

SOFTWARE AND MIND

Andrei Sorin

EXTRACT

Chapter 3: *Pseudoscience*

**This extract includes the book's front matter
and chapter 3.**

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 discusses the concept of pseudoscience, the principles of demarcation between science and pseudoscience developed by Karl Popper, and the value of these principles in studying the pseudoscientific nature of our mechanistic culture.

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.

CHAPTER 3

Pseudoscience

The mechanistic view of mind we studied in the previous chapter is only one of the many mechanistic delusions being pursued in universities today under the cloak of science. In the present chapter, I propose to study some of the other delusions, and to show that they all share a set of obvious characteristics.

This study has a dual purpose. First, we will expose the intellectual corruption of the academics – a corruption inevitable when mechanism changes from a mere hypothesis into a principle of faith. Second, we will establish methods for determining whether a given theory, or discipline, or research program, represents a legitimate scientific activity or is nothing more than a system of belief. In addition, this study will help us later, when we examine the greatest mechanistic delusions of all time – our software theories. For, software mechanism has grown out of the mechanistic culture that pervades the academic world.

Even more harmful than the promotion of pseudoscientific theories are the political consequences of this mechanistic culture. If we believe that complex phenomena of mind and society can be modeled with exact theories, we are bound to believe also the utopian promises of totalitarianism. (We will study this trend in chapter 8.) Thus, even though *failing* as scientific theories, the mechanistic notions promoted in universities are helping various elites – the software elite, in particular – to implement totalitarian ideologies.

The Problem of Pseudoscience

1

A pseudoscience is a system of belief that masquerades as scientific theory. The list of pseudosciences, ancient and modern, is practically endless: astrology is founded on the belief that the heavenly bodies influence human affairs on earth; phrenology claims that we can determine various personality traits from the shape of a person's skull; graphology claims that we can determine traits from a person's handwriting; dowsing maintains that it is possible to discover underground water just by walking over an area; alchemy holds that it is possible to transmute base metals into gold. Other pseudosciences are based on the belief in psychic phenomena, visits from aliens, faith healing, prophecy, magical objects, and so on.

Astrology has been with us for five thousand years, but most pseudosciences lose their popularity over time and are replaced by new ones. The continuing appeal of pseudoscience rests on its promise of simple solutions to difficult problems, as opposed to the relatively modest claims made by science. Widespread education has not eradicated what seems to be a basic human need – our craving for supernatural powers – and it has been noted that pseudosciences, superstitions, and the belief in the paranormal are actually on the rise throughout the modern world.¹

A distinguishing characteristic of pseudoscience is the acceptance of a hypothesis as unquestionable truth, and the refusal to review it later in the light of falsifying evidence. Whereas serious researchers insist on careful and objective tests of validity for their theories, pseudoscientific theories depend on the enthusiasm of the practitioners and the credulity of their followers. When subjected to controlled experiments, the success rate of these theories is usually revealed to be no better than chance. Pseudoscientific theories do not work, but believers interpret the chance successes as evidence of their truth, and belittle the significance of the failures. It is important to note that the practitioners' sincerity is often above suspicion; it is precisely their *belief* that prevents them from recognizing the falsity of their theories. But because there are no serious validity tests, pseudosciences also attract many charlatans – practitioners who knowingly deceive the public.

Despite their variety, the traditional pseudosciences have been addressing the same concerns since ancient times: our fears and desires, our longing for omnipotence and immortality. But today the mechanistic delusions are

¹ See, for example, Paul Kurtz, *The Transcendental Temptation: A Critique of Religion and the Paranormal* (Buffalo, NY: Prometheus Books, 1991).

fostering a new kind of pseudosciences: various academic pursuits that are part of modern disciplines and spheres of knowledge. And they are also fostering a new kind of pseudoscientists: researchers, professors, and theorists working in universities and other institutions. While these academic pursuits resemble scientific research, they belong to the pseudoscientific tradition insofar as they too are founded on a hypothesis that is taken as unquestionable truth. The hypothesis is that all phenomena can be explained with the mechanistic principles of reductionism and atomism. Although this belief is different from the beliefs upon which the traditional pseudosciences are founded, the ensuing pursuits acquire a similar character: they become systems of belief that masquerade as scientific theories. Thus, I call these pursuits *the new pseudosciences*. The new pseudosciences belong to the class of theories we examined in chapter 1 under scientism.

Like the traditional ones, the new pseudosciences do not work. Also like the traditional ones, blatant falsifications leave their supporters unperturbed. Instead of recognizing falsifications as a refutation of their theory, pseudoscientists think their task is to *defend* it; so they resort to various stratagems to make the theory appear successful despite the falsifications. Their work, thus, while resembling scientific research, is in reality a series of attempts to save from refutation an invalid theory.



We saw in chapter 1 how mechanistic delusions lead to futile pursuits (see pp. 104–106). If the phenomenon in question can only be represented with a complex structure – if, in other words, it cannot be usefully approximated by separating it into simpler, independent phenomena – the only way to explain it is by studying it as a whole. This is a difficult, often impossible, task. The researchers believe that a simple structure – in the form of a mechanistic theory, or model – can represent the phenomenon accurately enough to act as explanation. So they extract one of the simpler phenomena from the complex whole, hoping that a mechanistic model based on it alone will provide a good approximation of the whole phenomenon. They are committing the fallacy of reification, but they see this act as a legitimate method, sanctioned by science.

Science sanctions this method only for *mechanistic* phenomena. The researchers cannot know in advance whether their subject is indeed mechanistic, so the possibility of explaining the complex phenomenon by isolating the simpler phenomena that make it up is only an assumption. To validate this assumption, they must arrive at a successful explanation of the original phenomenon; specifically, they must discover a mechanistic approximation

that is close enough to be useful. But even when they find explanations for the isolated phenomena, the researchers fail to explain the original, complex phenomenon. We know, of course, why: the complex phenomenon includes the *interactions* between structures, and these interactions were lost when they separated the structures. They mistakenly assumed that the interactions are weak enough to be ignored, so the model based on reified structures does not represent the actual phenomenon accurately enough.

In their work, these researchers may be following the strictest methods. In their study of the isolated structures, their theories and procedures may be faultless. Thus, their activities may be indistinguishable from those of real scientists. The more complex the problem, the more opportunities there are to separate it into simpler problems, then to separate these into even simpler ones, and so on.

It is obvious, then, why the mechanists perceive these activities as important work. At any point in time, what they are doing resembles true research – the kind of work that in the exact sciences brings about great discoveries. Consequently, solving one of the isolated problems is seen as progress, as a contribution to the solution of the original problem. Besides, the theory does work in certain cases. It is in the nature of poor approximations to work in some cases and not in others, but the mechanists interpret the odd successes as evidence that their ideas are valid.

At this stage, they have forgotten that the entire project is based on the *assumption* that the original phenomenon can be explained mechanistically. The assumption is wrong, so all these activities – no matter how rational and scientific they may appear when judged *individually*, and no matter how successfully they may solve *isolated* problems – constitute a delusion. Not surprisingly, no theory that explains the original phenomenon is ever found. Modern mechanistic pseudosciences last several years, or several decades, and then they are quietly abandoned.

What is especially striking in pseudosciences, thus, is to see people engaged in activities that are entirely logical *individually*, even while the body of activities as a whole constitutes a delusion. All it takes is one wrong assumption; and if this assumption is never questioned, the research is nonsensical no matter how rational are the individual activities.

By its very nature, therefore, the mechanistic assumption engenders pseudosciences: If we assume that a non-mechanistic phenomenon can be explained by breaking it down into mechanistic ones, we will end up studying the latter. So, like real scientists, we will be engaged at all times in the exact work associated with mechanistic phenomena. We will be pursuing a delusion, but this will not be evident from the *individual* activities. The only way to recognize the delusion is by questioning the mechanistic assumption itself.



If it is so easy to fall prey to mechanistic delusions, how can we differentiate between those scientists engaged in important research and those who pursue hopeless, pseudoscientific ideas? Clearly, if we agree that science means simply the pursuit of mechanistic theories, regardless of whether they work or not, it is no longer possible to distinguish true scientists from crackpots and charlatans.

Note that it is not the *failure* of these theories that must concern us. Ambitious or revolutionary ideas often prove to be mistaken, so the risk that a theory may eventually fail should not prevent us from pursuing it. What we must question, rather, is whether the pursuit of a theory should be considered science simply because the theory is mechanistic. Science ought to mean the pursuit of *sound* theories: mechanistic ones for mechanistic phenomena, and non-mechanistic ones for complex phenomena.

Is there a way to avoid this enormous waste of resources – and, worse, its consequences? For, if we take the *software* theories as an indication of where this degradation can lead, the consequences are the destruction of knowledge, a return to the irrationality of the Dark Ages, and a totalitarian society. Once we recognize that software phenomena are non-mechanistic, any research program based on mechanistic software notions looks absurd, no different from the research program of the alchemists or the astrologers. It is only through the mechanistic hypothesis – namely, the assumption that any phenomenon can have a mechanistic model – that the software theories can be said to belong in the domain of science, rather than pseudoscience.

Thus, we have reached perhaps a critical point in history, where there is an urgent need to revise our conception of science. If the software delusions are an indication, the survival of our civilization may well depend on our decision whether or not to retain mechanism as an article of scientific faith.

2

Since mechanistic delusions undermine logical thinking in the same way that other delusions did in the past, the problem we are facing is a problem that has preoccupied philosophers for centuries: In our quest for new knowledge, how can we avoid irrational thinking, fallacious arguments, and unsound judgment? If what we discover is really new, how can we know whether it is true? For, the only way to be absolutely sure that something is true is by proving it on the basis of *previous* knowledge – knowledge whose truth is established. But then, a successful proof will also indicate that it depends entirely on facts

we knew before, leading to the conclusion that it is not really new. It seems, therefore, that we can gain new knowledge only if we do not also expect to be certain of its truth. This inference is very disturbing, as it suggests that the advance of knowledge depends entirely on something rather doubtful: the human capacity for faultless reasoning.

Many methods have been suggested for improving our thinking habits – methods ranging from rules of common sense to procedures of formal logic. In the seventeenth century, for instance, Francis Bacon, who stressed the importance of experimentation and logical thinking in scientific research, described four categories of reasoning errors (which he called “idols of the mind”). And in the nineteenth century, John Stuart Mill popularized a set of methods that can be used in any experimental inquiry to check the validity of hypotheses and to avoid drawing mistaken conclusions.

To the traditional principles we must add a new one if we want to guard against *mechanistic* delusions: Before attempting to explain a phenomenon by separating it into several independent phenomena, we must first prove that the interactions between these phenomena can be ignored. In other words, we must determine that the original phenomenon can indeed be modeled with simple structures. Since only mechanistic phenomena lend themselves to this treatment, if we commence our project by isolating structures we merely beg the question: we start by assuming the very fact that needs to be determined – the mechanistic nature of the phenomenon. Thus, if the project is to be considered science and not speculation, we must start by proving that the links *between* structures are weak relative to the links *within* structures; specifically, we must prove that they are weak enough to be ignored. And if such a proof is impossible, the phenomenon must be deemed non-mechanistic. Any search for a mechanistic theory is then known in advance to be futile, so it cannot be considered a serious scientific activity.

In particular, most phenomena involving human minds and societies consist of interacting structures, and weak links between these structures are the exception. Scientists isolate these structures precisely because they want to avoid the complexity generated by their interactions. They fail to see, though, that once they eliminate the interactions they are no longer studying the original phenomena. So we must not be surprised when, years later, they are still searching for a useful model. But we must remember that, by applying a simple logical principle, they could have avoided this futile work.



No system has been found that can guarantee sound reasoning while also permitting creativity, innovation, and discovery. The problem of reconciling

these conflicting ideals remains a difficult one. Descartes believed that the “geometrical method” is the answer: if we treat all knowledge as simple hierarchical structures, we will discover, without ever falling into error, everything the human mind can comprehend. Only in a world limited to mechanistic phenomena, however, could such a naive method work. And it is precisely his legacy – the belief that non-mechanistic phenomena, too, can be explained with the geometrical method – that engenders the new pseudosciences.

This problem has evolved into what is known today as the problem of *demarcation*: how to differentiate between scientific and pseudoscientific theories. The best-known and most successful principles of demarcation are those developed by Karl Popper. These principles can be used to assess, not only formal theories in the traditional disciplines, but any concepts, statements, and claims. In the domain of software, particularly, we can use them to assess notions like structured programming, the relational database model, and software engineering in general. These notions are in effect empirical theories, insofar as they make certain claims concerning the benefits of various methods or aids. The principles of demarcation will help us to determine whether these theories express important software concepts, or whether they are pseudosciences.

Popper's Principles of Demarcation

1

Sir Karl Popper, generally recognized as the greatest philosopher of science of the twentieth century, created a philosophy of knowledge and progress that can be applied consistently in all human affairs. It is useful for scientific theories as well as social and political ideas, for difficult decisions as well as common, everyday puzzles.¹

Popper held that it is impossible, in the empirical sciences, to *prove* a theory; so we can never be sure that our knowledge is correct or complete. The only way to advance, therefore, is through a process of trial and error, by learning from our mistakes: we must treat all ideas and theories as *tentative*

¹ See, in particular, these books by Karl Popper: *Conjectures and Refutations: The Growth of Scientific Knowledge*, 5th ed. (London: Routledge, 1989); *The Logic of Scientific Discovery* (London: Routledge, 1992); *Realism and the Aim of Science* (London: Routledge, 1985).

solutions, as mere conjectures, and we must never cease to doubt them. It is our responsibility, in fact, to attempt to refute our own theories – by subjecting them to severe tests. And we must always try to find better ones. In this way, our theories will keep improving, and we will get nearer and nearer to the truth. But, because the world is so complex, this process can never end. Indeed, even if one day we do arrive at the truth, we will have no way of knowing that we did.

Theories turn into worthless pursuits when their supporters choose to ignore the falsifying evidence. Unlike true scientists – who seek the truth and know that their theories, even when apparently successful, may be mistaken – pseudoscientists believe their task is simply to defend their theories against criticism.

Popper considered demarcation to be “the central problem of the theory of knowledge.”² It must be noted that he sought to distinguish the empirical sciences not only from pseudosciences, but also from metaphysics and purely logical theories. He recognized the value of these other types of knowledge; but, he said, they are different from the empirical sciences. For example, some theories considered scientific today originated in antiquity as pseudosciences, so even as pseudosciences they must have been useful; and purely logical systems like mathematics, while not part of the real world but our invention, can provide invaluable models (by *approximating* the real world). Any theory, thus, can be useful. But the theories of empirical science occupy a special position, because they alone permit us to develop knowledge that matches reality. So, if we want to improve our knowledge of the world, we must have a way of determining whether a given theory belongs or not to the domain of empirical science.

It may seem odd to place the rigorous theories of pure mathematics in the same category as pseudosciences. These theories *are* alike, though, when viewed from the perspective of empirical science; that is, when judged by their ability to represent the world. The mechanistic *software* theories provide a nice illustration of this affinity. The structured programming theory, for instance, and the relational database theory, are founded upon mathematical principles. But these principles reflect only minor and isolated aspects of the phenomenon of software development, not whole programming projects. In their pure form, therefore, these theories are useless for creating serious applications, because they do not approximate closely enough the actual software phenomena. They became practical (as we will see in chapter 7) only after renouncing their exact, mathematical principles and replacing them with some vague, informal ones. And this degradation is one of the distinguishing characteristics of

² Karl R. Popper, *The Logic of Scientific Discovery* (London: Routledge, 1992), p. 34.

pseudoscience: the experts continue to promote their theory on the basis of its exactness, even while annulling, one by one, its exact principles.

Mechanistic software theories, thus, can exist only as purely logical systems and as pseudosciences; and in either form they cannot be part of empirical science. Empiricism stipulates that theories be accepted or rejected through actual tests, through observation and experiment. As logical systems, the mechanistic software theories were tested in the real world, and failed; and in their modified form, as pseudosciences, these theories offer no exact principles to begin with, so they cannot be tested.

2

Popper's interest in a criterion of demarcation started in his youth, when he "became suspicious of various psychological and political theories which claimed the status of empirical sciences, especially Freud's 'psychoanalysis,' Adler's 'individual psychology,' and Marx's 'materialist interpretation of history.'"³ Popper was struck by the *ease* with which one could find confirming evidence for these theories, despite their dubiousness. A Marxist could find evidence of the class struggle in every event and every news item, and also in the *absence* of certain events or news items. A Freudian or Adlerian psychoanalyst could find confirming evidence of Freud's or Adler's psychological theories, respectively, in every act performed by every person; and had a person acted differently, that behaviour too could have been explained by the same theory. Any event seemed to fit quite naturally within these theories. In fact, one could not even *imagine* an event that would have contradicted them.

While these dubious theories were so easily *verified*, Popper was impressed by how easy it was to *falsify* a true scientific theory. Einstein, for example, boldly predicted several events from his theory of relativity, and declared that if they did not occur as stated he would simply consider the theory refuted.

Popper realized that it was precisely the *ease* with which a theory can be confirmed that *reduces* its scientific value, and this led him to his criterion of demarcation: "But were these theories testable?... What conceivable event would falsify them in the eyes of their adherents? Was not every conceivable event a 'verification'? It was precisely this fact – that they always fitted, that they were always 'verified' – which impressed their adherents. It began to dawn on me that this apparent strength was in fact a weakness, and that all these 'verifications' were too cheap to count as arguments.... The *method of looking for verifications* seemed to me unsound – indeed, it seemed to me to be the

³ Karl R. Popper, *Realism and the Aim of Science* (London: Routledge, 1985), p. 162.

typical method of a pseudoscience. I realized the need for distinguishing this method as clearly as possible from that other method – the method of testing a theory as severely as we can – that is, the method of criticism, the *method of looking for falsifying instances*.⁴

Several years later, Popper recognized that the problem of demarcation is closely related to the classical problem of induction, and that the two had to be considered together.⁵ The problem of induction is this: When we develop a theory in the empirical sciences, we draw general conclusions from a limited number of observations and experiments; we reason from singular facts to general statements; we believe that we can explain an infinite number of situations that have yet to occur, from the study of a finite number of situations that we observed in the past. This concept – induction – is indispensable in science, for we could have no theories without it. Logically, however, induction is invalid, because there is no justification for deriving general laws from the observation of unique events. The only way to practise science, therefore, is by trusting the principle of induction even as we know that it is invalid.

But there can be no doubt that induction *does* work: our knowledge *has* been increasing, and this shows that we *can* draw valid conclusions from past events, and we *can* have useful theories. We accept induction, therefore, simply because it works; and it works because there are regularities in the world: some future events *will* be similar to past ones, so it is possible to discover theories and to make predictions, especially if we are content with approximations.

Unfortunately, this expectation of regularities also tempts us to see patterns where there are none, leading us to fallacious thinking and irrational behaviour. Pseudosciences and superstitions are theories that predict future events from current knowledge, just like the theories of empirical science. For example, if we noticed once the position of the planets while a happy event took place, we will plan our activities based on their future position; and if we noticed once a black cat while a tragic event took place, we will avoid black cats in the future. With pseudosciences and superstitions, thus, we also use induction; we also draw general conclusions from the observation of a few events; so we also reason from particular facts to general statements. The only difference from science seems to be that our observations are less careful, so our conclusions are less accurate and our predictions less successful.

The belief in induction is closely related to the belief in causality. We must accept both principles in order to develop theories, and both stem from the way our mind works: we expect to find regularities in our environment. When

⁴ Ibid., pp. 162–163.

⁵ For Popper's views on induction, see *ibid.*, ch. I, and his *Objective Knowledge: An Evolutionary Approach*, rev. ed. (Oxford: Oxford University Press, 1979), ch. 1.

an event occurs simultaneously with another, or shortly thereafter, we tend to conclude that they must be related, that one caused the other or perhaps a third one caused both. This belief is reinforced by the belief in induction, when we observe a repetition of that pattern.

As with induction, though, no matter how often we notice the pattern, we have no logical grounds to conclude that there is a causal relation between the two events. We feel that such a relation is likely, of course, just as we feel that an event which occurred frequently in the past is likely to recur in the future. But these are strictly subjective notions, which spring from our habits of mind, from our natural tendency to expect regularities. According to the theory of probability, if we observed only a finite number of events and there are an infinite number of future events, the probability of predicting anything about those future events from the past ones is the first number divided by the sum of the two, which is practically zero.

Causality and induction, then, are hardly the solid and objective foundation we would like to have for our empirical sciences. It is true that science, unlike pseudoscience and superstitions, demands more observations before concluding that one event causes another; and it is true that scientific theories are more than just our expectation to see in the future a repetition of past events. Nevertheless, it is disturbing that our scientific knowledge has the same foundation as our superstitions: our habits of mind, our inclination to expect regularities, perhaps a propensity resulting from the evolution of the brain.



If you think these problems ought to concern only philosophers, remember the sad story of the chicken that believed in causality and induction. The chicken noticed, day after day, that the farmer sheltered it, fed it, and watched its health. After observing this pattern for many days, the chicken felt justified to conclude that the farmer's acts were motivated by love, and that it would enjoy the same comfort in the future. Soon after, though, the farmer killed the chicken and ate it – which had been his intent, of course, from the start.

What philosophers are trying to determine is whether, from the information available to it, the chicken could have known the truth.⁶ Or, rather, they are trying to determine whether *we*, from our current knowledge, can arrive at the truth. For, at any given time, we are in a position not very different from that of the chicken: we must make decisions about *future* events by using the doubtful theories we developed from the observation of relatively few *past* events. And when, recognizing the limitations of our personal knowledge, we

⁶ It is Bertrand Russell who first noted the chicken's quandary.

listen to scientists and experts, to corporations and universities, to governments and media, all we do is trust the doubtful theories that *others* developed from those few past events.

