quantities’ or ‘fundamental categories’: ‘distance, time, mass, temperature, electric charge, and number of molecules,’ whatever social interpretation is to be given to each of these ‘dimensions’ or ‘social quantities.’ We are told further that ‘this list [of six dimensions] makes social physics in its dimensional structure isomorphic with physical science,’ that is, ‘there is a complete and trustworthy analogy between two or more situations’ which entitles one ‘to transfer equations from physics to politics.’”⁴²

Some scientists propose theories that are based directly on the concepts of reductionism and atomism. They maintain that there exist social *atoms* – elementary social particles that are the smallest units of psychosocial phenomena. The social atoms are small groups of individuals (families, for instance), whose relations are simple and irreducible and can therefore be described with precision. (In a social atom, these scientists say, the relations between individuals are nothing but forces of attraction and repulsion, just like the forces operating within the atoms studied by physics.) We should be able, then, to discover an exact theory for the entire society simply by representing larger and larger groups as neat structures of social atoms. But Sorokin shows that these theories are mistaken: members of a society interact in complex ways, and no group can be isolated from the others and described with precision.⁴³

The modern mechanistic theories, says Sorokin, are not only worthless but also unoriginal. These scientists seem to be unaware of the mechanistic delusions of the past, and are merely repeating the work previously “performed by a legion of social and psychological scribes hoping to establish a new ‘social physics,’ ‘social mechanics,’ ‘social geometry,’ or ‘social energetics.’ ... Contrary to their claims to being revolutionary, contemporary ‘social physicists,’ econometrists, psychometrists, sociometrists, and ethicometrists are

⁴⁰ Ibid., pp. 246–247.

⁴¹ J. Q. Stewart, quoted *ibid.*, p. 188.

⁴² Ibid., p. 189, citing J. Q. Stewart (brackets in the original).

⁴³ Ibid., ch. 10.

merely continuing centuries-old operations.”⁴⁴ After analyzing hundreds of claims, Sorokin concludes that “the recent period has not produced anything remarkable in the field of general systems of sociology and psychology.”⁴⁵



This brief survey of mechanistic thinking will suffice for now. It is important to recognize, though, that similar idiocies are being pursued even today, by scientists working in famous universities. And we should perhaps ponder over the prospects of a society that not only tolerates this intellectual corruption, but holds the scientists and their institutions in high esteem. If you feel this judgment is unfair, you will see in “Popper’s Principles of Demarcation” (in chapter 3) that it is possible to determine whether a given activity constitutes true scientific research, or whether it is pseudoscientific. Accordingly, mechanistic theories like these can be shown to be worthless pursuits.

We will continue our study of scientism in the next three chapters, where we will examine the mechanistic delusions that concern us the most. By learning to recognize the common mechanistic foundation of these theories, we will be in a better position to understand our *software* delusions, which also spring from mechanistic thinking. The study of scientism, thus, can help us to see our software ideology in the broader context that is our mechanistic culture: given this tradition, the emergence of software mechanism was inevitable.

⁴⁴ Ibid., p. 110.

⁴⁵ Ibid., p. 310.