For example, when we accept the programming methods concocted by software theorists because they seem to work with some simple textbook examples, or when we judge the value of a software system from a few “success stories” or “case studies,” we are using in effect a few past events to make decisions about the future. But how can we be sure that we are not making the same mistake as the chicken?

So, if the problem of demarcation is how to distinguish our scientific from our pseudoscientific theories, the problem of induction is that *all* theories are logically unjustifiable, so there is no real difference between the scientific and the pseudoscientific ones in any case.

The problem of induction and its disturbing implications were first studied by David Hume, who resigned himself to complete skepticism. His conclusions had a profound influence on the development of Western thought, as they cast doubt on the possibility of rationality and objective knowledge: “The growth of unreason throughout the nineteenth century and what has passed of the twentieth is a natural sequel to Hume’s destruction of empiricism.... It is therefore important to discover whether there is any answer to Hume within the framework of a philosophy that is wholly or mainly empirical. If not, there is no intellectual difference between sanity and insanity.... This is a desperate point of view, and it must be hoped that there is some way of escaping from it.”⁷



Popper found a solution to Hume’s problem of induction, and to the skepticism engendered by it, through his solution to the problem of demarcation: “If, as I have suggested, the problem of induction is only an instance or facet of the problem of demarcation, then the solution to the problem of demarcation must provide us with a solution to the problem of induction.”⁸ He agrees that induction and past confirmations are insufficient to prove a theory; but he does not agree with the conclusion drawn by the earlier philosophers – namely, that this limitation will forever prevent us from distinguishing between our rational theories and our delusions.

What Popper proposes is to combine the methods of induction, which are indispensable for discovering new theories but cannot prove them, with the

⁷ Bertrand Russell, *A History of Western Philosophy* (New York: Simon and Schuster, 1972), p. 673.

⁸ Karl R. Popper, *Conjectures and Refutations: The Growth of Scientific Knowledge*, 5th ed. (London: Routledge, 1989), p. 54.

methods of *deduction*, which cannot create new knowledge but *can* prove statements. Deduction allows us to prove the validity of a statement by showing that it can be derived logically from other statements, which are known to be valid. Mathematical and logic systems, for example, are based on deduction: a conclusion is derived by combining premises; a new theorem is demonstrated by combining previous, simpler theorems. With strict deduction, there can be no knowledge in a new statement that is not already contained in the original ones (this is what guarantees the validity of the new statement). But, even though they do not create new knowledge, the methods of deductive logic are still important, because the new statements may express the same knowledge more clearly, more economically, and more usefully.⁹

Popper was impressed by the *asymmetry* between trying to *prove* a theory and trying to *refute* it. A theory is a universal statement that makes a claim about a large, perhaps infinite, number of events. Consequently, any number of confirmations are insufficient to prove its validity. At the same time, just one event that *contradicts* the theory is sufficient to *refute* it. Imagine, for instance, that we wanted to verify the universal statement “all swans are white” (one of Popper’s favourite examples). No matter how many white swans we observe, these confirmations would not verify the statement, for we could never be sure that we saw all the swans in the world; but observing just one black swan would suffice to *refute* the statement.

This is how Popper explains his idea: “My proposal is based upon an *asymmetry* between verifiability and falsifiability; an asymmetry which results from the logical form of universal statements. For these are never derivable from singular statements, but can be contradicted by singular statements. Consequently it is possible by means of purely deductive inferences (with the help of the *modus tollens* of classical logic) to argue from the truth of singular statements to the falsity of universal statements.”¹⁰

Modus tollens states that, if we know that whenever p is true q is also true, then if q is found to be false we must conclude that p is false. So what Popper says is this: if p stands for any one of the assertions that make up a theory, and q stands for any one of the conclusions derived from this theory, then just one instance of q being false will refute the theory.¹¹ In other words, while no number of “ q is true” claims that are true suffices to *prove* the theory, just one “ q is false” claim that is true suffices to *refute* it.

⁹ The induction discussed here must not be confused with the method known as *mathematical induction*, which employs in fact deduction.

¹⁰ Popper, *Scientific Discovery*, p. 41.

¹¹ *Ibid.*, p. 76.



The first thing we learn from Popper's discovery is how absurd is the popular belief that we must *verify* our theories, that we must search for *confirming* evidence. For, no matter how many confirmations we find, these efforts can prove nothing. Rather than attempting to show that a theory is valid, we must attempt to show that it is *invalid*; and the theory will be accepted as long as we *fail* in these attempts. It will be accepted, not because we proved its truth (which is impossible), but because we failed to prove its falsity.

Thus, if we sincerely attempt to refute our theories, if we agree to accept only those that pass the most severe tests we can design, our knowledge at any point in time is guaranteed to be as close to the truth as we can get. This, says Popper, is all we can hope to achieve: "Assume that we have deliberately made it our task to live in this unknown world of ours; to adjust ourselves to it as well as we can; to take advantage of the opportunities we can find in it; and to explain it, *if possible* (we need not assume that it is), and as far as possible, with the help of laws and explanatory theories. *If we have made this our task, then there is no more rational procedure than the method of trial and error – of conjecture and refutation*: of boldly proposing theories; of trying our best to show that these are erroneous; and of accepting them tentatively if our critical efforts are unsuccessful."¹²

With this method we combine, in effect, the benefits of induction and deduction. In our search for new theories, we can now use induction as often as we want. We need no longer worry about our habits of mind – about our inclination to expect regularities. Ideas revealed to us in our dreams are as good as those discovered through formal research methods. We can use our imagination and creativity freely, and we can propose theories that are as bold and original as we like. We can do all this because we need no longer fear that our thought patterns may be wrong, or that our conclusions may be mistaken. The discovery of a theory is now only the first stage. The theory is accepted provisionally, and it is in the next stage that the most important work is done: attempting to refute the theory by subjecting it to severe tests.

If we allowed the *uncertainty of induction* in order to *discover* the theory, we rely on the *certainty of deduction* in order to *refute* it. We benefit from deductive logic in two ways. First, as noted earlier, in the knowledge that the failure to pass even one test will prove that the theory is invalid. Second, we must use deductive methods – formal logic, mathematics, established theories, controlled experiments – in the tests themselves. It is pointless to devote any effort and to insist on deductive methods for tests that *verify* the theory; for, no

¹² Popper, *Conjectures and Refutations*, p. 51.

matter how scrupulous these tests are, each confirmation of the theory does not increase significantly the likelihood of its validity (since there will always remain an infinite number of unverified instances). Instead, we must devote this deductive effort to tests that try to *falsify* the theory. Logically, we can learn little or nothing from any number of instances that confirm it, but we can learn a great deal from just one instance that falsifies it.

Popper's solution, thus, has rescued the principle of empiricism – the requirement that theories be accepted or rejected on the basis of observations and experiments – from the destructive consequences of induction. All we must do is replace the principle of *accepting* a theory on the basis of *confirming* evidence, with the principle of *rejecting* the theory on the basis of *refuting* evidence. Empiricism “can be fully preserved, since the fate of a theory, its acceptance or rejection, is decided by observation and experiment – by the result of tests. So long as a theory stands up to the severest tests we can design, it is accepted; if it does not, it is rejected. But it is never inferred, in any sense, from the empirical evidence. There is neither a psychological nor a logical induction. *Only the falsity of the theory can be inferred from empirical evidence, and this inference is a purely deductive one.*”¹³

What Popper's solution amounts to, in essence, is a trade. We agree to give up the dream of knowing with certainty whether a theory is true or false; in return, we save the ideals of empiricism, the possibility to distinguish rationality from irrationality, and the hope for intellectual progress.

3

If the correct way to judge theories is by subjecting them to tests that try to falsify them, it follows that we cannot even consider theories that do not lend themselves to tests and falsification. This quality, then, is the criterion of demarcation that Popper was seeking: “Not the *verifiability* but the *falsifiability* of a system is to be taken as a criterion of demarcation.... I shall require that its logical form shall be such that it can be singled out, by means of empirical tests, in a negative sense; *it must be possible for an empirical scientific system to be refuted by experience.*”¹⁴

Most people think that to test a theory means to show that it works, so they choose for their tests situations that confirm the theory. But such tests are worthless: “It is easy to obtain confirmations, or verifications, for nearly every theory – if we look for confirmations.”¹⁵ The criterion of demarcation

¹³ Ibid., p. 54.

¹⁴ Popper, *Scientific Discovery*, pp. 40–41.

¹⁵ Popper, *Conjectures and Refutations*, p. 36.

prescribes the opposite; namely, for a theory to be included in the domain of empirical science, there must exist tests that, if successful, would *falsify* it. Thus, scientific theories are falsifiable; theories that are unfalsifiable are pseudoscientific.

It is important to understand the difference between the two qualities, *falsifiable* and *falsified*. The criterion of demarcation is not concerned with the theory's validity; it only determines whether the theory should be considered part of empirical science. If our tests – our attempts to find falsifications – are successful, the theory is rejected; if unsuccessful, it is accepted. But it must be falsifiable to begin with, so that we can *apply* the tests; and this quality is what makes it scientific.

A scientific theory is always falsifiable, but it may or may not be eventually falsified by tests (and even if falsified, and then abandoned, it does not lose its scientific status). Pseudoscientific theories, on the other hand, are unfalsifiable, so they can never be falsified by tests. They are, therefore, untestable. The fact that they are never falsified makes them appear successful, but in reality they are worthless; for, they do not earn their success by *passing* tests, as do the scientific theories, but by *avoiding* tests. We will examine shortly how theories can be made unfalsifiable, but we can already see the simplest way to accomplish this: by keeping their predictions vague and ambiguous, so that any event appears to confirm them. (This is typically how pseudosciences like astrology manage to appear successful.)

The principle of falsifiability can also be expressed as follows. A scientific theory makes a statement about a universe of events, dividing them into two categories: those events it permits and those it forbids. The more specific the statement (i.e., the less it permits and the more it forbids), the more valuable the theory: "Every 'good' scientific theory is a prohibition: it forbids certain things to happen. The more a theory forbids, the better it is."¹⁶ A falsification of the theory takes place when one of the forbidden events is observed to occur. So, a good scientific theory is also a theory that is relatively easy to falsify: because it forbids many more events than it permits, it actually helps us to specify tests that, if successful, would refute it.¹⁷

A good theory, therefore, makes a bold statement and takes great risks:

¹⁶ Ibid.

¹⁷ Thus, for an object moving at a certain speed in a given time period, a theory stating *The distance is the product of speed and time* is better than one stating *The greater the speed, the greater the distance*. The number of events permitted by the theory (i.e., the *correct* combinations of values) is much smaller in the first case than in the second; and the number of events forbidden by it (i.e., the *incorrect* combinations of values) is much larger. This difference is what makes the first theory easier to test, and hence, if invalid, to falsify. So this difference is what makes it more valuable.

“Testability is falsifiability; but there are degrees of testability: some theories are more testable, more exposed to refutation, than others; they take, as it were, greater risks.”¹⁸ (We hope, of course, that these tests will fail and the theory will be accepted. But the failure or success of tests, and the consequent acceptance or rejection of the theory, is a separate issue. The criterion of demarcation merely prescribes that such tests be possible.) Whereas a good scientific theory forbids a great deal, a pseudoscientific theory forbids little or nothing: any conceivable event belongs to the category of permitted events. Thus, it takes no risks. Nothing can falsify it. It is worthless precisely because it appears to work all the time: “A theory which is not refutable by any conceivable event is non-scientific. Irrefutability is not a virtue of a theory (as people often think) but a vice.”¹⁹

Recall the problem of the growth of knowledge (the fact that we can never be certain of the validity of our current knowledge) and the conclusion that the only way to progress is by trial and error. Since we cannot *prove* our theories, we must accept them with caution; we must doubt them, try to show that they are wrong, and continue to search for better ones. Seen from this perspective, a theory that cannot be falsified is a dead end: because we cannot show that it is wrong even if it is, we can never reject it; we must accept it on faith, so it is not a scientific idea but a dogma.



Published in 1934, Popper's principles of demarcation were misunderstood and misinterpreted from the beginning. Nevertheless, these principles are well known today, and are often used to expose pseudosciences. Most philosophers and scientists respect them. At the same time, we notice that few of us actually use these principles to decide whether to accept or reject a theory; that is, few of us seriously attempt to falsify our theories by subjecting them to severe tests. The mistaken belief that we must prove a theory by searching for confirmations continues to guide our decisions; and, incredibly, it affects even academic research.

It is easy to see the reason for this delusion. We tend to fall in love with our theories. We cannot bear to see them criticized. And it is even more difficult to accept the idea that it is *our* responsibility, if we are serious workers, to attack our theories. It takes a great deal of intellectual integrity, which most of us lack, to consciously design tests through which we may refute our own ideas. So, although we appreciate the falsification principle, we find it hard to adhere to it. In the end, we succumb to the temptation of *confirming* evidence.

¹⁸ Popper, *Conjectures and Refutations*, p. 36.

¹⁹ *Ibid.*

Another reason why we cannot trust verifications is that our observations are subjective and open to interpretation: "Observations are always collected, ordered, deciphered, weighed, in the light of our theories. Partly for this reason, our observations tend to support our theories. This support is of little or no value unless we consciously adopt a critical attitude and look out for refutations of our theories rather than for 'verifications.'"²⁰ In other words, we must design our tests in such a way that their success would constitute a *falsification* of the theory, not a confirmation. The observations collected in a particular test are significant only if that test sought to *falsify* the theory; they are meaningless when the test sought to *confirm* it. Thus, "every genuine *test* of a theory is an attempt to falsify it, or to refute it."²¹

Moreover, we must specify the nature of the tests, and which results should be interpreted as confirmation and which ones as falsification, at the time we propose the theory – and then *stay* with these criteria. This reduces the temptation to avoid tests found later to falsify the theory, or to modify the theory to fit the results of tests: "*Criteria of refutation* have to be laid down beforehand; it must be agreed which observable situations, if actually observed, mean that the theory is refuted."²²

Popper stresses an important aspect of the testing procedure: the requirement for "*the severest tests we have been able to design*" and for "*our sincere efforts to overthrow*" the theory.²³ Only if we resort to such severe tests and sincere efforts does their failure count as an indication of the theory's validity. Popper calls these results *corroborating evidence*: each failed test provides additional support for the theory (although, of course, not a proof). The qualities "severe" and "sincere" in these requirements are not subjective assessments of the researcher's attitude; they are exact, technical concepts.²⁴ Specifically, they mean that only *comprehensive* attempts to falsify the theory count as tests; that is, only tests which, given all current knowledge, are the most likely to falsify the theory.

²⁰ Popper, *Aim of Science*, p. 164.

²¹ Popper, *Conjectures and Refutations*, p. 36. Popper appears to be using the terms "falsify" and "refute" interchangeably. Although the difference is often subtle, in this book I use "falsify" for the individual tests, and "refute" for the theory as a whole. Since one falsification suffices to refute it, a theory that is "falsifiable" is also "refutable," and if "falsified" it is also "refuted"; but the two terms still refer to different aspects of this argument.

²² *Ibid.*, p. 38 n. 3.

²³ Both quotations are from Popper, *Scientific Discovery*, p. 418.

²⁴ Karl R. Popper, "Replies to My Critics," in *The Philosophy of Karl Popper*, vol. 2, ed. Paul A. Schilpp (La Salle, IL: Open Court, 1974), p. 1079.

4

Before continuing this study, let us pause for a moment to reflect on the significance of what we have learned. For, we can already recognize how far Popper's principles are from the *actual* way we accept new ideas and theories. We have been aware of these principles for many years, and it is an indication of the irrationality and corruption of our present-day society that we continue to base our decisions on confirmations rather than on falsifications.

It should be obvious that we must apply these principles, not only to scientific theories, but also to everyday personal and business decisions; for example, to the adoption of a new product. Products are in effect theories, not unlike the theories of empirical science, insofar as they make certain claims – claims that can be verified or falsified through experiments. So, if we want to make the best decisions possible from the knowledge available to us, we must follow the same methods when considering a new product as we do when considering a scientific theory. Since many of these products greatly affect our life, there is no reason to treat them less seriously than we do our scientific theories.

The methods employed in promotional work like advertising and public relations offer a striking example of fallacious decision-making principles. Promotions are based entirely on *confirming* evidence – typically in the form of testimonials, or case studies, or success stories. These promotional devices describe a few applications of a product, asking us to interpret them as evidence of its usefulness. Most people believe that the issue here is one of veracity: if the claims are honest, the product must indeed be as useful as it appears. But the honesty and accuracy of the claims are irrelevant, since, from Popper's principles, the very idea of assessing the usefulness of a product by means of confirming evidence is unsound. (Still, it is worth noting that, technically, the use of isolated testimonials or success stories *is* dishonest, even if the claims themselves are true. It is dishonest because it does not include the *whole* truth – i.e., all pertinent cases; and this omission is, logically and legally, equivalent to lying. Thus, the well-known principle of the whole truth is similar to Popper's falsification principle, which states that confirmations do not count as proof because they cannot include all pertinent cases.)

We see this type of promotion everywhere: in books and periodicals, on radio and television, for ordinary consumer products as well as major corporate and government issues. From pain remedies to management theories, from fitness gadgets to software systems, this type of promotion is so prevalent because it is effective; and it is effective because it exploits our natural tendency

to draw general conclusions from the observation of a small number of events – the same tendency that leads us, as we saw, to develop superstitions as readily as we develop sound theories.

But from Popper's principles we know that confirming instances prove nothing, that it is the *falsifying* instances that we must examine. What this means in practice is that the successes may be due to some unusual conditions. So we could learn a lot more by studying the *failures*. We might find, for example, that the failures exceed by far the successes, or that our situation resembles more closely those situations where the product failed than those where it succeeded.

Instead of being deceived by these promotional tricks, then, we could use them to our advantage. For, now we know that the promoters select a few confirming instances precisely because this is the only evidence they have, because the product is *not* as useful as they claim. We know that if they were honest, they would seek and discuss the *falsifying* instances – of which there are always thousands.

The link between promotions and theories is easy to recognize when we examine the *way* the promoters present their products and the *way* we assess them. The promoters propose, in effect, a theory – the theory that a given product has certain qualities and provides certain benefits; and we, on our part, develop in our mind a similar theory about its qualities and benefits. Like all theories, this theory makes certain claims and predictions regarding future situations and events; for example, the prediction that certain operations would be performed faster, or better. Logically, therefore, both the promoters and we must accept or reject this theory, not by searching for confirmations, but by subjecting it to severe tests: by sincerely attempting to falsify it. We would then accept it – we would adopt, that is, the product – only as long as we cannot falsify it, only if it survives the harshest possible criticism.

Not only do we not observe this principle, but we ignore the many falsifications (situations where the product does not work as expected) that present themselves even without deliberate testing. By ignoring these falsifications, or by belittling their significance, we render in effect the theory unfalsifiable: we accept its claims and predictions in an act of faith. Our decision-making process when adopting a product on the basis of confirming instances is, thus, an irrational act, just like accepting superstitions.

Even more disturbing is that we find this fallacy – relying on confirmations – in the most respected sources. What is the most common method of deception in advertising is also found in professional, business, and even academic publications. Articles that purport to inform or educate us, for example, are little more than stories about specific situations. Decades after Popper has shown us why we must base our decisions on falsifications, our

entire culture continues to be founded on the absurd search for confirmations. It seems that we have given up the quest for knowledge and reason, and have resigned ourselves instead to our natural tendency to irrationality.

The most blatant demonstration of this irrationality can probably be found in the world of software and programming, which, because of widespread ignorance, resembles the world of primitive man. (We studied the similarity of software-related beliefs to primitive beliefs in “Anthropology and Software” in the introductory chapter.) Respected trade and business publications routinely extol the merits of software concepts on the strength of isolated success stories. Thus, while Popper’s principles state that *one* falsification suffices (logically, at least) to refute a concept, thousands of falsifications lie all around us (instances where a software concept was *not* useful) without even being mentioned in those publications. What ought to be the most important evidence in assessing a given concept – the failures – is deliberately ignored.



If the method of selecting ideas and theories through criticism – by attempting to falsify them rather than confirm them – appears to us too severe, it may help to remember that we only feel this way about our *current* theories. We find this method perfectly logical when judging *old* theories, which have already been discredited. It is when recalling those theories that we appreciate the wisdom of Popper’s principles, because with old theories we have no difficulty recognizing how absurd is the method of searching for confirmations.

Consider, for example, geocentrism – the theory that the earth is the centre of the solar system and the universe. When we believed that the planets, the sun, and the stars revolve round the earth, we had no difficulty *confirming* this theory. After all, everything in the sky appears to move, and the ground under us appears stationary. For centuries the idea that the earth is rotating and flying through space was ridiculed. So how did we eventually reject the wrong theory and accept the heliocentric one? We did that by noting the *falsifications* of geocentrism, not its confirmations; that is, not by dismissing, but by *studying*, the discrepancies between the phenomena predicted by the theory and those actually observed. Looking back, we can easily see now that the *only* way we could progress past our geocentric delusion was by ignoring the confirmations and accepting the falsifications. Had we continued to test the theory by searching for confirmations, we would be discovering confirming instances to this day, and we would still believe that the planets and the sun are moving round the earth. And the same is true of *all* knowledge: we can only make progress by taking the falsifications of our theories seriously – indeed, by *searching* for falsifications.

It is also interesting to note that serious programmers, even if they have never heard of Karl Popper, scrupulously apply the falsification principle when testing their software. A new piece of software is similar to a theory in empirical science, in that it makes certain claims about some events – claims that can be tested through experiments and observation. Specifically, we predict that, given certain data, certain effects will occur when using the software. Thus, similarly to a theory, we assess a new piece of software by subjecting it to tests: we *accept* it as long as our tests *fail* – fail, that is, to contradict the predictions; and we *reject* it if the tests *succeed* – succeed, that is, in refuting its claims. (The rejection is only temporary, of course: we modify the software to correct the errors – creating, as it were, a new theory – and then we repeat the tests.)

It is easy to see that this testing procedure amounts to an implementation of Popper's falsification principle: we don't test the software by searching for confirmations, but by trying to falsify it. Even when an application has many errors, there are countless situations where it runs correctly; in other words, situations that *confirm* the claims made by the software. But, while it is gratifying to see our new software run correctly, we understand that it is silly to restrict testing to these situations. We all agree that the only effective way to verify software is by specifically searching for those situations where deficiencies may be found; in other words, those situations most likely to *falsify* the claims made by the software. Imagine testing software by searching for confirmations; that is, restricting ourselves to situations where it runs correctly, and avoiding situations where it may fail. We would never find errors, so the application would appear perfect, when in reality it would be unverified, and hence worthless.

The reasons for accepting or rejecting theories, or concepts, or products are very similar logically to the reasons for accepting or rejecting new software. Thus, to recognize the absurdity of accepting concepts and products on the strength of confirmations – testimonials, case studies, success stories – all we need to do is imagine what it would be like to accept a new piece of software by testing only those situations where we already know that it is correct.

5

To summarize, two principles make up Popper's criterion of demarcation between scientific and pseudoscientific theories: first, the theory must be *falsifiable* (an unfalsifiable theory cannot even be considered, because we have no way to test it); second, we accept a theory because it passes tests that attempt to *falsify* it, not because we find confirming evidence. If we remember

these principles, it is not difficult to recognize pseudoscientific theories and irrational ideas, because what their defenders do is cover up the fact that they are being falsified; and the only way to accomplish this is by disregarding the two principles.

We all wish our theories to be proved right; that is, to remain unfalsified when exposed to the reality of tests and criticism. But unlike good theories, which remain unfalsified because they are useful and make correct predictions, the unscientific ones remain unfalsified thanks to the dishonest stratagems employed by their defenders: they are made unfalsifiable from the start, or become unfalsifiable later.

The simplest way to avoid falsifications is to make the theory unfalsifiable from the start. This is typically done by formulating the claims and predictions in such a manner that they cover most eventualities, so the theory cannot be effectively tested. Thus, the claims are so vague that almost any subsequent event appears to confirm them. The fact that the theory cannot be tested – and therefore is never falsified – makes it look successful, but we already saw the fallacy of accepting a theory when all we have is confirmations. A theory is successful when it *passes* tests, not when it *avoids* tests.

Popper uses Freud's and Adler's psychoanalytic theories as examples of theories that were unfalsifiable from the start.²⁵ It is important to emphasize again that the issue here is not whether these theories are valid, but whether, in the absence of any means to test them, they are scientific; in other words, whether we can rely on their interpretations. There probably is a great deal in them that is important. Few question, for instance, the concept of an unconscious mind, or that childhood experiences affect us later in life. And, on the whole, no one denies that these theories have contributed greatly to our understanding of human behaviour. However, "those 'clinical observations' which analysts naively believe confirm their theory cannot do this any more than the daily confirmations which astrologers find in their practice. And as for Freud's epic of the Ego, the Super-ego, and the Id, no substantially stronger claim to scientific status can be made for it than for Homer's collected stories from Olympus. These theories describe some facts, but in the manner of myths. They contain most interesting psychological suggestions, but not in a testable form."²⁶

The most common stratagem, however, is not to make a theory unfalsifiable from the start, but to make it unfalsifiable later. Most pseudoscientific theories start by being falsifiable, and thus indistinguishable from the scientific ones. They are, therefore, testable. But when in danger of being falsified by certain events, their defenders find a way to save them. One can save an invalid theory

²⁵ Popper, *Conjectures and Refutations*, p. 37.

²⁶ *Ibid.*, pp. 37–38.

by avoiding tests, or by testing it without sincerely attempting to refute it, or by studying only situations that confirm it, or by ignoring the falsifications (claiming that the tests were wrong, or belittling their significance).

Although crude, these stratagems are quite effective. I will not dwell on them, though, for it is the more sophisticated stratagems that we want to examine: those employed, not by propagandists, advertisers, or irrational people, but by the academics and the experts who create the new pseudosciences. The trick they use is to suppress the falsifications as they occur, *one at a time*. And they suppress them by *modifying* the theory; specifically, they *expand* the theory so as to make the falsifying situations look like a natural part of it.

Thus, while the theory remains testable and falsifiable *in principle*, it is rendered unfalsifiable *in fact*, by incorporating into it every falsifying situation. What the pseudoscientists are doing is *turning falsifications of the theory into new features of the theory*. This stratagem may be difficult to detect, because the theory appears, at any given moment, very similar to the serious, scientific theories. It only differs from them when threatened by a falsifying situation. At that point, rather than being abandoned, it expands so as to swallow that situation – thus eliminating the threat. This task accomplished, it appears again to be a serious theory – until threatened by another falsifying situation, when the same trick is repeated.

Popper called the tricks used to avoid falsifications “immunizing tactics or stratagems,”²⁷ since their purpose is to immunize the theory against falsifications. Popper anticipated some of these tactics, but recognized that new ones can be easily invented.²⁸ He singled out the stratagems that modify a theory in order to make it correspond to the reality that would have otherwise refuted it – the trick I have just described. We will examine these stratagems in detail later, when discussing specific pseudosciences.

To combat these stratagems, Popper added a third principle to his criterion of demarcation: a theory, once formulated, cannot be modified. If we want to modify our theory (to save it from being falsified by evidence), we must consider the original theory refuted and treat the modified one as a *new* theory: “We decide that if our system is threatened we will never save it by any kind of *conventionalist stratagem*.... We should agree that, whenever we find that a system has been rescued by a conventionalist stratagem, we shall test it afresh, and reject it, as circumstances may require.”²⁹

²⁷ Popper, “Replies to My Critics,” p. 983. (Popper attributes this phrase to Hans Albert.)

²⁸ Popper, *Scientific Discovery*, pp. 81–82.

²⁹ *Ibid.*, p. 82. “Conventionalist stratagem” is the term Popper used earlier, before adopting “immunizing stratagem.” It derives from the conventionalist philosophical doctrine, which holds that a theory may be used even if falsified by observations (*ibid.*).

Recall the interpretation of theories as statements that permit certain events and forbid others. Recall also that a good theory makes very specific claims, and hence permits relatively few, and forbids most, events. Falsifying events are those events that are forbidden by the theory but do occur. Since a pseudoscientific theory forbids little or nothing, almost any event is compatible with its predictions; and consequently, it has little empirical value. Viewed from this perspective, stratagems that modify a theory in order to suppress the falsifying events reduce the number of events the theory forbids. They succeed in rescuing the theory from refutation, but at the price of reducing its value. A theory may start by making bold claims, but if it is repeatedly expanded so as to transfer previously forbidden events (which are now found to falsify the claims) into the category of permitted events, its empirical value is no longer what it was originally. It becomes increasingly unfalsifiable (that is, permits more and more events), and eventually worthless – no different from those theories which are unfalsifiable (that is, permit most events) from the start.

Popper uses Marxism as an example of theories that start by being falsifiable but are later modified by their defenders in order to escape refutation. Some of Marx's original ideas were serious studies of social history, and as such they made predictions that were testable. It is, in fact, because they were testable that they were falsified by subsequent historical events. The events, therefore, *refuted* the theory. "Yet instead of accepting the refutations the followers of Marx re-interpreted both the theory and the evidence in order to make them agree. In this way they rescued the theory from refutation; but they did so at the price of adopting a device which made it irrefutable ... and by this stratagem they destroyed its much advertised claim to scientific status."³⁰



It is always unscientific to trust a theory unconditionally; and it is this dogmatic belief that prompts its defenders to try to rescue the theory, even at the risk of turning it into a pseudoscience. We can understand now even better the requirement to doubt and criticize our own theory, to subject it to tests that sincerely attempt to refute it. Clearly, the immunizing stratagems – which aim to suppress falsifications – violate this requirement, and hence exclude the theory from the domain of science. Scientists know that they must *doubt* and *attack* their theory; pseudoscientists think their task is to *defend* their theory.

Because they do not question the validity of their theory, the pseudoscientists are bound to interpret a falsification as an insignificant exception. They

³⁰ Popper, *Conjectures and Refutations*, p. 37.

feel justified then to modify the theory to make it cope with that situation. They do not deny that the theory is deficient; what they deny is that it has been refuted. They don't see the modification of the theory as a dishonest move, but as an improvement. They believe that only a few such exceptions exist, and that their effort to make the theory match reality constitutes serious research work.

This delusion is enhanced by the fact that the falsifications are discovered *one at a time*; so each falsification looks like a small problem, and also like the only one left. But in the case of pseudoscientific theories there is no end to falsifications. The reason these theories keep being falsified is that their claims are fantastic, and thus unattainable. Pseudosciences typically attempt to explain a complex phenomenon through some relatively simple concepts. Since the simple concepts do not work, the falsifying events are not exceptions but an infinity of *normal* occurrences. By the time the theory is modified to cope with them all, there is nothing left of the simplicity and exactness it started with.

The only way the pseudoscientists can deal with these "exceptions" is by incorporating them into the theory. And they accomplish this by contriving various extensions, which they describe as enhancements, or new features. The extensions, thus, are only needed in order to bring the falsifying events into the realm of events that the theory can be said to explain. So their true effect is not to *improve* the theory, but to *degrade* it – by reducing its rigour and precision. In the end, the patchwork collection of features ceases to be a theory. Its defenders, though, still fascinated by the beauty of their original fantasies, continue to believe in it and to expand it.



All mechanistic *software* theories, we will see in this book, start by being testable and falsifiable but are later modified in order to suppress the falsifications. Consider, for example, the theory behind the relational database model (we will study it in detail in chapter 7). This theory started by claiming that, if we separate the database from the rest of the application, and if we agree to "normalize" our files, we will be able to represent the database structures with a formal system. This will ensure that the result of database operations reflects accurately the stored data. Moreover, we will access the data from a higher level of abstraction: every database requirement will be expressed simply as a mathematical combination of relational operations. Ultimately, the relational model will turn database programming into an exact, error-free activity.

Now, an experienced programmer would immediately recognize the absurdity of this concept, without even using it in a real application: since

the database structures interact with the other structures that make up the application, and since most of these interactions occur at the low level of records and fields, it is impossible to separate the database operations from the other types of operations, or to raise their level of abstraction. So the relational model was essentially an academic concept, unconcerned with reality. For example, in order to restrict programming to high-level operations on normalized files, it had to assume that processors and disk drives have infinite speed.

Still, despite its absurdity, the relational model was a falsifiable theory. And if one does not expect practicality from software concepts, it was even an interesting theory. It became testable when software companies decided to implement it in actual database systems – systems intended to serve real business requirements. At that point, needless to say, it was refuted. But instead of studying the reasons and admitting that the relational model has no practical value, its advocates started a long series of “enhancements”: they suppressed the falsifications by expanding the model so as to include the falsifying situations. This made the theory unfalsifiable, and the relational database model became a pseudoscience.

Strict data normalization, for instance, was found to be impractical, so the concept of “denormalization” was introduced; now a file could be either normalized or not – a reversal of a fundamental relational principle, and a return to the informal criteria of the traditional design methods. Then, separating the low-level database entities from the other entities of the application was found to be impractical, so various alternatives were introduced in the guise of new relational features.

In the end, as the invention of a new relational feature to suppress each new falsification proved too cumbersome, SQL, a simple database language originally intended just for queries, was expanded into a *programming* language. Although quite primitive in its new role, SQL allows us to deal with most falsifications through the traditional expedient of programming. In other words, SQL restored in a complicated and roundabout manner the low-level links between database entities, and between database entities and the other types of entities – links which the traditional programming languages had been providing all along. So SQL, while described as a relational database language, is not used to *implement* the relational principles, but to *override* them.

Thus, there is nothing left of the idea of high-level operations in today's relational database systems. Nor is there anything left of the promise to separate the database structures from the application's other structures. So the relational model could be rescued only by expanding it to include the very features it had originally excluded, and from the exclusion of which it had derived its simplicity and precision.

If this fact escapes the notice of software practitioners, it is because the features were given new names. Thus, the relational theory also exemplifies another practice frequently found among pseudoscientists: the use of new and pompous terminology for the features introduced to suppress falsifications. This trick serves to mask the fact that these are not new features at all, but reversals of claims: reinstating well-known and indispensable concepts, which the theory had attempted to eliminate in its quest for formality.

Recall Popper's principle that a theory which was modified to escape refutation must be treated as a different theory: its claims must be assessed afresh, and it must be subjected to the same severe tests as a new theory. Imagine now that the relational concepts were introduced for the first time today, as they are found in the latest relational systems. It would be obvious to everyone that their only features that are useful are those taken from the *traditional* database model. The relational theory established its reputation through its *original* claims – precisely those claims that had to be abandoned later in order to save it.

The relational theory, thus, exemplifies yet another deceptive practice employed by pseudoscientists: advertising the original benefits even after the theory was modified, its principles were forsaken, and those benefits were lost.

6

Popper's criterion of demarcation is one of the most important contributions to the theory of knowledge and to the philosophy of science. It provides an excellent solution to the problem of distinguishing between science and pseudoscience – between true research and the pursuit of delusions. It is a sad reflection on our civilization, therefore, that in the period of time since these principles were discovered, our delusions have been multiplying and flourishing at an even higher rate than before. We can observe this not only in the traditional pseudosciences (which attract the ignorant and the gullible), or in the useless products and concepts (which are, as we saw, similar logically to pseudosciences), but also in the new pseudosciences, including the software pseudosciences, which are practised in universities and corporations. We must take a moment to investigate this phenomenon.

When there exists a simple method for recognizing worthless pursuits, one would expect scientists to adopt it and to rely on it in their work. And when this method can detect delusional thinking, not only in scientific research but in all our affairs, one would expect scientists, philosophers, and teachers to explain it to the rest of us. The intellectuals have failed in this task, however, and have left the education of society to advertisers, propagandists, and

charlatans. But, what is worse, they have ignored the method even in their own work. As we will see later in this chapter, academic research means today, more often than not, simply the pursuit of a mechanistic idea. And when the idea turns out to be a fantasy, research becomes the pursuit of a pseudoscience: looking for confirmations and for ways to cover up the falsifications. What masks this degradation is the formal tone in which these worthless activities are reported.

The mechanists are in an awkward position: they cannot reject Popper's principles, because their validity is obvious; but they cannot accept them either. So they resolve this conflict by debating their philosophical meaning instead of studying their practical applications. In addition, they deliberately misinterpret the principles.

To understand why the academics prefer to ignore the true significance of Popper's principles, let us imagine what would happen if they applied these principles to their theories. Clearly, if they did, most research work in disciplines like psychology, sociology, linguistics, economics, and programming would have to be classified as pseudoscientific. They cannot *afford* to recognize the importance of these principles; for, they would then have to admit that most of their work is not scientific research but a pursuit of mechanistic fantasies.



The misrepresentation of these principles manifests itself in three ways. First, the academics present them as philosophical, rather than practical, principles: instead of being used to assess theories, concepts, and products, they have become a topic of debate in the philosophy of science.³¹ The debates involve issues such as these: When facing a falsification, how can we decide what is closer to the truth, the claim that the theory was falsified or the claims made by the theory? Or, when two theories remain irrefutable, or when both are refuted but in different degrees, how can we determine which theory is the better one?

Popper recognized these problems, and he addressed them on many occasions. But these philosophical subtleties are irrelevant when we employ his principles only as a criterion of demarcation, which was his original intent. We don't have to concern ourselves with any fine points when all we want is a way to recognize as early as possible a worthless idea. Thus, no philosophical subtleties can vindicate the frauds we are studying in this book. (Could

³¹ For examples of these debates, see Paul A. Schilpp, ed., *The Philosophy of Karl Popper*, 2 vols. (La Salle, IL: Open Court, 1974).

the settlement of such issues correct the fallacies of universal grammar, or structured programming, or the relational database model? The principles of demarcation expose the uselessness of these theories no matter how we choose to interpret the finer points.)

The most common error is to interpret his criterion of demarcation as a requirement to *actually* falsify a theory through experiments, when in fact it is the requirement for a theory to be falsifiable *in principle*: we must be able to specify some conditions that, *if* they occurred, would falsify the theory. (Of course, if we can design *actual* tests, so much the better.) Popper agrees that it may be impossible to actually falsify a theory, but he stresses that this is not required by his demarcation criterion: "An entire literature rests on the failure to observe this distinction.... And the difficulties, in many cases the impossibility, of a conclusive practical falsification are put forward as a difficulty or even impossibility of the proposed criterion of demarcation.... A flippant attack on a misunderstood logical-technical term [falsifiability] has thus led some people to far-reaching and disastrous philosophical and even political conclusions."³²

A second way to misrepresent Popper's principles is by reducing them to a weaker version. Scientists extract one aspect or another from the criterion of demarcation, and criticize that aspect alone. The isolated aspect is, indeed, insufficient to function as criterion of demarcation, so the scientists conclude that Popper's principles have no value and can be disregarded.

One claim, for example, is that the concept of falsifiability ends up treating even silly theories as scientific, simply because they are falsifiable. Astrology, for example, doesn't work. But it is falsifiable, and therefore scientific, if we accept Popper's principles. This claim, however, is invalid: it ignores the requirement that a theory be tested by looking for falsifications, not for confirmations. Astrologers always note the successful predictions and dismiss the wrong ones, contrary to this requirement. Falsifiability is not the only principle; it is complemented with the principle of searching for falsifications. Both are needed for the criterion of demarcation.

A second claim is that the requirement to abandon a theory that has been falsified would have forced us in the past to reject certain theories that were later found to be, in fact, correct. But this claim too is invalid. Popper's principles do not prevent us from modifying the falsified theory. What they say is that the new version must be considered a new theory, and subjected to the same rigorous tests as any new theory. As we saw, this is meant to prevent us from relaxing a theory's tenets over and over, which would render it in effect unfalsifiable. In practice, we may agree to ignore, say, one falsification, if this

³² Popper, *Aim of Science*, pp. xxii–xxiii.

results in a successful theory that is slightly weaker than the original one. But this is very different from dismissing one falsification after another, indefinitely, and ending up with a theory that is so weak that it has no practical value. The academics single out a few situations in the past where a particular falsification was indeed mistaken, and use them to justify their research activities, where the falsifications never end.

A third claim is that the requirement to employ severe tests and sincerely attempt to refute the theory shows that Popper's principles are not a criterion of demarcation, but merely an informal guideline for scientific research. This claim is invalid because terms like "severe" and "sincere" do not describe the scientist's attitude but the proper way to search for falsifications (see p. 217). They are technical concepts that help to make the principles of demarcation a formal, exact criterion.

Finally, some academics misrepresent Popper's principles by saying that we don't have to follow them rigorously. They generally praise them and even adhere to them, but when their theory is threatened by these principles they feel free to ignore them. Thus, it is not uncommon today for a scientist to claim that we can disregard certain falsifications, and that we can embrace certain confirmations – even as he maintains that Popper's principles of demarcation are important. The fallacy in this attitude, of course, is that these principles become meaningless if obeyed selectively, only when convenient.

So Popper had to spend much of his career explaining again and again his demarcation principles, correcting the misinterpretations, and repeating over and over what he had said clearly in his original writings. He called the mistaken views "the Popper legend": "There were books in which I had protested against the various parts of the legend, and older books and papers to which I referred in these protests, and which needed only to be read to disprove the legend. Nevertheless, the legend grew, and it continues to grow."³³

The academics are distorting and misrepresenting Popper's principles because this permits them to engage in easy mechanistic pursuits in fields where mechanism does not work. Pseudoscientific preoccupations look then like serious research, and worthless theories seem important. Few outside academia understand these principles, or even heard of them, so there is no one to prevent this fraud. In the case of software theories, moreover, the pseudoscientific notions have spread outside academia and are allowing now workers in other fields – programming, in particular – to engage in senseless mechanistic pursuits.

³³ Popper, "Replies to My Critics," p. 963.

The New Pseudosciences

The Mechanistic Roots

In the following subsections, we are going to examine some of the greatest mechanistic delusions of our time. I ignore here the fads that emerge continually in the human sciences – fads like those we encountered in “Scientism” in chapter 1. Although any one of these delusions can be shown to display the characteristics of a pseudoscience, examining them would be an interminable task. I will single out, instead, three major theories – or, rather, systems of theories – which are among the most influential intellectual movements of our time: the psychological theory of behaviourism, the social theory of structuralism, and the linguistic theory of universal grammar. The first two are now defunct, but the third one is still drawing a large number of believers.

Unlike the lesser fads, which last only a few years and attract relatively few scientists, the three theories I have selected for study dominated their respective fields for many decades. Also, their founders and supporters are world-famous scientists: men like Noam Chomsky, B. F. Skinner, Jean Piaget, and Claude Lévi-Strauss are among the best-known intellectuals of the twentieth century. As these three delusions became major research programs, they are good examples of the new pseudosciences. The discussion, however, is not meant to be a complete study of their fallacies. What I want is only to bring out their common characteristics (which they also share with the *software* pseudosciences, as we will see in chapter 7): their mechanistic foundation, and their dishonest methods. Here is a summary of the common characteristics.

Scientists who uphold these theories regard mechanism as undisputed truth. That is, the possibility of arriving at a useful solution or explanation through reductionism and atomism is not taken as hypothesis, but as established fact. It is this dogmatic attitude that prevents them from accepting the evidence later, when their theories fail. They notice a structure – a certain regularity, or uniformity, or pattern – in the phenomenon they are investigating, and immediately conclude that this structure can form the basis of a mechanistic theory.

The structure they noticed is, of course, one of the structures that make up the complex phenomenon. Their mistake is to assume that its interactions with the other structures can be ignored. They believe that a *simple* structure (their mechanistic theory, which reflects the one structure they noticed in the phenomenon) can provide a useful approximation of the *complex* structure

(the whole phenomenon). When they base their theory on one structure, when they assume that it alone can represent the phenomenon, these scientists commit the fallacy of reification: they extract that structure from the complex whole, and thereby sever its interactions (which are, in fact, the most important part of the phenomenon). And even when they do recognize that one structure alone cannot represent the phenomenon, they still expect to find a mechanistic theory, by somehow combining several structures.

The mechanistic nature of the theory can manifest itself in one or more of these features: the use of atomistic and reductionistic concepts; the use of hierarchical concepts or diagrams, of neat systems of things within things; the use of other precise diagrams, or rules, or methods, or mathematical representations. As we know, all these models are logically equivalent to a simple hierarchical structure. Mechanistic theories, in the end, always claim the same thing; namely, that a precise and relatively simple diagram, or formula, or procedure can describe and explain a complex phenomenon. They claim, in other words, that it is possible to find a deterministic representation for an indeterministic phenomenon.

Up to this point, the scientists are only guilty of wishful thinking. They are convinced that a mechanistic approximation can explain their phenomenon, so they naively emulate the methods employed in fields like physics or astronomy, where mechanistic approximations are indeed useful. But when their theory proves to be inadequate, instead of abandoning it, they forsake their responsibility as scientists and turn it into a pseudoscience: they search for confirmations; they ignore or suppress the falsifications; and, to deal with those falsifications that cannot be denied, they incorporate them into the theory.

Specifically, the scientists repeatedly expand the theory by adding various features, principles, and conditions to make the falsifying situations appear to be part of it. They coin pretentious terms for these modifications, to make them look like novel and important concepts, when in reality their function is to reinstate old, informal concepts – precisely those concepts that the original theory had tried to exclude. Often, they describe the modifications with terms like “transformation” or “normalization,” borrowed from mathematics; but, whereas in mathematics these are exact operations, in pseudoscience they are makeshift, artificial conversions, invented in order to bring the falsifying instances into the range of events that the theory can be said to account for.

It is the mechanistic dogma, in the final analysis, that fosters these pseudosciences. Even when an idea starts as an honest attempt to explain a phenomenon, even if it starts as a falsifiable and testable concept, the belief in mechanism is bound to make it unfalsifiable. If a theory is grounded on mechanism and mechanism is accepted unquestioningly, a falsifying instance

is necessarily interpreted as an anomaly, a rare exception. Its supporters, therefore, see nothing wrong in ignoring the falsification, or in modifying the theory to cope with it. They verify their theory and confirm that it does indeed obey the mechanistic principles. So the theory, they conclude, cannot possibly be wrong. If its predictions are contradicted by a certain event, something must be wrong with that event; or perhaps the theory needs an adjustment.

Now, if these falsifications were limited to a few cases, they would indeed be exceptions, and the mechanistic theory would provide a useful approximation. But in mechanistic delusions the falsifications never cease, and the theory must be modified again and again to match reality. It then becomes an unfalsifiable, and hence worthless, concept. Its supporters, though, do not consider this activity to be dishonest, or unprofessional, or illogical. On the contrary: because the mechanistic ideology has redefined science to mean the pursuit of mechanistic concepts, even when these concepts are useless, an activity that tries to save a mechanistic theory from refutation is seen as the very model of scientific work.

Behaviourism

1

The first of the modern mechanistic pseudosciences was the psychological theory known as *behaviourism*. There aren't many behaviourists left today, but for more than half a century, and as late as the 1960s, behaviourism was the dominant school in academic psychology, especially in American universities. In addition, behaviourism had a profound influence on sociology and the other human sciences.

Described as *behavioural science* – the science of human behaviour – behaviourism was seen by its advocates as an effort to turn psychology into an exact science, like physics. Psychological theories, the behaviourists claimed, will not be as successful as the theories of the exact sciences as long as they deal with the subjective and unscientific concept known as the mind. The exact sciences deal with real entities – entities that can be observed and measured. So, if psychology is to become an exact science, we must stop searching for theories of the mind, and confine ourselves to the study of human *behaviour*; namely, those human acts that can be observed, measured, and subjected to experiments.

Behaviourism, thus, rejected the traditional subjects of psychology – consciousness, knowledge, intelligence, memory, volition, emotions, beliefs, desires, fears, etc. These phenomena, the behaviourists say, are nothing but the

combination of some elementary units of behaviour; and, once we identify those units, we will be in a position to describe with precision all human acts. As in physics, therefore, we must trust the principles of reductionism and atomism, and search for the smallest bits of behaviour, the simplest human acts that can be observed: reflexes, blinking of eyelids, the movement of a finger or limb, and the like. These elementary acts are the behavioural atoms: the building blocks from which all human acts are made up, including those complex acts we attribute to intelligence. Behaviourism asserts, in other words, that there are no hidden, private, internal processes – processes requiring the invention of a concept like the mind. All human acts can be explained as a combination of simple mechanical processes, which can be observed and assessed objectively by an experimenter.

Human beings learn to display a particular combination of behavioural atoms by interacting with their environment. The basic unit of interaction is the *stimulus-response* mechanism, or S-R: an event in the environment provides the stimulus, and the organism produces the response. The responses are the behavioural atoms just mentioned; and the stimuli are the simplest events that can be perceived by the organism with its senses (the presence of a certain object, or light, or sound). When the organism is exposed to various stimuli and tries various responses, it gradually discovers certain associations between the stimuli and the responses. The associations it discovers are those that produce pleasant experiences or prevent unpleasant ones. The phenomenon whereby the organism establishes these associations is called *reinforcement*, and is taken to be a propensity of all organisms. The process whereby an organism acquires a certain set of S-R units is called *conditioning*.

The connections between stimuli and responses – the S-R units – are thus the basic elements from which all interaction between the organism and its environment is made up. The interaction is assumed to be strictly sensori-motor: the stimuli affect the senses, and the responses are muscular or glandular reactions. Ultimately, all human acts can be explained through a reduction to combinations of S-R units. There is nothing else.



The era of behaviourism started in 1913, when John B. Watson, generally viewed as its founder, published his “behaviourist manifesto”: the proclamation that psychology must be practised as an objective science, and that its goal must be, not just to observe, but to predict and control human behaviour. In other words, the task of psychologists is to study and manipulate human minds, just as other scientists study and manipulate physical entities.

Historically, behaviourism was a continuation of the mechanistic theories

of mind originated two centuries earlier by John Locke, David Hume, and David Hartley. These theories, which later became known as associationism, maintained that all knowledge can be explained as combinations of various knowledge atoms connected through stronger or weaker associations.

The early behaviourists gained additional confidence from the work of physiologist Ivan Pavlov, who investigated the process of conditioned reflexes in dogs: after learning to associate the presence of food with a certain stimulus (a specific sound or image), the dog would salivate even when this stimulus alone, without any food, was presented. This seemed to prove the stimulus-response theory – for *reflex* responses, at least. The behaviourists also liked the theory of Edward Thorndike, who experimented with cats and explained their learning behaviour as nothing more than trial and error and conditioning – a process that requires no intelligence.

Although even this limited evidence was later shown to be tenuous, the behaviourists saw nothing wrong in extrapolating it to explain, not only *all* animal behaviour, but also *human* behaviour, and even human *intelligence*. Thus, experiments with animals – especially rats – and a preoccupation with their simplest acts became the distinguishing features of behaviourism. For several decades, scientists were trying to understand human intelligence by studying the behaviour of rats through trivial experiments in which the animals were rewarded with food for performing some simple acts.

An endless variety of such experiments were designed, all for the purpose of studying and measuring with precision the process of animal conditioning. One ingenious device, for instance, invented by Skinner and known as the Skinner box, consists of a small cage equipped with a mechanism that releases a food pellet into a tray when the animal inside presses a lever. It also includes means for automatically controlling this reward and for counting and recording the animal's attempts. Through such experiments, scientists can determine the rate of learning and extinction of various patterns of behaviour under different conditions.

The behaviourists justified their experiments with animals by claiming that human behaviour, while more complex than the behaviour of rats, is not *qualitatively* different; it is only a more complicated combination of the same atoms of behaviour. The purpose of their experiments is to restrict the animal's environment so as to isolate and study these atoms. Obviously, we cannot subject *people* to experiments in a laboratory. But, whether we study animals or humans, at the lowest levels we are observing the same phenomena. Just as the same bricks are used to build both small and large buildings, the atoms that make up animal behaviour can also be used to explain *human* behaviour.

2

Behaviourism did not work, of course. It failed to explain even the behaviour of rats, let alone that of humans. It was successful only in those experiments that created a highly impoverished, artificial environment – an environment in which the animals were almost forced to display the kind of responses the experimenters expected of them. When observed in their natural environment, the animals' behaviour remained quite unpredictable, and the behaviourist theories were useless.

Like all mechanistic delusions, behaviourism extracted from the complex structure that constitutes the phenomenon in question (animal or human intelligence, in this case) a simple structure (the patterns of visible behaviour, in this case), assuming that one structure could explain the whole phenomenon. Now, there is no doubt that animals and humans are affected by their environment, that they sense stimuli and exhibit responses, and that there exists a process of associations and reinforcement which occurs somehow in conjunction with their experiences. But these patterns and regularities cannot be extracted and studied in isolation. They are only *some* of the structures that make up the existence of animals and humans, and when studying them on their own we ignore their interactions with the other structures. This is why the behaviourist model can only account for the simplest kind of behaviour – the kind for which the interactions with the other structures are indeed weak enough to be ignored.

Even for a mechanistic theory, behaviourism was very naive. In particular, it tried to explain everything with *chains* of S-R units, rather than structures of elements within elements. It claimed, in effect, that a trivial two-level hierarchy (S-R units as terminal elements and behaviour as the top element) can account for all knowledge and intelligence. Unlike most mechanistic theories, it did not try to build large, multilevel hierarchical models, so it did not even exploit fully the concepts of reductionism and atomism.

We must not be surprised that such a naive theory did not work. But let us see how, instead of admitting that it was refuted by evidence, its supporters turned it into a pseudoscience. First, they adopted the simple tactic of looking for confirmations and ignoring the falsifications. They designed their experiments not as severe tests, not as attempts to *falsify* the theory, but as means to *verify* it; specifically, as means to produce the results they wanted to see. Since they wanted to confirm that behaviour can be reduced to simple elements, their experiments consisted in creating restricted environments, in which the rats could perform *only* simple acts (pressing a bar, for instance). When the

environment was more complex (finding their way in a maze, for instance), the rats frequently displayed unexpected and more intelligent behaviour, which could not be readily explained. Since Thorndike wanted to prove that the only way cats can learn is by trial and error, he designed his experiments so that the only way to solve the problem was by trial and error. In other experiments, when confronted with different challenges, cats were shown to act more intelligently.¹

The tactic, thus, consisted in simplifying and restricting the experimental environment until the animals' behaviour was reduced to a small number of trivial, isolated acts, at which point the scientists could indeed confirm their hypothesis of behavioural atoms. In this artificial environment, models based on S-R chains did indeed provide a useful approximation of behaviour, but only because, out of the whole range of normal behaviour, the animals were restricted to isolated S-R structures. It was this limited behaviour that the model explained. When used to explain their normal, natural behaviour, which includes many interacting structures, the model failed.

Another tactic used by behaviourists to confirm their conditioning theories was, obviously, the choice of animals. Rats and pigeons were the preferred subjects in their experiments precisely because it was found that these creatures, being particularly docile and rather stupid, were most likely to display the kind of behaviour these theories postulated.

While the behaviourists were busy confirming over and over their theories with contrived experiments, their critics had no difficulty finding falsifications. The most common problem was the failure to reproduce in the real world the results observed in artificial laboratory conditions. If exposed to conditioning experiments while in their natural environment, animals ranging from pigs to whales were found to behave unpredictably, contradicting the laboratory theories.² These falsifications were ignored by behaviourists. Also ignored were the "experiments on experimenters," which showed that the laboratory measurements of rat performance that were so confidently accepted by everyone were in fact biased, and merely reflected the expectations of the individual experimenters.³

So for half a century, while the world believed that these scientists were studying human psychology, what they were studying was not even animal psychology, but some technicalities related to experiments designed to confirm

¹ Arthur Koestler, *The Act of Creation* (New York: Macmillan, 1964), pp. 568–571.

² Kellar Breland and Marian Breland, "The Misbehavior of Organisms," cited in Lawrence LeShan, *The Dilemma of Psychology: A Psychologist Looks at His Troubled Profession* (New York: Dutton, 1990), pp. 76–78.

³ R. Rosenthal and K. L. Fode, "The Effect of Experimenter Bias on the Performance of the Albino Rat," cited in Koestler, *Creation*, p. 568.

their fantasies: “In spite of the impressive mathematical apparatus, and the painstaking measurements of ‘rates of response,’ ‘habit-strength,’ ‘fractional anticipatory goal-responses,’ and the rest, rarely in the history of science has a more ambitious theory been built on shakier foundations.”⁴

3

Let us examine next how behaviourists used the other pseudoscientific tactic to make their theory unfalsifiable: repeatedly modifying the theory by incorporating into it the falsifying situations. The original theory postulated that all behaviour can be reduced to chains of S-R units, and that both the stimuli and the responses are small, atomic units, which can be observed and measured experimentally. Apart from trivial experiments, however, behaviour could not be reduced to S-R chains, and responses could not be reduced to elementary movements. Thus, because the evidence did not agree with the theory, behaviourists made the theory agree with the evidence – by expanding it to account for those situations that it could not explain. Also true to the pseudoscientific tradition, they invented impressive terms to describe the extensions. This served to mask the fact that the extensions were not new features but reversals of the original claims. What the extensions accomplished, essentially, was to reinstate the complex and inexplicable capabilities traditionally attributed to a mind – capabilities which had been specifically excluded earlier, because they could not be measured or reduced to atomic units.

For example, to account for the unaccountable responses, Edward Tolman held that there are two kinds of behaviour: higher levels, or *molar*, and lower levels, or *molecular*; and only the molecular levels can be explained with S-R units. Behaviour at the molar level is an *emergent* phenomenon and cannot be reduced to, or explained in terms of, molecular units. Edwin Guthrie invented a similar concept: the movements of the organism are low levels of behaviour, while the complex acts are high levels; and acts cannot be explained in terms of movements alone. These extensions introduced some mysterious processes between the stimulus and the response, which explained previously unexplainable responses only by remaining unexplained themselves, and were therefore a radical departure from the original goal of strict reductionism.

Tolman also introduced the concept of *intervening variables*. These variables – described as subjective and unexplainable phenomena that somehow occur between the stimulus and the response – served to revive the traditional,

⁴ Koestler, *Creation*, p. 568.

informal concept of *mental* acts. The informal concept of *drives* was also revived, and was profitably employed to explain certain types of behaviour. And to combat the limitations of the atomic behavioural units, Tolman introduced “sign-Gestalt expectations,” which used the *holistic* concepts of Gestalt psychology – a reversal of the *atomistic* principles of behaviourism.

So, little by little, the traditional psychological concepts were reinstated, and were incorporated into behaviourism in the guise of new features. The behaviourists continued to use S-R chains to explain trivial responses, and reverted to the traditional, informal concepts whenever they had to describe complex forms of behaviour.

By the time of B. F. Skinner, the last and best known of the great behaviourists, the countless “enhancements” made the theory sufficiently different from its original version to earn it the title *neobehaviourism*. Skinner added his own enhancements, of which the most important was a complete obliteration of the original meaning of stimuli and responses. And, although in his experiments he never progressed beyond chains of simple S-R units with rats and pigeons, he confidently extrapolated these results into the most fantastic theories of human knowledge and human society.

Thus, in his Skinner boxes he managed to shape the behaviour of pigeons so as to make them perform some relatively complex and unusual acts; for example, walk to a certain wall of the box and peck at a coloured disk there. He achieved that by reinforcing, in several stages, various movements which the bird had performed randomly in the direction of the disk, thus creating a chain of conditioned S-R units that looked like one purposeful act. From successes such as this, Skinner boldly concluded that everything human beings learn is also in the form of simple S-R chains, and human acts that appear purposeful or intelligent are only illusions, just as the pigeon’s act was an illusion.

He could not confirm this hypothesis, nor describe how various intelligent or creative acts can be reduced to chains of S-R units. What he did instead was modify the meaning of “stimulus” and “response” to match whatever acts had to be explained. For his rats and pigeons, these terms retained their original meaning of elementary sensations and movements. But for human behaviour, the terms expanded to include, respectively, such complex acts as reading a letter and then reacting emotionally to its contents, or being threatened with a weapon and then surrendering one’s wallet, or noticing merchandise displayed in an alluring fashion and then purchasing something. Thus, the concept of stimulus and response became so vague that it could account for any human act, thereby rendering the whole theory unfalsifiable. Moreover, the requirement to reduce complex acts to combinations of behavioural atoms – to the movement of a finger or an eyelid, for example – was forsaken. By now behaviourism had completely abandoned its original goal of being an exact

science of behaviour. Judged by their own standards, the behaviourists were now mere charlatans.

Using the new, high-level concept, Skinner even managed to describe linguistic performance (which behaviourists called “verbal behaviour”) as nothing but stimuli and responses. Again, he makes no attempt to reduce language-based communication to elementary S-R units (which might be the movement of the tongue or lips, or the vibration of the eardrum). Instead, stimulus and response refer now directly to such complex behaviour as creating and understanding sentences, formulating a challenging question, or returning an intelligent answer. Skinner’s naive views of language attracted a scathing criticism from linguist Noam Chomsky, in a review that became somewhat of a classic.⁵

Some say that the demise of behaviourism was hastened by Chomsky’s criticism and the rising popularity of his own theories of mind; others say that it was the rise of cognitive science and the theories that depict the mind as a computing device. Either way, the shift exemplifies a spectacle common in the academic world: one pseudoscience is replaced with another; one popular theory is displaced by another, which seems very different, when in reality both are rooted in the mechanistic culture and suffer therefore from the same fallacy – the belief that non-mechanistic phenomena can be represented with mechanistic models.

Structuralism

1

The movement known as *structuralism* was popular in one form or another for much of the twentieth century, especially in Europe. It flourished in the 1960s and 1970s, and had adherents even in the 1980s. Few remember it today. Structural linguistics, however, which acquired a life of its own through the work of Noam Chomsky, continues to dominate the study of language; I treat it, therefore, as a separate pseudoscience (see the next subsection, “Universal Grammar”).

The structuralists noticed that, despite their immense variety, human activities, languages, societies, customs, and institutions display many regularities. The reason for this uniformity, the structuralists say, is that all human acts are governed ultimately by the working of the brain. Thus, since human brains

⁵ Noam Chomsky, “A Review of B. F. Skinner’s *Verbal Behavior*,” *Language* 35, no. 1 (1959): 26–58.

are the same everywhere, from the most primitive societies to the most advanced, we should not be surprised to find the same patterns in the various aspects of their cultures.

Up to this point, the structuralist idea is quite sensible. When expressed informally, it is neither ambitious nor original. This modest idea, however, is only the *basis* of the structuralist philosophy. The important claim is that the biological characteristics of the brain can be described mathematically. These characteristics constitute, as it were, an alphabet of human propensities; and, once we discover this alphabet, we will be able to depict with precision every human accomplishment as a function of the human propensities.

The structuralists claim, in other words, that it is possible to represent mathematically all human capabilities; and, since the various types of human activities are in the end combinations of these capabilities, they too can be represented mathematically. Human activities, therefore, are no different from the phenomena studied by physics or chemistry. Thus, anthropologist Claude Lévi-Strauss, the most famous structuralist, claimed that the customs of all societies that ever existed are nothing but “certain combinations from a repertoire of ideas which it should be possible to reconstitute [and depict as] a sort of periodical chart of chemical elements, analogous to that devised by Mendeleev. In this, all customs, whether real or merely possible, would be grouped by families and all that would remain for us to do would be to recognize those which societies had, in point of fact, adopted.”¹

We recognize structuralism as one of those mechanistic theories that attempt to reduce to mathematics the complex phenomena studied by the human sciences. Structuralism is especially ambitious, though, in that it does not limit itself to one discipline, but claims that *all* human activities can be reduced to *the same* mental operations. Disciplines like anthropology, linguistics, psychology, sociology, political science, and philosophy can be turned into exact sciences, no different from physics or chemistry, simply by discovering the elementary human propensities. One day, the structuralists say, we will be able to explain everything in the human universe – every sentence we utter, every custom and tradition, every piece of literature and folklore, every work of art, every musical composition, every type of social organization, and even our clothes fashions and our cooking and eating habits – with equations as precise as the equations of physics.²

¹ Claude Lévi-Strauss, *Tristes Tropiques*, p. 60, quoted in Howard Gardner, *The Quest for Mind: Piaget, Lévi-Strauss, and the Structuralist Movement* (New York: Knopf, 1973), p. 118.

² It must be stressed that these were actual claims, made as late as the 1970s and 1980s. Respected scientists were actually working on theories that attempted to represent mathematically these aspects of human life.



Historically, structuralism has its roots in some of the linguistic theories proposed in the 1930s. Roman Jakobson, among others, showed that all languages share a set of common features. This, however, becomes evident only when studying the smallest elements of language: the sounds that make up phonemes. These sounds (the atoms of verbal communication) are based on a small set of elementary features. Moreover, it is possible to describe these features in terms of *binary opposites*: a phoneme is voiced or unvoiced, nasal or oral, etc. This discovery gave linguists hope that the phenomenon of language can be represented with a mechanistic model: since any sentence, in any language, can be expressed as a combination of phonemes, we should be able to reduce sentences to exact structures of sounds, and hence explain the phenomenon of language mathematically.

No one has achieved this, of course, and we know why. Language is a complex phenomenon, a system of interacting structures. The mechanists isolate these structures and study them separately, hoping to find one that can explain, alone, the complex phenomenon. The structure created by sounds plays indeed a part in language, but it interacts with the others: the meaning of words, the context in which we use a sentence, syntax rules, voice stress, and various knowledge structures present in the mind. It is impossible to explain the whole phenomenon of language with *one* structure, no matter how accurate that structure is. Thus, a theory that tries to represent language as sound structures alone is very naive – as naive as one, like Chomsky’s, based on syntactic structures alone. If we view the richness of language as the large set of alternatives for the top element of a complex structure, then an isolated structure cannot explain language because it cannot account for all the alternatives: when we separate the structures we lose their interactions, and with them many of the alternatives.

But without waiting for a confirmation of the phoneme theory with actual languages, the structuralists extrapolated it to cover, not only language, but all human capabilities. Thus, Lévi-Strauss maintained that all aspects of culture, all human activities, can be seen as forms of communication, and hence as languages; and if all *languages* seem to be based on a small set of common elements, we should also expect to find an analogous set of common elements in all *cultures*. He then proceeded to analyze hundreds of myths and customs collected from primitive societies, searching for their common elements. This analysis, according to Lévi-Strauss, is a process of *decoding*. The various myths or customs may look very different from one another, and may appear disordered, complicated, or illogical, but this is because we only see their *surface structures*; it is their *deep structures* that we must study, and it is

at these low levels that we will discover the common elements. As in the phoneme theory, the atomic concepts of human knowledge form pairs of binary opposites: left/right, good/bad, male/female, day/night, up/down, cold/warm, and so forth. Myths and customs, and all other aspects of a culture, can be reduced to combinations of such binary concepts.

We can start with any myth or custom, therefore, and through a process of *transformations* we will arrive at a structure similar to that of another myth or custom, belonging perhaps to a different society. The transformations convert a surface structure – the story told by a myth, the costume worn by a woman, the painting or carving of an object, the rules observed in a certain social setting, the food eaten on a certain occasion – into the common, atomic concepts; that is, into one of the two parts of various pairs of opposites. At this low level, all myths, customs, traditions, and institutions reveal similar structures; and the chief component of these structures is a play between opposing themes. Since the atomic concepts, according to the structuralist theory, are a reflection of the basic capabilities of the mind, it seems that an important function of the brain is to classify experiences into opposite categories.

While anthropologists like Lévi-Strauss were analyzing myths and customs, other structuralists were using similar techniques to analyze works of art, of literature, or of historiography. Their goal was the same: to find the set of basic elements (the building blocks, or alphabet) from which a whole body of works is constructed. They tried to show, for instance, that the meaning of a novel, or poem, or painting is only a surface structure; that it can be reduced, through a process of transformations, to a deep structure; and that, at this level, we find the same atomic concepts as in another novel, or poem, or painting.

Psychologist Jean Piaget believed that all human intelligence can be reduced to a small set of binary operations that are very similar to the basic operations of mathematical logic. As we grow up, our mind acquires new operations and learns to combine them into more and more complex logical structures. This theory, he claimed, is all we need in order to explain how humans perform intelligent acts of increasing complexity.

To reduce intelligent behaviour to binary operations, Piaget suggested various transformations, analogous to those defined in modern algebra: “The algebra of logic can help us to specify psychological structures, and to put into calculus form those operations and structures central to our actual thought processes.”³ He tried to prove his theory by subjecting children of various ages to intelligence tests of increasing levels of difficulty. The purpose of these

³ Jean Piaget, *Logic and Psychology*, p. xvii, quoted in Margaret A. Boden, *Piaget* (London: Fontana, 1979), p. 80.

experiments was to explain the intellectual development of the child in terms of basic logical operations. If, for example, a four-year-old child correctly solves a problem which a three-year-old child cannot solve, Piaget explains this progress by identifying some logical operators or transformations that are required to arrive at the solution, and concludes that they are only acquired by the mind at four.

The structuralists are fascinated by a rather trivial quality of binary operations: they can be combined to generate complex patterns while starting with operands that have only two values (*yes* and *no*, *0* and *1*, etc.). For example, certain problems and solutions can be represented with a hierarchical structure, if we employ this structure as a decision tree; that is, as decisions within decisions, where each decision involves two alternatives. We know that only *mechanistic* knowledge can be represented with simple hierarchical structures; but the structuralists believe that *all* knowledge can be reduced to such decision trees, and hence to the binary elements and operations known as Boolean logic (the same elements and operations used in digital circuits).

The inanity of the structuralist theories is evident in these silly analogies of minds to computers (which are far more naive than the ideas of artificial intelligence – themselves futile mechanistic pursuits, as we saw in chapter 2). Computers do indeed perform complex tasks by reducing them to simple binary operations, but the use of the word “binary” is the only thing that computer logic has in common with structuralism. Unlike the vague transformations of structuralism, computer operations can be explained completely and precisely, down to the last bit.

Thus, using terms and concepts borrowed from logic, Piaget describes the “Boolean operations” and “truth tables” that allegedly can be employed to explain human intelligence.⁴ An important set of logical operations, for instance, which appears only in adult intelligent behaviour, is the “quaternary group” of operations called INRC (which stands for Identity, Negation, Reciprocity, and Correlativity, or inversion): “What we have here is a group of four transformations of which the operations of a two-valued propositional logic supply as many instances as one can form quaternaries from the elements of its set of subsets.... The group INRC has for its elements, not the 4 cases of a truth table for 2 variables, but the 16 combinations of its set of subsets (or, for 3 variables, the 256 combinations, and so on). Because of its greater complexity, the INRC group does not make its appearance psychologically until early adolescence, whereas ... simpler models of groups of 4 elements are accessible to 7 and 8 year olds.”⁵

⁴ See chapter 4, p. 332 for a brief discussion of Boolean operations and truth tables.

⁵ Jean Piaget, *Structuralism* (New York: Harper and Row, 1971), pp. 31–32 n. 9.

What Piaget is saying here is that, as our mental capabilities develop, we can handle problems that involve more facts and more combinations of facts, because we can process larger decision trees. This is undoubtedly true, but it doesn't follow that we can represent mental processes with mathematical logic. The reason, again, is that mental processes are complex structures: when our mind develops, we gain the capacity to handle, not just increasingly large decision trees, but *interacting* decision trees. The simple structures suggested by Piaget constitute, in effect, a mechanistic mind model; and, like all mechanistic approximations, in simple situations this model may well be adequate.

Lévi-Strauss, too, takes the binary operations of computers as evidence of the validity of structuralism. For example, after struggling to find some connection between the wind and a flatfish in a certain myth, he concludes that they both function as "binary operators," because both have *yes/no* qualities (the flatfish can be seen from one angle but not from another, and the wind can either blow or not). So, "we could only understand this property of the myth at a time when cybernetics and computers have come to exist in the scientific world and have provided us with an understanding of binary operations which had already been put to use in a very different way with concrete objects or beings by mythical thought."⁶ It is hardly necessary to point out the absurdity of this comparison of myth logic to computer logic.

Edmund Leach is another structuralist fascinated by the binary operations of computers: "In some respects and in some circumstances, the products of expressive action (e.g. ritual sequences, mythological texts, poems, musical scores, art forms) show marked pattern similarity to the output of a digital computer, and when we attempt to decode such message-bearing systems we usually find that binary discriminations of the *yes/no* type are very prominent."⁷ But the only "decoding" that Leach manages to perform through his analogy to computers is some speculative interpretation of a few isolated cultural elements, no better than the interpretation reached through any other type of analysis.

2

As pseudoscientific theories go, structuralism is not very sophisticated: it belongs to the category of pseudosciences that are unfalsifiable from the start. These theories, we saw earlier, manage to escape refutation by making claims

⁶ Claude Lévi-Strauss, *Myth and Meaning* (New York: Schocken Books, 1979), p. 23.

⁷ Edmund Leach, *Culture and Communication: The Logic by which Symbols are Connected* (New York: Cambridge University Press, 1976), p. 57.

so vague that any event appears to confirm them. In the case of structuralism, it is the concepts of transformations and binary opposites that are vague and make the theory unfalsifiable.

In mathematics, transformations are well-defined operations, but the structuralists employ this term freely, whenever they want to show that one aspect of culture is related to another. In particular, they don't restrict themselves to a *fixed* set of transformations; rather, for every pair of stories, customs, or works of art which they wish to relate, they feel free to invent, if necessary, a new type of transformation. Clearly, we can always find some common elements in different aspects of culture. So, if what we seek is *any* relation, with just a little imagination we can relate any stories, customs, works of art, and so forth. The transformations are meaningless, therefore, precisely because they are guaranteed to work: there are no aspects of culture that *cannot* be related through one transformation or another. This guarantee makes the concept unfalsifiable, and hence worthless.

This weakness was pointed out by many critics. Philip Pettit, for example, after analyzing structuralism in general and Lévi-Strauss's work in particular, concludes: "The objection to Lévi-Strauss's method ... is that the sort of hypothesis that he puts up in the analysis of [myths] is just not falsifiable."⁸ "The method is hardly more than a licence for the free exercise of imagination in establishing associations between myths."⁹ Lévi-Strauss divides a myth into a number of elements, selecting those elements that best fit his purpose. Then, he relates them to the elements of another myth in any way he chooses: he may call them "equivalent," or "inverted," or "symmetrical," or anything else. In the end, "if the only constraints put on transformation are that it be achieved by a set of rules then anything can be transformed into anything: you make up the rules as you go along. Thus with a modicum of ingenuity, any two myths could be presented as transformations or versions of one another."¹⁰

Similarly, the concept of binary opposites is not restricted to a set of well-defined attributes, like left/right, male/female, or light/dark, but is extended to fit any situation. As a result, any number of contrasts can be found between the elements of a myth, or between the elements of two different myths. A particular animal, for instance, can be contrasted with a human; or, if a land animal, with a fish or bird; or, if it hunts by day, with one that hunts by night; or, if it has coloured stripes, with one that has no stripes; and so on. These pairs of attributes are indeed binary opposites, and they represent valid relations; but this doesn't mean that myths can be analyzed mathematically. The elements of myths have *many* such attributes, so myths are connected through many

⁸ Philip Pettit, *The Concept of Structuralism: A Critical Analysis* (Berkeley: University of California Press, 1975), p. 88.

⁹ *Ibid.*, p. 96.

¹⁰ *Ibid.*, p. 90.

structures at the same time, one for each attribute. The totality of myths constitutes a complex structure.

The structuralists use terms like “isomorphism,” “dimensions,” “axes,” and “matrices” to describe how the individual structures overlap and interact. But, while having a precise meaning in mathematics, these terms are only vague analogies in structuralism. Thus, Lévi-Strauss claims that “the algebra of the brain can be represented as a rectangular matrix of at least two (but perhaps several) dimensions which can be ‘read’ up and down or side to side like the words of a crossword puzzle.”¹¹ The use of terms like “matrix,” however, is the only thing that the “algebra of the brain” has in common with real algebra.

Using this sort of mathematics, Leach attempts to show that three stories from the Bible “have the same general structure and ... reflect the same narrative impulse.”¹² He presents the various elements of these stories in a complex diagram full of blocks and arrows that suggest binary opposites in three or four dimensions.¹³ Most of these opposites are contrived, as usual, but even if we accept them the diagram has no mathematical meaning. It is hard to see the point in this kind of analysis, since those conclusions that make sense – the recurring theme of good and evil, for instance – can be reached without structuralism’s mathematical pretences.

Lastly, Piaget’s reduction of intelligence to mathematical logic has been shown by more than one critic to be inconsistent and ambiguous, and hence meaningless.¹⁴ Thanks to the vagueness of his mathematics, however, the theory appears to be confirmed by almost any experiment. Thus, “it is often possible to amend Piaget’s claims so as to take account of new, apparently conflicting, evidence. But this possibility may sometimes seem too strong for comfort, suggesting that his theory is so vague as to be virtually unfalsifiable.”¹⁵

One reason why the structuralists fail to note the failure of their theories is that they always look for *confirmations*. We saw earlier that the correct way to assess a theory is by looking for *falsifications*; that is, by subjecting it to tests that attempt to *refute* it. Confirmations are worthless because, no matter how many we find, they cannot prove that the theory is valid.

The structuralists, thus, are convinced that all myths, customs, literature, etc., can be reduced to common structures through transformations, so they approach a new case, not by trying to show that it *cannot* be so reduced, but by *expecting* to find a common structure. As a result, when no meaningful interpretation is forthcoming, they keep analyzing it until they find *some*

¹¹ Edmund Leach, *Claude Lévi-Strauss* (Chicago: University of Chicago Press, 1974), p. 55.

¹² Howard Gardner, *The Quest for Mind: Piaget, Lévi-Strauss, and the Structuralist Movement* (New York: Knopf, 1973), p. 152.

¹³ *Ibid.*, p. 153.

¹⁴ Margaret A. Boden, *Piaget* (London: Fontana, 1979), pp. 82–83.

¹⁵ *Ibid.*, p. 153.

similarity or contrast to another case. The harder it is to find a meaningful transformation, the closer that situation is to being a *falsification* of the theory. But the structuralists interpret even the most contrived analysis as confirmation: the structure is especially subtle, they say, and it takes a more complex transformation to decode it. Others, though, did not fail to notice the dishonesty of this procedure: “Not a few critics complain that Lévi-Strauss is *overly* clever; that he makes distinctions and syntheses where data are lacking or ambiguous; that he ignores information incompatible with his theories and overemphasizes the limited amount of information in their favour.”¹⁶

In general, both Lévi-Strauss and Piaget have been criticized for employing imprecise terms, descriptions, and methodologies, for presenting as facts what are in reality subjective assessments, and for their inclination to interpret the results of experiments as confirmations of their theories when other explanations are also possible.¹⁷



The structuralist movement is a particularly morbid manifestation of our mechanistic culture, and a vivid demonstration of the resulting corruption. It is not surprising that serious workers were outraged by structuralism’s inane theories and its unwarranted claims to scientific status. Stanislav Andreski, for example, in his harsh criticism of Lévi-Strauss’s ideas, calls his meaningless symbols and transformations “crazy formulae” and “pseudo-mathematical decorations,”¹⁸ and the graphic depictions of sexual matters from the life of primitive peoples, with their transformation into pseudo-mathematical representation, “surrealist pornography.”¹⁹

Andreski is especially annoyed by the immense popularity that such worthless theories have among intellectuals: “No doubt the chief reason why Lévi-Strauss’s inconsequential musings about applications of mathematics to the study of culture have found such a wide acclaim is that they affect many people as hallucinogenic incantations. . . . One of the great attractions of this kind of poetry masquerading as science is that it would be very difficult to invent a topic more remote from everything that matters in social life, and better fitted for a non-committal conversation among pseudo-intellectual international bureaucrats of most divergent outlooks and loyalties.”²⁰

¹⁶ Gardner, *Quest for Mind*, p. 158.

¹⁷ *Ibid.*, pp. 219–221.

¹⁸ Stanislav Andreski, *Social Sciences as Sorcery* (London: André Deutsch, 1972), p. 133.

¹⁹ *Ibid.*, p. 135.

²⁰ *Ibid.*, pp. 135–136.

Universal Grammar

1

The linguistic theory of Noam Chomsky, based on the concept known as *universal grammar*, is seldom mentioned without being called revolutionary; even its critics agree that it has revolutionized the study of language. More than that, its influence has spread into related fields, notably psychology and the philosophy of mind.

Although it has its origins in earlier theories of structural linguistics, Chomsky's theory, first developed in the 1950s, is much more rigorous – and much more ambitious. Chomsky is searching for a theory, or model, that would account for each and every grammatical sentence in a particular language; in other words, a formal system of rules that can generate (just like a native speaker familiar with that particular language) all correct sentences, while avoiding the incorrect ones. This kind of formal grammar, which emulates a native speaker's knowledge, he called *generative grammar*. Due to the nature of its rules, it is also known as *transformational grammar*.

The study of grammar is, for Chomsky, the most important part of linguistics, and he believes that the traditional and structuralist theories failed to provide an adequate explanation of language because they were not formal enough. His project calls for a *mathematical* analysis of grammar, which would eventually allow any sentence to be formally described as a precise structure of linguistic elements: “Mathematical study of formal properties of grammars is, very likely, an area of linguistics of great potential.”¹

After more than half a century, however, Chomsky's theory still doesn't work. It has gone through innumerable versions; it has spawned countless sub-theories; it has grown into a fantastic array of rules and principles; but it still has not achieved its goal – a mechanistic model of the phenomenon of language. It can account for many aspects of language, of course, but this means very little: we know how easy it is to find mechanistic *approximations* of non-mechanistic phenomena. And the ultimate goal of Chomskyan linguistics remains as ambitious as ever: not an approximation, but a complete, formal description of all natural languages.

The fact that a theory which doesn't work can be so popular and influential in academic circles; its foundation on nothing more substantial than the observation of a few patterns and regularities; the practice of avoiding refutation by constantly expanding it to incorporate the falsifying instances; the

¹ Noam Chomsky, *Aspects of the Theory of Syntax* (Cambridge, MA: MIT Press, 1965), p. 62.

preoccupation with isolated mechanistic problems, the individual solution of which is interpreted as progress toward the explanation of the original, complex phenomenon – these characteristics make universal grammar an excellent example of the new pseudosciences.

A commonly expressed view is that, even if it will ultimately turn out to be mistaken, this theory will have made an invaluable contribution to linguistics by showing that it can be studied with the same methods as the exact sciences: “We must at least envisage the possibility that Chomsky’s theory of generative grammar will be dismissed one day, by the consensus of linguists, as irrelevant to the description of natural languages.... I personally believe, and very many linguists will share this belief, that even if the attempt he has made to formalize the concepts employed in the analysis of languages should fail, the attempt itself will have immeasurably increased our understanding of these concepts and that in this respect the ‘Chomskyan revolution’ cannot but be successful.”²

The fallacy of this view, of course, is that if the theory turns out to be mistaken it is precisely because mechanistic theories cannot explain the phenomenon of language. In this case, then, it will have made no contribution whatever to linguistics, nor to the understanding of the mind. Even more serious, we will see in the next section, is the fact that mechanistic delusions of this kind are causing great harm to society, by promoting a diminished view of our capabilities and responsibilities as individuals.

2

Chomsky maintains that our linguistic capacity has little to do with learning or culture. It is a biological trait, an *innate* human faculty: “The structure of particular languages may very well be largely determined by factors over which the individual has no conscious control and concerning which society may have little choice or freedom.”³

Thus, Chomsky says, our language faculty is akin to an organ, and we must study it in the same way we study the function of organs. Every sentence we utter or comprehend is a reflection of this language organ, and it is possible to describe with mathematical precision the working of this organ by analyzing the structure of sentences. The task of linguistics, therefore, is to discover a model that can represent all the sentences that humans utter and comprehend when they use natural languages. This model will then help us to understand

² John Lyons, *Chomsky*, 3rd ed. (London: Fontana, 1991), p. 153.

³ Chomsky, *Theory of Syntax*, p. 59.

how our mind processes language. And, since we probably have similar mental organs for performing other intelligent acts, the language model will also increase our general knowledge of the mind.

Chomsky is basing his hypothesis of an innate language faculty on a number of observations. For example, while the thousands of spoken languages and dialects appear very different from one another, on closer analysis they reveal common characteristics. Thus, sentences in all languages seem to have a neat hierarchical structure: they can be divided into distinct grammatical units (noun phrases, verb phrases, prepositional phrases, etc.), which can be further divided into parts (component phrases), then into words (nouns, verbs, adjectives, etc.), and finally into morphemes and phonemes (the smallest speech elements). Also, sentences in all languages can be modified to yield related forms: past or future tense, negative or passive meaning, etc. Languages may differ in the way the elements are combined into hierarchical structures, or in the way the modified forms are derived, but it seems that a small number of categories can account for all possible variations.

Another observation is how quickly and effortlessly children learn the particular language spoken in their community: without consciously studying the language, they acquire by the age of four or five a significant subset of the adult language, and by the age of twelve or fourteen practically the whole adult language. Thus, despite its complexity, children are capable of acquiring a language simply by being exposed to it – without having to learn its rules of grammar, and without even knowing that such rules exist. This fact contrasts, for example, with the lengthy and arduous learning process we must undergo to acquire a second language as adults. It also contrasts with the *general* mental development displayed by children: at an age when they are already proficient language users, their logical and mathematical abilities, for example, are still poor and can only be improved through painstaking learning.

We also note that all normal adults in a certain community manage to acquire the same language, despite otherwise great variations in level of education or in intellectual capabilities. It is also well known that a child will acquire whatever language he is exposed to: an English child growing up in a Japanese-speaking community will acquire Japanese just like a Japanese child.

But perhaps the most striking phenomenon is the *creativity* inherent in the knowledge of a language: individuals who acquired a language without even being aware of its rules of grammar can, nevertheless, produce an infinite number of original sentences that are grammatically correct. Also, they can instantly recognize whether a sentence they hear is grammatical or not (without being able to explain why), and they can understand the meaning of complicated sentences they have never heard before. Moreover, they accomplish this although the sentences they hear spoken in their community while

growing up, and through which they presumably learned the language, are usually an impoverished and incorrect sample of that language.

All these facts, says Chomsky, can be explained only if we assume that human beings possess, as part of their genetic structure, a *language faculty*. There is no obvious reason for different languages to share so many important characteristics, or, for that matter, to have those particular characteristics in the first place. But this is readily explained if we assume that they are all governed by the same factors: certain limitations of the human mind. An innate language capacity also explains why all humans acquire a language so quickly and easily: we become proficient language users without having to consciously learn the language because, in a sense, we already know the language. We don't learn to grow arms, or to breath, or to digest food. Our organs develop and perform specific functions without any participation from us, so why should language be different? Since verbal communication confers such obvious evolutionary advantages, the human body has evolved a specific language capacity, just as it has evolved so many other functions and organs.

The language faculty is unique to human beings; it is a species-specific aptitude, like dam building for beavers or navigation for migratory birds. We are born with the capacity to acquire language, but at the same time, because this aptitude is part of our genetic structure, we are severely restricted in the type of languages that we can acquire naturally. The similarities we observe in the various languages are a reflection of these restrictions.

Also, the fact that it is so much easier for a child to acquire at an early age the complex system of rules that make up a natural language – while having such a hard time acquiring a system like mathematics, which is simpler – points to the special position occupied by language in our mental capabilities. Our brain is wired, so to speak, for natural languages, but not for other knowledge systems. Actually, acquiring a language is not a learning process at all, but more akin to the growth of an organ. Although there are variations among individuals, just as there are variations in height or lung capacity, the basic language faculty is the same for all human beings. And since any human being can acquire any language, we must conclude that it is not the features specific to a particular language, but the characteristics common to all languages, that form the innate language faculty.

Chomsky calls this set of common characteristics *universal grammar*: “Let us define ‘universal grammar’ (UG) as the system of principles, conditions, and rules that are elements or properties of all human languages not merely by accident but by [biological] necessity.... Thus UG can be taken as expressing ‘the essence of human language.’ UG will be invariant among humans.”⁴

⁴ Noam Chomsky, *Reflections on Language* (New York: Pantheon Books, 1975), p. 29.

Children acquire so easily whatever language they happen to be exposed to because they don't actually have to *learn* the language: since they already possess the knowledge of universal grammar, all they have to do is find out, as it were, how universal grammar is implemented in that particular language.

Chomsky believes that one day we will discover the physiological roots of these innate mental functions in the brain. In the meantime, we should be able to discover the principles of universal grammar – discover, that is, a theory, or model, that exactly represents it – simply by studying the languages themselves.⁵ In fact, it doesn't even matter which language we study: whether we start with English or Chinese or Latin, we should reach the same model, because universal grammar includes only what is common to all languages. The comparative study of languages can perhaps help us to discover their common characteristics, but otherwise we may as well search for the model of universal grammar by studying the language we know best. Thus, Chomskyan linguistic concepts are derived largely from English sentences.

3

Chomsky's notions of a language faculty are, of course, pure speculations. His entire theory is grounded on the innateness hypothesis, but few people notice that the hypothesis itself is necessary only in order to account for a mechanistic theory of language. Typical of mechanistic question-begging, Chomsky started with the *assumption* that there exists a mechanistic theory of language, was then compelled to contrive an innateness hypothesis to explain linguistic phenomena mechanistically, and finally used this hypothesis as warrant for his research program. (Grounding a theory on biological and evolutionary hypotheses, instead of presenting it as a body of speculations, makes it more respectable.) The idea whose truth needs to be proved – the existence of a mechanistic explanation of language – is used as the starting point, as an assumption. This circularity is blurred by the enormous number of technical and complex aspects, and by their formal and rigorous treatment, which make the theory look like a serious scientific pursuit when in reality it is just another mechanistic delusion.

It is because people don't appreciate how fantastic its claims are that this theory is taken seriously at all. It would be instructive, therefore, to analyze its fallacies in some detail. And there is a second reason why we must devote more time to this pseudoscience than we did to behaviourism and structuralism: since the mechanistic language delusions have contributed to our mechanistic

⁵ *Ibid.*, p. 36.

software delusions, this analysis will help us later to understand the fallacies of *software* mechanism. Language and software fulfil a similar function – allowing us to mirror the world in our mind and to communicate with it; so it is not surprising that they engender the same type of delusions. (We will study this similarity in chapter 4.)

Linguistics is concerned with the study of the various aspects of language, especially phonology, morphology, syntax, and semantics. Some theories stop at the level of phonemes, morphemes, or words, but Chomsky's generative grammar, like other modern linguistic theories, is concerned with the structure of entire sentences. Significantly, linguists do not attempt to study elements of language that are more complex than sentences; they do not try to interpret, for example, the meaning of an argument encompassing several sentences. This, they say, is the task of philosophy, not linguistics.

But in normal discourse the meaning of sentences depends usually on the context in which they are used. Thus, if linguistics restricts itself to the study of isolated sentences, it must admit that there are certain aspects of language which necessarily lie beyond its range of explanations. And indeed, most linguistic theories are content to study only *some* aspects of language. Chomsky, though, claims that it is possible to discover a formal model that provides a complete and exact explanation of *all* possible sentences; specifically, a model that generates all the grammatical sentences in a given language and avoids the ungrammatical ones. In other words, he claims that we can account for all possible uses of a language from its grammar alone, without being concerned with the contexts in which the language might be used. But does this claim make sense?

In normal speech we rarely use words in isolation, so we rarely express a simple, rigid meaning of a word. When used in sentences, words can have more meanings than one could deduce by studying the words in isolation; it is the interactions between words – the complex structures generated in the mind when we interpret sentences – that provide the additional information. Similarly, we seldom use isolated sentences; a sentence is normally part of a context, and its meaning is affected by the meaning of the other sentences, by the interaction between its words and those of the other sentences, and also by any number of factors involving the persons who utter and interpret the sentences.

Thus, while there is much that can be learned about language by studying individual words and sentences, we cannot expect to detect all the information that a sentence can convey by studying it in isolation, any more than we can detect all possible meanings of a word by studying it in isolation. Yet this is precisely what Chomsky is attempting to do. He criticizes those linguistic theories that are content with an incomplete and informal analysis of sentences,

and claims that it is possible to find an exact, mathematical model that accounts for all the information conveyed by a sentence. But how can a model based on isolated sentences accomplish this?

Chomsky studies isolated sentences because he knows that it is impossible to find a mechanistic theory for the whole phenomenon of language – which would be tantamount to searching for a mechanistic theory of all human knowledge. To recognize the futility of searching for a mechanistic representation of knowledge, we only need to recall the many attempts made by philosophers to find an exact correspondence between language and knowledge (we will examine some of these attempts in chapter 4). By studying isolated sentences, Chomsky reifies in effect small portions of language, and hence small portions of knowledge, from the complex phenomenon of human intelligence. By severing the interaction of these sentences with other knowledge structures, he gets closer to a mechanistic representation of language. But what he is studying now is no longer the whole phenomenon of language.

And Chomsky goes even further: not only does he extract individual sentences from their context, but he separates the syntax of the reified sentences from their semantics. Thus, he makes the bold claim that the syntax and the meaning of a sentence are independent structures and can be analyzed separately. As evidence, he notes the following two sentences: “colorless green ideas sleep furiously” and “furiously sleep ideas green colorless.”⁶ As speakers of English we recognize both sentences as meaningless, but for different reasons: the first sentence, although meaningless in many ways, is perfectly grammatical, while the second one is not; we can easily recognize certain syntactic elements in the first sentence, while in the second one we recognize none and end up treating each word as a separate phrase. It is as if we had a feeling of familiarity with the first sentence, but not with the second one, even though we hear both for the first time; we can memorize, for example, and recall the first sentence more easily than the second one.⁷ This and other facts give Chomsky the confidence to postulate the independence of syntax from meaning. It is chiefly the syntactic structure of a sentence that determines how we interpret it: we feel more comfortable with the first sentence, although both are meaningless, because, being grammatical, our language organ can more readily cope with it.

Chomsky, thus, decided to ignore the meaning of sentences – their semantic aspect – altogether: universal grammar is independent of meaning, and we should be able to discover a precise and complete model of the language

⁶ Noam Chomsky, *Syntactic Structures* (The Hague: Mouton, 1957), p. 15.

⁷ *Ibid.*, p. 16.

faculty without getting involved with the semantic interpretation of sentences. He agrees that we use both syntax and semantics to create and interpret sentences; but he argues that we can develop separately theories of syntax and of semantics.⁸ In any case, syntax is the more important component, and it is the syntactic structure of sentences that is the essential element in a scientific study of language: “Despite the undeniable interest and importance of semantic and statistical studies of language, they appear to have no direct relevance to the problem of determining or characterizing the set of grammatical utterances.”⁹ “Grammar is best formulated as a self-contained study independent of semantics. In particular, the notion of grammaticalness cannot be identified with meaningfulness.”¹⁰

The independence of syntax from meaning is, of course, just another hypothesis Chomsky had to adopt in order to find a mechanistic model of language. Thus, he observes that all attempts made by previous linguists to include aspects of semantics led to vague and unsatisfactory theories.¹¹ But, apart from a few examples and arguments, he made no serious attempt in his original theory to show why the two can be separated. He made it clear, in fact, that the main reason he prefers to view syntax as an independent subject is that this approach offers the only hope for a rigorous study of language: “The motivation for this self-imposed formality requirement for grammars is quite simple – there seems to be no other basis that will yield a rigorous, effective, and ‘revealing’ theory of linguistic structure.”¹²

So, like the man who is looking for his keys under a streetlamp, not because that is where he lost them but because that is where there is light, Chomsky candidly admits that he is searching for a mechanistic theory simply because mechanistic theories are exact and “revealing.” This they are, of course; but a revealing theory of *language* can be discovered only if there is something to reveal – only if language is indeed a mechanistic phenomenon.



Whether it is the reification of individual sentences from a discourse or the reification of syntax or semantics from a sentence, the goal is to break down a complex knowledge structure into several simple ones – which can then be represented with mechanistic models. The phenomenon of language is the result of many interacting structures (see p. 110). It is easy to identify some of the structures that make up a sentence, but just because we can identify them it doesn’t follow that we can explain language by studying them separately.

⁸ *Ibid.*, ch. 9.

¹¹ *Ibid.*, pp. 93–94.

⁹ *Ibid.*, p. 17.

¹² *Ibid.*, p. 103.

¹⁰ *Ibid.*, p. 106.

Thus, structures like the syntax of a sentence, or the meaning of its words, or the context in which it is used, occur together; and they interact, because they share their elements. Moreover, their elements are not just the words, but also pieces of knowledge that, while not part of language, affect our interpretation of the sentence.

To convey the flavour of these issues, I will mention just one of the problems studied by Chomskyans – the problem of ambiguity. The sentence “John lost his book” can mean either that John lost his own book or that he lost another man’s book. A generative grammar based on syntactic rules, like the one developed by Chomsky, can indeed resolve this ambiguity (by treating the sentence as one phonemic string generated from two different syntactic structures, one for each meaning). This may tempt us to conclude that we can account for multiple interpretations of a sentence with a model based on syntax alone, without depending on word meaning or the context in which the sentence is used. But the sentence “John lost his way,” although syntactically identical to the previous one, can have only one meaning: losing his own way. And we can only account for this discrepancy with a model that uses *both* syntax and word meaning in the interpretation of sentences.¹³

The difficulties encountered by Chomsky and his followers, with the original theory as well as its innumerable variations, are due to the fact that the impoverished model of language he reached through repeated reifications cannot explain all possible sentences. His theory does indeed provide a mechanistic model of language, but only by failing to explain the *whole* phenomenon of language. The model ignores the interactions between structures, and it is these interactions that give language its richness. As is the case with all mechanistic delusions, Chomsky wishes to have both the richness of a complex phenomenon and the simplicity of a mechanistic model – an impossible goal. When he separated the complex phenomenon of language into simpler ones – when he severed the interactions – he renounced, in effect, the original project.

Chomsky’s mechanistic theory of language is a fantasy, and we must not be surprised that it doesn’t work. We should examine, though, how Chomsky and his followers react to its falsifications. The original concepts were expressed in the form of a falsifiable theory, and Chomsky himself recognizes the importance of falsifiability as a criterion of demarcation.¹⁴ But, while *introduced* as a testable and falsifiable theory, universal grammar became unfalsifiable soon thereafter, when its defenders started to modify it in order to suppress the falsifications. The theory was turned into a pseudoscience, thus, by the

¹³ For this example, as well as other, similar problems, see Chomsky, *Reflections on Language*, ch. 3.

¹⁴ *Ibid.*, p. 37.

decision to *expand* it, rather than abandon it, each time an aspect of language was found that could not be accounted for through the existing principles.

4

Let us briefly review the original concepts. Chomsky's first model of a generative grammar consisted of three components: the phrase-structure component, the transformational component, and the morphophonemic component. The phrase-structure component provides the rules for generating *phrase markers*; these are the *underlying strings*, or *deep structures*, of linguistic elements. The transformational component provides a set of *transformational* rules, which convert the underlying strings into *surface structures* – the final, grammatical sentences. The morphophonemic component provides the interpretation rules for converting the surface structures into the phonemic strings that make up speech.

The rules of the phrase-structure component show us how to generate an underlying string as a hierarchical structure of lexical elements. Thus, a sentence is built from certain elements, those elements from smaller ones, and so on, down to the lexical atoms – the words and morphemes that make up the underlying strings. There are only about a dozen phrase-structure rules. Thus, the top element of the hierarchy is a sentence, *S*, and is derived by concatenating a noun phrase, *NP*, and a verb phrase, *VP*; *VP* is derived by concatenating a *Verb* and an *NP*; *Verb* is composed of an optional auxiliary, *Aux*, and an actual verb; *Aux* is a morpheme like *will* or *may*, or a form of the verbs *have* or *be*; and so on. By combining and repeatedly applying these phrase-structure rules, it is possible to generate an infinite number of underlying strings. And any string generated in this fashion will eventually result in a grammatical sentence.

An underlying string may have to be further modified, by applying one of the transformational rules. The transformations manipulate words and morphemes in various ways; for instance, they modify their relative position in the string. Transformations are required in order to generate sentence constructions like negation, passive voice, and past tense, which cannot be generated directly by the hierarchical phrase-structure rules. In other words, a transformational rule must be defined for each surface structure that cannot be derived directly from a deep structure.¹⁵

It must be emphasized that all these rules were specified in a formal and precise manner – precise enough, for example, to be implemented as

¹⁵ Chomsky, *Syntactic Structures*, pp. 111–114.

a computer program. Chomsky recognized that the rules he described in his original model were imperfect, that they did not adequately define all grammatical English sentences; but he was convinced that a perfect model was attainable. In particular, he described only a small number of transformations. It was chiefly through transformations that the model was expected to improve in the future, as this concept seemed versatile enough to generate any sentence. We only need to analyze all possible sentence constructions, he believed, and determine the transformations that generate them, and we will end up with a formal definition of the whole English language.

5

The origin of the Chomskyan delusion is not without interest, and is worth therefore a short digression. The study of formal grammars and languages, along with the study of automata (abstract machines that are mathematically related to formal languages, in that they can generate or process statements expressed in these languages), formed a new and exciting field in the 1950s. The theories discovered in those days had immediate applications in the emerging discipline of computer science, in both hardware and software design. The theories of programming languages, in particular, and of compilers (the software tools that translate them into the lower-level languages of the hardware), were a direct application of the theories of formal languages.

Scientists saw great potential in the fact that a relatively simple system of specifications was all they needed in order to define a grammar or a machine, which could then generate an infinite number of different strings of elements. The principle behind this power is recursion: performing an operation with certain elements, then with the resulting elements, and so on. By nesting elements within elements hierarchically, scientists could build mathematical models of grammars or automata that displayed very complex behaviour while their definition remained completely specifiable and relatively simple.

It was natural perhaps to think that nothing lay beyond the capabilities of such mechanistic models. Reassured by the mathematical foundation of these concepts (established in the preceding two decades by pioneers like Alan Turing), and fascinated by the first computers (which were already demonstrating the practicality of these ideas), many scientists concluded that they had finally found the answer to the great mysteries of knowledge and mind: the capabilities of the models they had already built resembled some of the simpler capabilities of the mind; computers afforded the means to build models of any complexity; therefore, to attain a model with the full capabilities of the mind, they only needed to apply the same principles on higher and

higher levels. Mind mechanism – the belief that reductionism and atomism can explain the concept of mind – had entered the computer age.

Viewed from this perspective, Chomsky's fantasy is the linguistic counterpart of the other mechanistic mind fantasies of that period – fantasies which became known as artificial intelligence. The naive optimism of that period has been preserved for posterity through the ludicrous statements made by a number of scientists; namely, that computer models of the whole phenomenon of human intelligence would be attained within a few years (see p. 143).

It is significant that, although not directly involved, Chomsky always approved of the principles and goals of artificial intelligence. And it is quite irrelevant that Chomsky himself only worked on *models* of grammar: since his project calls for a complete and precise definition of natural languages, this definition could always be used to develop a computer program. Thus, his project too is, in effect, a search for a mechanistic model of mind, an attempt to replace human intelligence with software (the delusion we discussed in chapter 2).

Now, Chomsky had done some of the original work on formal languages, so he was familiar with the properties and capabilities of a series of grammars that had already been investigated – grammars called *regular* (or *finite-state*), *context-free*, *context-sensitive*, and *phrase-structure*.¹⁶ Each one of these grammars is more powerful than the preceding one, in that it can generate a greater variety of statements. Context-free grammars, for instance, are more versatile than regular grammars, and are powerful enough to serve as the foundation of programming languages. The neat hierarchical structures of elements generated by context-free grammars are well-suited for the construction of software statements, modules, and applications, as they can grow to any size while remaining unambiguous and basically simple (and hence easily processed by compilers).

Chomsky showed that these grammars are too weak to generate all the sentences people use in a natural language like English, and he assumed that all we need for this task is a grammar that is even more powerful than the phrase-structure type. He also assumed that a formal grammar powerful enough to describe a natural language would be an extension of the existing grammars, just as each one of *those* grammars was an extension of the preceding one. His original model clearly reflects this belief: the phrase-structure component is the implementation of a grammar that was already understood, while the new, transformational component provides the extension (it modifies the resulting

¹⁶ For a discussion of these grammars (including Chomsky's early contributions), see, for example, John E. Hopcroft and Jeffrey D. Ullman, *Formal Languages and Their Relation to Automata* (Reading, MA: Addison-Wesley, 1969).

strings so as to generate new types of sentences). The transformational rules were expected, in other words, to cope with all the differences between natural languages like English and simple, formal systems such as programming languages. The few transformations that Chomsky proposed were precisely specified, using a mathematical representation, just like the other rules that define formal grammars. He evidently hoped that, with additional work, it would be possible to discover a complete set of transformations, and the English language would be shown to be merely a more complex system than the others – something akin to a sophisticated programming language.¹⁷

This background also accounts for his view that grammar is independent of the *meaning* of words and sentences. A programming language, after all, can be defined without a knowledge of the actual applications that will eventually be created in that language. Similarly, a natural language must be studied “as an instrument or a tool, attempting to describe its structure with no explicit reference to the way in which this instrument is put to use.”¹⁸

For Chomsky, then, there is a difference in degree, but not in kind, between human minds and the human environment that gives rise to natural languages, and the machines controlled by means of programming languages. This diminished view of humanity is an inevitable consequence of the mechanistic dogma.

6

Let us return to the main issue: the pseudoscientific nature of Chomsky’s theory, the practice of modifying and extending it in order to escape refutation. The principal feature of the original theory was the claim that a natural language can be fully specified without taking into account the meaning of words and sentences or the context in which they are used. This idea, and hence the possibility of a formal definition of an entire language with a

¹⁷ A computer program is a system of interacting structures, so what is completely specifiable is only the individual structures. The program’s run-time performance depends on the interactions between these structures, and is therefore a non-mechanistic phenomenon. It is silly, therefore, to strive to reduce natural languages to a formal system resembling our programming languages, seeing that even computer programs, whose language *already is* a formal system, cannot have mechanistic models. What the mechanists fail to understand is that the software entities which make up a program, as much as the linguistic entities which make up a sentence, belong to several structures at the same time; and mechanistic models cannot represent the resulting interactions. We will study this problem in chapter 4. The mechanistic *software* theories are failing, therefore, for the same reason the mechanistic *language* theories are failing.

¹⁸ Chomsky, *Syntactic Structures*, p. 103.

relatively simple system of rules and principles, is what made the theory famous. The subsequent development of the theory, however, consisted mainly in the discovery of types of sentences that *cannot* be explained without resort to meaning, followed by the modification of the theory to make it explain these sentences too. And this was usually accomplished by reinstating some *traditional* grammatical concepts, which do take meaning into account. The response to each falsification, in other words, was to turn it into a new feature of the theory. The following discussion is only a brief survey of this evolution, as it is impossible to mention here all the theories and sub-theories that have formed, at one time or another, the school of Chomskyan linguistics.

Just a few years after proposing his original theory, Chomsky introduced a series of major modifications.¹⁹ (The new model became known as the *standard* theory.) There were changes in the phrase-structure component (now called the base component) and in the transformational component, but the most startling change was the introduction of a *semantic* component: deep structures were now processed both syntactically and semantically, so the resulting surface structures had both a syntactic structure and a meaning.

The new theory was more complicated than the original one, and more obscure. Neither theory worked – that is, neither went beyond a few examples and suggestions for future research – so both were mere speculations. But even as speculation, the new theory was a step backward: not only was its claim that semantics plays a role in the interpretation of sentences a blatant reversal of the original principles, but it left more questions unanswered. What was left for future research was not just some rules or transformations, as was the case earlier, but major problems in all sections of the model. We were now further away from a formal model of language, but this situation, instead of being recognized as a refutation of universal grammar, was interpreted as progress. What impressed people was, again, Chomsky's authoritative tone and the formal treatment of the problems; in other words, the fact that issues involving phrases, verbs, or pronouns were studied like issues in the exact sciences. The fact that few solutions were actually offered, and that most problems were merely stated, without even an attempt to solve them, made no difference.

Chomskians allowed semantics into their grammatical model because they believed that a set of rules can be found to define with precision the relations between word meaning and syntax. No such rules exist, of course, but the search for them has been a major issue ever since. Chomskians still do not admit that the interpretation of a sentence is related to the entire knowledge structure present in the mind, so in the new theory (and in all subsequent ones) they isolate various aspects of syntax, and search for ways to relate

¹⁹ Chomsky, *Theory of Syntax*.

them formally to the meaning of words. To pick just one example, Chomsky proposed at one time a system of *concept categories* (animate or not, abstract or not, etc.) to determine whether the use of certain types of words is valid in specific situations.

The application of semantic rules to deep structures was eventually abandoned, and was replaced by a new model (known as the *extended standard theory*), in which the semantic functions are performed mainly on surface structures. But to retain the links to the syntactic structures, a complicated *trace* sub-theory was developed to allow the transfer of such information as the position of words in the sentence, from the deep structures to the surface structures. In the meantime, other linguists proposed a theory of *generative semantics*, which tried to build the meaning of sentences from the meaning of smaller elements.

None of these theories worked, so the next step was to replace the entire transformational philosophy, which was based chiefly on systems of rules, with a new model, based on *principles and parameters*. Chomsky argues now that languages can be described as sets of principles, where each principle can be implemented only as one of the alternatives permitted by universal grammar. All languages are basically the same, the only difference being in the implementation of these principles; and language acquisition consists in the unconscious discovery of the correct alternatives for a particular language. It is as if our language organ had a number of switches, all set at birth in a neutral position and ready to accept any value (from among the values permitted by universal grammar). What we do when acquiring the first language is set these switches to one value or another.

This is how Chomsky describes the new concept: “The principles are language-independent and also construction-independent; in fact, it appears that traditional grammatical constructions (interrogative, passive, nominal phrase, etc.) are taxonomic artefacts, rather like ‘terrestrial mammal’ or ‘household pet.’ These categories, with their special and often intricate properties, result from the interaction of fixed general principles, with parameters set one or another way. Language acquisition is the process of determining the values of parameters. There are no ‘rules of grammar’ in the traditional sense: rather, language-invariant principles and values for parameters of variation, all indifferent to traditional grammatical constructions.”²⁰

This text is typical of Chomsky’s writing style: he is describing some linguistic fantasies, but by presenting these fantasies in an authoritative tone he makes them look like a scientific revolution.

²⁰ Noam Chomsky, “Chomsky, Noam” self-profile, in *A Companion to the Philosophy of Mind*, ed. Samuel Guttenplan (Oxford: Blackwell, 1995), p. 161.

The new theory, Chomsky declares, is “a conception of language that [departs] radically from the 2500-year tradition of study of language.”²¹ Unfortunately, while elements “of the picture seem reasonably clear” (to Chomsky, at least), “a great deal is unknown, and clarification of principles regularly opens the doors to the discovery of new empirical phenomena, posing new challenges. Though much less is understood, something similar must also be true of the lexicon, with the links it provides to the space of humanly accessible concepts and signals.”²²

Thus, Chomsky admits, what is “much less understood” than the part of which “a great deal is unknown” is (as always) the interaction between language structures and the other structures that make up human knowledge; in other words, the actual, complex phenomenon of language, as opposed to the reified, mechanistic phenomena studied by linguists.

Note again his authoritative tone, even as he is describing what are, in fact, mere speculations. For example, while admitting that we know practically nothing about a certain phenomenon, he confidently asserts that certain aspects are “reasonably clear,” and that “something similar must also be true” of others. This is the same confidence that brought us the previous theories, all now forgotten. So now we have a new revolutionary theory that is mere speculations and doesn’t work, to replace Chomsky’s other theories that revolutionized linguistics though they were mere speculations and didn’t work.²³

Note also, in that passage, the statement about “new empirical phenomena” being regularly discovered and “posing new challenges.” This assertion illustrates how pseudoscientific thinking distorts the idea of research – from an effort to discover the truth, to an effort to save a theory from refutation: “new empirical phenomena” is a euphemistic term for the falsifications of the theory, while the “challenges” constitute the search for ways to turn these falsifications into new features; that is, ways to expand the theory so as to account for them and thus escape refutation.



It is instructive to take a look at some of the principles that make up the new model:²⁴ *X-bar* theory deals with phrase structure and lexical categories and

²¹ *Ibid.*, pp. 160–161.

²² *Ibid.*, p. 161.

²³ Like the previous theories, the new one did not last long. In the following years many of its principles were abandoned, and by the late 1990s another linguistic revolution – another batch of speculations – was being promoted: the so-called minimalist program.

²⁴ See, for example, Noam Chomsky, *Knowledge of Language: Its Nature, Origin, and Use* (Westport, CT: Praeger, 1986).

their mutual relationships. *Theta* theory deals with the thematic roles (agent, patient, goal) played by elements in a sentence. *Case* theory deals with the assignment of case (nominative, accusative, genitive) to noun phrases. *Control* theory deals with the subject of infinitival clauses (the relation between the missing subject and the other elements in the sentence). *Binding* theory deals with the problem of expressions that refer or not to the same entities as other expressions in the sentence (as in constructions involving pronouns or anaphors). *Bounding* theory deals with the movement of grammatical units from one place in the sentence to another (as when deriving passive or interrogative constructions).

What is immediately striking about these principles, or sub-theories, is that each one deals with a single, isolated aspect of grammar. There are many other, similar principles in the new model, and additional ones are known but little has been done to study them. New principles, Chomskyans say, will undoubtedly be discovered in the future. And to cover any grammatical cases that may remain unexplained no matter how many principles and sub-theories will be discovered, the concepts of *core* and *periphery* have been introduced.²⁵ Every language, it appears, has two types of grammatical constructions: the core is that part of language explained by universal grammar, while the periphery includes those aspects of language that somehow evolve outside the scope of universal grammar.

The theory, thus, has become blatantly unfalsifiable, as any conceivable sentence and any aspect of grammar is now guaranteed to be accountable: either it is explained by the known principles, or it will be explained by principles yet to be discovered, or it doesn't need to be explained at all, because it belongs to the periphery. Little by little, Chomskyan linguistics has turned into a full-fledged pseudoscience.

If we compare the new principles to the original theory, what we notice is the evolution from a simple and elegant model that made bold and sweeping claims, to a collection of distinct and rather complicated theories that deal with isolated and minute aspects of grammar. It is also interesting that these aspects are not unlike those studied by *traditional* grammars. So, if we ignore the new terminology, many of these concepts are in fact a reinstatement of older grammatical concepts, which had been excluded by the original theory when it claimed that a relatively simple system of rules can explain a whole language. And we must recall that it was its simplicity and elegance that made the original model so attractive in the first place. Thus, Chomskyan linguistics continues to benefit today from its original prestige, even though its current features and claims are, in many respects, the exact opposite of the original ones.

²⁵ Ibid., p. 147.

Chomskyans stress now the benefits of the “modular” approach to the study of language: each sub-theory forms an independent module, which can be studied separately, while the modules also interact and work together as one system – the language faculty. Chomskyans draw block diagrams to depict these mental language modules and their interactions; and they connect the blocks with arrows, and use terms like “input” and “output” to describe the alleged data flow in the mind. The entire language faculty is treated then as one module among the many modules of the mind (the other faculties), which are believed to be relatively independent while interacting and working together to produce intelligence. It is hardly necessary to point out the mechanistic nature of this model: Chomsky’s study of language and mind looks now just like an engineering project whose difficulty was originally underestimated.

This evolution is typical of mechanistic delusions: Chomsky started with a fantastic claim – the claim that a fairly simple model can provide an exact and complete explanation for the phenomenon of language. To make such a claim, he had to *assume* that the phenomenon is mechanistic in nature; namely, that it can be explained by explaining separately the simpler phenomena which appear to make it up. This led to the reification of language from the whole phenomenon of human knowledge, the reification of syntax from the phenomenon of language, and, finally, the reification of individual aspects of syntax. The reductionistic procedure looks perfectly logical – if we forget that the mechanistic nature of the phenomenon is only a hypothesis. With this hypothesis, we can always break down a phenomenon into simpler and simpler ones. Eventually, we are certain to reach phenomena that are simple enough to explain with mechanistic models – with rules, diagrams, mathematics, etc.

It is clear, then, why Chomskyans believe that they are making progress. They keep finding explanations for isolated grammatical phenomena, and they believe that these explanations will one day be combined into an explanation of the original phenomenon. But language is a complex phenomenon. So even if one day they manage to identify all its constituent structures, their model will still not work, because mechanistic models cannot represent the *interactions* between structures.

It is interesting that the new theory specifically depicts language as the result of many *interacting* principles of grammar, all sharing the same linguistic elements. The theory describes, therefore, a complex structure; and these principles are, quite correctly, some of the simple structures that make up the phenomenon of language. Chomskyans, however, still fail to see that it is impossible to explain a complex structure by explaining separately its constituent structures. And they still fail to see that the phenomenon of language involves, not only grammatical structures, but many knowledge structures present in the mind.

7

Chomsky compares our current linguistic knowledge with our knowledge of physics before Galileo. He modestly admits, with each new model, that these are only beginnings, that there is much work left to be done. He believes that, just as Newton synthesized the knowledge of his time and discovered the laws of universal gravitation, a similar breakthrough will take place one day in linguistics, when someone will discover a unified theory of language.²⁶

Chomsky's belief in language mechanism is unshakable: he does not doubt for a moment that the phenomenon of language can be explained, just like gravitation, through reductionism and atomism. Viewed in this light, the practice of modifying the theory to account for contradicting empirical evidence may look like a legitimate research method – a way to improve the theory. Recalling Popper's principles, however, the scientist must sincerely attempt to *refute* his theory. If he modifies it to avoid the falsifications, he does the opposite: he attempts to *save* the theory. The scientist must specify, when proposing his theory, what events or situations, if observed, would refute it. And if subsequent tests reveal such events or situations, the correct response is to declare that theory refuted, propose a *new* theory, and specify what events or situations would refute *it*.

If we keep this principle in mind, it becomes clear that Chomsky is not trying to refute his theory, but to save it. We must not be confused by his endless models; these models are not really new theories that replace previously refuted ones, but different versions of the *same* theory. Chomsky's theory is not just a formal model of grammar, but the system comprising a model of grammar *and* the idea of an innate universal grammar. One cannot exist without the other. The search for a mechanistic model of grammar is motivated by the innateness hypothesis – the hypothesis that humans possess a language faculty which is akin to an organ. *This hypothesis* is, in the end, Chomsky's thesis, what has made the whole theory unfalsifiable and hence pseudoscientific. The innateness hypothesis never changed, and it is in order to save *it* from refutation that all those models of grammar – all the theories, sub-theories, and principles – had to be invented, modified, and extended.

But why is the innateness hypothesis so important? Why does Chomsky defend it at all costs? Because, he frequently asserts, it is the only logical alternative. An innate language faculty is the only way to account for the ease

²⁶ See, for example, Noam Chomsky, *Language and Politics* (Montréal: Black Rose Books, 1988), p. 418.

and speed with which children learn a language, especially when we consider the impoverished sample they are exposed to; it is the only way to account for their ability to create correct sentences which have little resemblance to those they heard before; and so on. Since we can think of no other explanation, says Chomsky, we must accept the hypothesis of an innate language capacity.

But is it true that there are no alternative explanations? Only if we assume that language is a mechanistic phenomenon do we have to resort to an innateness hypothesis. If we admit that there are complex phenomena in this world – phenomena which cannot be explained through reductionism and atomism – then an alternative hypothesis is that the linguistic capability of humans is a complex phenomenon.

The circularity characteristic of mechanistic thinking is, again, obvious. Because he wishes to explain language with a mechanistic theory, Chomsky must conceive a second mechanistic theory: the innateness hypothesis (which is, in effect, the notion that there exists in the mind a thing whose operation can be described with precision). Then, he uses this hypothesis as warrant for his linguistic theory. Chomsky must assume both language mechanism and mind mechanism at the same time. One mechanistic assumption is adduced to justify another. The mechanistic philosophy is invoked to defend the mechanistic philosophy.



Since the entire Chomskyan project is grounded on the innateness hypothesis, we should perhaps investigate the soundness of this hypothesis. In our discussion of skill acquisition, we concluded that it makes no sense to postulate the existence of specialized high-level mental functions (see “Tacit Knowledge” in chapter 2). We saw that the same model of mind can account for any skills: general skills acquired simply by belonging to a human society (using language, interpreting visual sensations, recognizing social contexts), and specific skills selectively acquired by each individual (playing chess, interpreting X-rays, programming computers). We develop the necessary knowledge by being exposed to the phenomena – that is, the complex structures – which embody that knowledge. Our mind discovers the simple structures (the regularities) in the information captured by the senses, and creates an approximate replica of the complex structures by discovering also the interactions.

Complex structures can exist only in the phenomena themselves and in the mind; they cannot be transferred directly into a mind, because our senses communicate with our environment only through simple structures (through systems of symbols or sounds, for instance). The complex structures formed in the mind manifest themselves as non-mechanistic knowledge: we can *use* this

knowledge, but we cannot precisely describe what we know. In other words, we cannot reduce this knowledge to simple structures. Non-mechanistic knowledge is the type of knowledge we possess when we reach expertise in a particular skill.

The human brain may well have some specialized low-level innate functions, like those found in simpler animals. And such functions may even take part in our verbal acts. But it is both absurd and unnecessary to postulate innateness in order to explain *high-level* mental capabilities; that is, to assume specialized faculties to account for *particular* skills, as Chomsky does.

It is absurd, first, from an evolutionary perspective: low-level functions, or instincts, play a dominant part in the behaviour of simple organisms, and the brain has evolved precisely in order to confer the advantages of learning. It makes no sense to assume that language – the most human-specific faculty, perhaps our most complex capability – is handled mostly by innate functions, while the *learning* functions of the brain, which have evolved specifically as an improvement over innate functions, play only a secondary part.

Another reason why the innateness hypothesis is absurd is that it leads to the conclusion that we possess a specialized faculty for each skill we can acquire. We might perhaps accept the innateness hypothesis for those skills acquired early in life by all humans – using language, recognizing faces, etc. But we saw that there is no fundamental difference between these natural skills and the skills related to a particular culture or occupation, which can be acquired at any age. All skills can be accounted for through a mind model based on complex knowledge structures.

Recall the skills we studied in “Tacit Knowledge.” No one would suggest that we possess a specialized faculty for playing chess, or for interpreting X-rays, or for distinguishing chicks. Humans can acquire thousands of different skills, so we must conclude that the *same* mental capabilities are used in all of them. And if we can acquire so many skills using some generic mental capabilities, why do we have to assume that some other skills – like the use of language, which also can be accounted for by the same model of mind – are innate? The innateness hypothesis is unnecessary if we accept the existence of complex mental structures. Chomsky postulates specialized mental faculties, not because of any evidence that such faculties exist, but because this is what he needs for his mechanistic mind model.

And what about those linguistic phenomena Chomsky says can only be explained by an innate language capability? The fact that languages have so many common features, and the fact that children learn a language so quickly and easily, can be explained, just as Chomsky says, by an innate characteristic: our brain has developed the capability to process hierarchical knowledge structures. So this characteristic may well be reflected in our languages too: in

each one of the various structures, including the grammatical structures, that make up the phenomenon of language. What Chomsky chooses to interpret as a specialized *language* capability – the mind’s capacity for hierarchical structures – is a *general* capability. It is a capability that can be observed in *all* mental acts.

For example, when we see a face we perceive its elements and attributes as structures, not as isolated parts; we don’t notice one eye, then the chin, then the nose, then the other eye; we don’t study the elements randomly, or left to right, but unconsciously perceive them as several facial structures that exist at the same time. An expert chess player doesn’t perceive a position by examining the pieces on the board in a certain sequence, but by unconsciously recognizing many overlapping and interacting logical structures. Similarly, we don’t make sense of a sentence by combining words randomly, or left to right, but by detecting structures of grammar and of meaning. We perceive *everything* as structures, but this is masked by the fact that these structures share their elements, so we perceive them simultaneously. As we saw in “Tacit Knowledge,” only when *inexperienced* in a certain domain do we notice the individual structures separately.

As for the creative aspect of language – our capability to utter and comprehend an infinite number of sentences that only resemble, and only in unspecifiable ways, those we heard before – it too is not peculiar to linguistic skills, but common to all skills. The distinguishing aspect of expertise, we saw, is the capability to recognize new situations intuitively. As novices, we can only cope with a new situation mechanistically; that is, by following rules and by decomposing it into familiar elements. After a great deal of practice, however, when we reach expertise, we can cope with new situations directly, holistically. Expertise, therefore, permits us to cope also with *complex* situations, which cannot be precisely described as a combination of familiar elements.

Thus, we can recognize a familiar face from any angle, or from any distance, or in any light, or in a photograph; that is, when the image formed on the retina only *resembles* the previous images. And we cannot describe with precision how we recognized the new image, nor in what ways it resembles the previous ones. An experienced radiologist correctly interprets X-ray pictures that are necessarily only *similar* to others he saw previously. Expert chess players recognize positions that are only *similar* to previously encountered ones. To drive a car we must be able to handle random traffic situations, which at best *resemble* previously encountered ones. Moreover, we need this capability not only with specific skills, but to perform any intelligent act. We need it, in fact, just to behave normally in everyday situations; in this case we need it in order to recognize *contexts* that only resemble previous ones.

So, if this capability is used in all mental acts, why not also in the acquisition

of language? Language is a complex structure, and the child discovers its constituent structures, including the grammatical structures and their interactions, simply by being exposed to it – as is the case with other skills. Also like other skills, the child manages to cope with novel situations; that is, he can create and understand sentences that only *resemble* previous ones. The complex knowledge the child acquires in the case of language includes the grammar; so, as is the case with any non-mechanistic knowledge, he can benefit from his knowledge of grammar without being able to describe what he knows.

In conclusion, the innateness hypothesis – the foundation of Chomskyan linguistics – is not as solid as Chomsky believes. And without this foundation, his theory is left as just another mechanistic mind delusion: just another system of belief, and no more of a science than behaviourism or structuralism. Despite the preciseness observed by most Chomskysans in their work, their models can lead nowhere if the project itself is unsound. No matter how many rules or principles they study, or how successful they are in reducing each one to a formal representation, these reified structures cannot improve our understanding of the phenomenon of language, nor of the human mind. There is no innate language faculty, and universal grammar is a mechanistic fantasy.²⁷

Consequences

Academic Corruption

In the previous section we studied some of the more influential mechanistic delusions of our time – modern pseudosciences pursued in universities and accepted by large numbers of scientists. By discussing these pseudosciences here I am making a statement; namely, that I view our *software* delusions as a social phenomenon belonging to the same tradition. (The *language* delusions, as a matter of fact, have contributed *directly* to the software delusions. We will study this link in chapter 4.)

The theories of software engineering – the relational database model, structured programming, object-oriented programming, and the like – are in

²⁷ More than a few thinkers have criticized Chomskyan linguistics, of course, sometimes with arguments very similar to those presented in the foregoing discussion. No criticism, however, and no falsifications, can affect the popularity of a mechanistic idea among the academic bureaucrats.

the domain of programming what behaviourism, structuralism, or universal grammar are in the human sciences: mechanistic delusions, naive attempts to represent complex phenomena with exact models. They are the work of academic bureaucrats: individuals who cannot make a real contribution to their discipline or to society, and who hide their incompetence by imitating the methods of the exact sciences. Through this imitation they appear to be engaged in serious research, while pursuing in fact a pseudoscience.

One consequence of the mechanistic dogma, thus, is the intellectual corruption it fosters. These theories do not work, and they cannot possibly work; but because mechanism is taken as unquestionable truth, each falsification is seen as a challenge – the challenge to find ways to *deny* that it is a falsification. The theory becomes then unfalsifiable: it changes from a naive hypothesis to a full-scale system of belief, a pseudoscience.

This mechanistic culture is what allows now the *software* elites to deceive society, with *their* mechanistic concepts. For it is in universities and other research institutions that the software fantasies emerge: among individuals whose programming experience is limited to textbook examples, to trivial problems and neat solutions. To them, the possibility of finding exact models for complex, real-world software phenomena is as certain as is the possibility of finding exact models for complex psychological, social, or linguistic phenomena to their colleagues in the human sciences. The software fantasies are not so extraordinary once we recognize their grounding in the mechanistic ideology, and their similarity to the other academic fantasies; nor is extraordinary the dishonesty of their promoters and the evolution of the theories into pseudosciences.

It is impossible to assess the price we pay for these mechanistic obsessions and their ramifications. We cannot even imagine the progress we might have made in the human sciences, had the effort wasted on futile mechanistic theories been invested in other directions, more likely to increase our understanding of human minds and human relations; specifically, in theories that attempt to explain whole human phenomena, rather than break them down into simple and independent processes as if they were engineering projects.

Another consequence of our mechanistic obsessions is the prevention of expertise and responsibility. Workers in all fields are expected to follow blindly the principles of reductionism and atomism, rather than to search creatively for solutions and explanations. Instead of seeking to increase and broaden their knowledge, these two principles – which are taken as “the method of science” – allow them to equate expertise with narrow specialization: knowing as little as possible is perceived as a virtue, as a sign of professionalism. And if a narrow domain still requires too much knowledge, workers invoke these principles again and again, until they finally reach those low levels in the structure of

knowledge where even the most ignorant people can be experts – levels where they only need to deal with trivial and isolated problems.

This trend has affected all occupations that involve knowledge and skills, but is especially noticeable in research work. Science has been redefined: an individual is considered a great scientist simply for discovering a mechanistic theory, regardless of whether the theory works or not. Thus, a mechanistic culture rewards mediocrity and discourages creativity. To be successful in academia, an individual must think like a bureaucrat and must accept blindly the mechanistic doctrine. Moreover, creative individuals who could make an important contribution are ignored, or see their work branded as “unscientific,” simply because they reject the mechanistic principles and try to deal holistically with complex phenomena.

And this trend is just as widespread in our software-related activities – in universities, in business, and now even in our personal affairs. An individual is considered knowledgeable simply for accepting the latest mechanistic software concepts, regardless of whether these concepts are valid or not. To be successful in a software-related career, an individual must have the temperament of a bureaucrat, must restrict himself to mechanistic practices, and must display an unwavering allegiance to whichever authority is supporting the doctrine of software mechanism.

Psychologist Abraham Maslow¹ suggests that mechanistic beliefs are a sign of immaturity and insecurity. Instead of seeking to understand the complex reality, the mechanists prefer the comfort of artificial, narrow domains, where it is easy to find theories: “Science, then, can be a defense. It can be primarily a safety philosophy, a security system, a complicated way of avoiding anxiety and upsetting problems. In the extreme instance it can be a way of avoiding life, a kind of self-cloistering. It can become – in the hands of some people, at least – a social institution with primarily defensive, conserving functions, ordering and stabilizing rather than discovering and renewing.... The greatest danger of such an extreme institutional position is that the enterprise may finally become functionally autonomous, like a kind of bureaucracy, forgetting its original purposes and goals and becoming a kind of Chinese Wall against innovation, creativeness, revolution, even against new truth itself if it is too upsetting.”²

In many academic disciplines, and in our software pursuits, our culture increasingly resembles the culture of primitive societies, or of the West during the Dark Ages, or of totalitarian states. What characterizes these cultures is their dogmatic value system, grounded on belief instead of logic. In our culture the dogma is mechanism. This is a scientific rather than religious or

¹ Abraham H. Maslow, *The Psychology of Science: A Reconnaissance* (South Bend, IN: Gateway, 1966).

² *Ibid.*, p. 33.

political dogma, but its consequences are the same: intellectual stagnation; an ignorant population susceptible to irrational ideas, and hence to deception and propaganda; and, in the end, a society dominated by corrupt elites that exploit these weaknesses.



In the following subsections, I want to discuss the social and political consequences of mechanistic thinking. Specifically, I want to show that the mechanistic theories promoted in our universities, even though invalid, are shaping the future of our society – by fostering totalitarianism.

The reason we must study the consequences of mechanism is that the belief in *software* mechanism has created the conditions for mechanistic theories to be actually implemented in society. The concepts promoted by mechanistic theories in sociology, psychology, and linguistics have undoubtedly influenced our world view, but they were never implemented on a large scale. These theories may pervade the academic world, but they have no direct application in business, or in politics, or in our social or personal affairs.³ The mechanistic *software* theories, on the other hand, promise immediate benefits to everyone. They are appealing because they address the use and programming of computers, and we now depend on computers in practically every activity.

For example, corporate managers who have never heard of structuralism, and who would probably dismiss its fantastic claims, accept software ideas described as “solutions,” without realizing that these ideas are based on the same mechanistic delusions as the structuralist theories. And liberal politicians who have never heard of behaviourism, and who would never endorse its totalitarian policies, accept the utopian promises of the “information revolution,” without realizing that these promises are based on the same vision as the behaviourist theories.

So we accept mechanistic *software* theories, which are just as worthless as the traditional ones, not because we understand their principles better, but because their claims address immediate concerns. And we do not recognize their common totalitarian aspects any better than we do their common mechanistic principles. Thus, to recognize the totalitarian tendencies of the *software* theories, we must start by examining the totalitarian tendencies of the *traditional* mechanistic theories. (These two types of theories are the subject of the next two subsections.)

³ We must remember, though, that totalitarian systems like Nazism and Communism were founded on mechanistic social and economic theories. And, as we will see in chapter 8, the democratic systems too are moving in this direction.

The Traditional Theories

1

What do all mechanistic delusions have in common? If we represent as structures the phenomena they try to explain, then what they all claim is that it is possible to account for all the values of the top element from the values of the starting elements. For behaviourism, the starting elements are bits of behaviour, and the top element comprises all possible behaviour patterns and intelligent acts. For structuralism, the starting elements are bits of knowledge or logic, and the top element comprises all human knowledge, accomplishments, social customs, and institutions. For universal grammar, the starting elements are words and elementary sounds, and the top element comprises all valid sentences and some of the knowledge embodied in sentences. So these theories claim that we can explain precisely and completely, starting with some simple elements, all possible manifestations of the human phenomenon in question – mental acts, social behaviour, linguistic competence, etc.

The significance of the claims is evident, therefore, when the phenomena are seen as structures. We immediately notice that the theories describe *simple* structures. They may use diagrams or equations rather than a structure; but we know that if they attempt to provide a precise explanation, they are deterministic models, so they could also be represented with a simple structure. And we also know why these theories do not work: because the phenomena they try to model can be usefully represented only with *complex* structures.

The fact that they do not work, thus, is not surprising. It is important to note, however, the *claim*, or the *expectation*, that they work. The scientists who defend these theories *wish* them to work. Specifically, they *wish* the phenomena to be simple structures, and the top element to be precisely describable in terms of the low-level elements. But if the phenomena are in reality complex structures, if they are the result of interactions between the simple structures the scientists recognize and some *other* structures, then what these scientists do in effect is deny the importance of those other structures; that is, they deny their bearing on the value of the high-level elements. And what are those other structures? They are the phenomena created by human minds: the knowledge, the experience, the creativity, the intuition of individual human beings.

When behaviourists say that intelligent behaviour can be computed from elementary units of behaviour, or when structuralists say that knowledge and social customs can be computed from elementary bits of logic, or when linguists say that sentences can be computed from words, what they claim in effect is that there is nothing between the low levels and the high levels that is

unpredictable. They claim, thus, that we can describe the high levels in terms of the low ones just as we describe the operation of a machine in terms of its subassemblies and elementary parts.

But in the case of human beings and human societies, the high levels *are* unpredictable; and this unpredictability is what we understand as creativity, free will, and indeterminism. The indeterminism is caused by the complexity of interacting structures: the knowledge structures formed in individual minds, and the structures formed by many minds in a society. The structures studied by behaviourists, structuralists, and linguists are indeed among the structures that make up minds and societies. Taken alone, though, these structures cannot explain entire human phenomena; and this is why their theories do not work. In the end, the failure of the mechanistic theories constitutes corroborating evidence (to use Popper's principle) for *non-mechanistic* social and psychological theories: for theories that endow human beings with free will and unbounded creativity.

Mechanistic theories fail because they do not recognize the unique knowledge structures that can develop in a mind. Thus, the failure of these theories ought to enhance our respect for the potential of human beings, for the creativity of each individual. Instead, the scientists insist that these are not failures but merely setbacks, that they *will* eventually find mechanistic theories of mind and society.

Our mechanistic culture has given rise to this incredible spectacle: in a democratic society, in the name of science, renowned professors working in prestigious universities believe it is their duty to prove that human beings have no value. For, by denying the bearing that the knowledge structures present in our minds have on the structures studied by their theories, these scientists deny the unique contribution that each individual can make. By claiming that human phenomena can be explained with mechanistic theories, they claim in effect that these phenomena can be explained without taking into account the knowledge structures developed by individual minds.



And this is not all. Although these theories do *not* work, the mechanists use them to draw sweeping conclusions about man and society – the kind of conclusions that one would draw if the theories *did* work. Specifically, they maintain that human freedom and creativity are only illusions, prescientific notions, not unlike the ancient beliefs that the earth is the centre of the universe, or that Man was created in God's image. Hence, just as science has shown that the earth is merely another planet, and that Man is merely a higher animal, we must trust science again and resign ourselves to the fact that we

cannot be truly creative: everything we do is dictated by our genetic structure, or by our environment, or by other factors over which we have no control. Human beings are in reality nothing but machines – complicated ones perhaps, but machines nevertheless.

Thus, Skinner could only confirm his behaviourist theory in contrived laboratory experiments with rats and pigeons, but concluded that there is an urgent need to apply this science of behaviour to the shaping of human minds and societies. All human acts are the result of external influences, he says, and it is a mistake to believe that we are free and responsible agents. So, rather than allowing ourselves to be controlled by whoever has the power to influence us – parents, teachers, friends, advertisers – it is best to allow a hardheaded government and expert behaviourists do that. These elites would use objective principles and rigorous methods to shape the personality of each individual starting from birth, and thereby create a society of perfect citizens: “What we need is a technology of behavior... But a behavioral technology comparable in power and precision to physical and biological technology is lacking.”¹ This is where the science of behaviourism can help: the conditioning techniques that seem to work for the rats and pigeons trapped in a Skinner box in a laboratory must now be used for the people that make up modern society.

What prevents us from creating this progressive system is our democratic prejudices; that is, our naive belief in human freedom and dignity – notions that the science of behaviourism, according to Skinner, has shown to be illusory anyway: “The conception of the individual which emerges from a scientific analysis is distasteful to most of those who have been strongly affected by democratic philosophies.”² Skinner was so confident in the potential of behaviourism to solve our social problems that he wrote a science-fiction novel to depict the kind of society we could create through “behavioural technology.”³

It is significant that Skinner’s ideas were very popular and became somewhat of a cult in America, especially among intellectuals. It is also significant that most of those who *rejected* Skinner’s utopia did so because they found his behavioural technology objectionable on humanistic, not scientific, grounds: how outrageous that a professor from Harvard University is promoting totalitarianism. Few realized that the first objection to behaviourism must be that it is a pseudoscience, that it does not work, that it is founded on fallacious concepts of mind and society. And what ought to be outrageous is that our universities foster the corrupt environment where pseudoscientists like Skinner can peddle their theories.

¹ B. F. Skinner, *Beyond Freedom and Dignity* (New York: Knopf, 1972), p. 5.

² B. F. Skinner, *Science and Human Behaviour* (New York: Free Press, 1965), p. 449.

³ B. F. Skinner, *Walden Two* (New York: Macmillan, 1948).

The structuralist theories work no better than the behaviourist ones, but their defenders do not hesitate to conclude that human freedom and creativity, in the sense in which we generally understand them, are mere illusions. When we acquire skills and knowledge, when we invent something or solve a problem, when we develop social customs and institutions, all we do in reality is select various acts from a predetermined range of alternatives – the range for which our brains are biologically wired.

Thus, Lévi-Strauss held that the set of possible social customs is analogous to the periodic table of chemical elements: all a society does when adopting a certain custom is select, perhaps randomly, one of the slots available in this table. And Piaget held that the mental development of children is analogous to an increase in the number of levels in a hierarchical structure of binary operations – the structure which is built into the human brain, and which, ultimately, determines our mental capabilities. As individuals or as societies, human beings can be no more creative or free than programmed computers. What we like to think of as creativity and free will is only an illusion caused by the large number of available alternatives.



The case of Chomsky and his universal grammar is especially interesting, because Chomsky himself draws attention to the harmful influence that theories of mind can have on social and political ideologies. He stresses that the innateness hypothesis behind his linguistic theory postulates a view of human nature in the *rationalist* tradition. Rationalist philosophers, starting with Descartes, held that we possess certain mental capabilities simply by being born human; and, although we acquire much knowledge later, our innate capabilities restrict and structure forever what we can know. The rival philosophical school of *empiricism*, on the other hand, holds that human minds are empty at birth, that everything we know comes from interacting with our environment, and that there are no innate restrictions on the kind of knowledge we can acquire.

Chomsky points out that the rationalist view is conducive to ideologies that defend freedom, equality, and respect for the individual, whereas the empiricist view is conducive to ideologies that support authoritarianism, inequality, and exploitation. Specifically, if human beings are empty organisms at birth, as the empiricists say, this means that by their very nature they have no rights; so there is nothing wrong in moulding them to fit a certain policy. Thus, theories like behaviourism, ideologies like Nazism and Communism, and even democratic systems where various elites are permitted to control society, demonstrate the danger of the empiricist view: the leaders can invoke

the idea of human nature to justify the control of knowledge. Rationalists, on the contrary, respect the fundamental rights of the individual – the right to live free, to develop any personality, to pursue any lifestyle – simply by recognizing that human beings possess from birth some important and immutable faculties.⁴

Accordingly, says Chomsky, while rationalism appears to postulate a more limited view of human capabilities, it is in fact the one philosophy that defends individual freedom and creativity. He admits that the hypothesis of an innate language faculty restricts the types of languages that humans can acquire, and the types of sentences – and hence also the types of ideas – that they can create; and he admits that, if everything we can know is governed by innate faculties, similar restrictions apply to all other kinds of knowledge. But, he reassures us, we needn't worry that these restrictions limit our creativity or freedom, because we still have a very large number – indeed, an infinite number – of alternatives *within* the boundaries of innate capabilities. We can create an infinite number of sentences, for instance, despite the severe restrictions imposed by universal grammar.

Chomsky's thesis, however, is fallacious and dangerous. We can agree with him that the concept of empiricism has been distorted and abused by certain ideologies. And, even without accepting his hypothesis of innate faculties, we can agree that our mental capabilities are structured and restricted by certain biological characteristics. But these facts do not warrant his conclusions.

Chomsky's mistake is to assume that, if our mental capabilities lie within a certain range, we should be able to discover a deterministic model that accounts for all possible human acts (because these acts necessarily derive from that range of capabilities). His mechanistic theory of mind compels him to *degrade* the definition of creativity: from the capacity to perform *unpredictable* acts, to the capacity to *select* an act from a predetermined range of alternatives. The traditional view is that creativity gives rise to an infinity of alternatives, and, in particular, to alternatives that were not known in advance. Chomsky, though, believes that we can *account* for these alternatives – simply by inventing a deterministic model that generates an infinity of sentences, ideas, and types of knowledge. But the infinity displayed by a deterministic model is only a fraction of the *real* infinity of alternatives that human minds can develop. (This, obviously, is why his theory doesn't work.)

Chomsky speaks eloquently of human freedom and creativity, but at the same time he attempts to determine with precision all the manifestations of creativity. He seems oblivious to the self-contradiction. For, if there were ways

⁴ See, for example, Chomsky's *Language and Politics* (Montréal: Black Rose Books, 1988), pp. 594–595, and *Reflections on Language* (New York: Pantheon Books, 1975), pp. 127–133.

to account for all possible human acts, it would be absurd to call the quality involved in performing these acts “creativity.” No surprises would be possible – no exceptions, no novelty, no originality. Anything an individual would do could be shown to be derivable independently of that individual. To put this differently, if Chomsky’s theory worked, we could implement it with software; a computer would then perform exactly the same acts as human beings (would generate, for example, the same sentences and ideas), but could the computer be said to be free or creative in the human sense? Determinism is the opposite of freedom, but Chomsky wants to have both: Chomsky the humanist is concerned with freedom, while Chomsky the scientist is searching for a theory that would make a mockery of freedom by showing that a programmed machine can be identical intellectually to a free human being.

Just like the mechanistic theories of mind in the field of artificial intelligence, Chomsky’s theories are, in effect, an attempt to replace human minds with software. And, with microprocessors becoming more and more powerful, some of these theories can already be implemented with just one semiconductor chip. They may not state it, and they may not even realize it, but what all these researchers are claiming, essentially, is that we will soon be able to replace human beings with inexpensive devices. The important point, again, is that although these theories do *not* work, the researchers, and the lay people who trust them, are convinced that soon they *will* work, and are therefore developing a world view that reflects these theories. Nor should we forget that our society is already dominated by political and business elites who hold the same conviction, and who are planning our future accordingly. It is not difficult to imagine the kind of future these elites are preparing for us, if they believe that human beings are not very different from expendable semiconductor chips.

In conclusion, Chomsky’s preference for a rationalist theory of mind, rather than an empiricist one, is irrelevant when rationalism is supplemented with a *mechanistic* theory of mind. It makes little difference which philosophy one starts with, if one ends by claiming that deterministic models of mind are possible.

Chomsky’s case, then, is a good example of the corruptive effect of the mechanistic dogma. Through writings and lectures, he has become known throughout the world as a humanist. His efforts as a scientist working in the mechanistic tradition, however, are harming the humanistic cause more than his efforts as a humanist can help it.

Geoffrey Sampson⁵ notes that Chomsky’s impoverished definitions of freedom and creativity provide the common philosophical foundation for both

⁵ Geoffrey Sampson, *Liberty and Language* (Oxford: Oxford University Press, 1979).

his linguistic and his political theories. As language users, Chomsky says, we are restricted genetically to certain types of grammatical constructions; hence, the creativity we display in speech is in reality only the capacity to select utterances from a certain range of alternatives. Similarly, as citizens, we are restricted genetically to certain types of achievements; hence, there is nothing wrong in defining freedom as merely the right to pursue any ideas within a prescribed range of alternatives: “Chomsky has misappropriated the term ‘creative’ as he misappropriated the term ‘free.’ In each case he uses the term in a sense that conflicts with its standard usage; but, by contrasting ‘freedom,’ or ‘creativity,’ in his impoverished sense with something that is even further removed from the notion usually associated with the respective word, he invites us to overlook the fact that what we usually mean by the word is something different from both his alternatives.”⁶

Chomsky contrasts his theories to those of authoritarian ideologies, which deny freedom and creativity altogether, and which hold that human nature must be moulded to fit a general plan. When presented from this perspective, his own views appear enlightened and liberal. His theories are dangerous precisely because he appears to be defending freedom while defending, in fact, not the traditional concept of freedom, but an impoverished version of it: “The adverse consequences of scientism stem from its assumption that all human phenomena can be analysed by the scientific method; creativity is an exception, since acts which are truly creative cannot, by definition, be predicted. To the question ‘Who, in the contemporary intellectual world, most stresses the importance of human creativity?’, the answer must undoubtedly be Noam Chomsky.... Yet, when we ask what Chomsky means when he calls men creative, he turns out to refer to our ability to behave in conformity to certain fixed, rigorous rules.”⁷

2

By way of summary, I want to show how the mistaken conclusions drawn by these scientists can be traced to the mechanistic fallacies. The scientists are evidently fascinated by the fact that simple structures, when used as mechanistic models, can generate a large number of different values for the top element. They can generate, in fact, an *infinite* number of values. Depending on the theory, these values represent the different alternatives displayed by individuals or societies in their knowledge, their behaviour, their traditions, their sentences, etc. So, the scientists conclude, mechanistic models can

⁶ Ibid., p. 106.

⁷ Ibid., pp. 107–108.

account for *all* possible alternatives: all knowledge and behaviour of an individual, all customs and language uses of a society, and so forth.

But a large number of values – even an infinite number – does not necessarily mean that the model can account for *all* possible alternatives. We can demonstrate this with a simple example. It is easy to create a device (a piece of software, for instance) that generates numeric values according to certain rules: values limited to a given range, or values that are prime numbers, or integers. Consider now the set of all numeric values. Although the subset of integers is an infinite number, in practice there are infinitely more fractions than integers (there are, in fact, an infinite number of fractions between any two integers). Consequently, while accounting for an infinity of values, a device that generates integers accounts for an infinitely small subset of all possible numeric values. There are many similar examples, so it is important to bear in mind that an infinite number of alternatives may not mean *all* the alternatives.

What these scientists are seeking is a fully specifiable model that can account, nevertheless, for all the alternatives displayed by human minds. They hope to find, in other words, a deterministic model that can account for indeterministic phenomena – for the creativity and unpredictability of human acts. They are misled by the infinity of alternatives that their mechanistic systems can generate, and conclude that they have discovered such a model. They misinterpret this infinity of alternatives as equivalent to creativity and unpredictability – equivalent, that is, to all possible alternatives. As we just saw, it is easy for a mechanistic system to generate an infinity of values in a given domain while failing to account for all values possible in that domain. Thus, the infinity that the scientists notice in their models is *not* the infinity that gives rise to indeterminism, to creativity and unpredictability. Mechanistic theories fail because, even though explaining an infinity of acts, this infinity is a small subset of all possible human acts. (See also the related discussion in chapter 8, pp. 814–817.)

The scientists start with a reasonable hypothesis: the idea that human beings are restricted in the types of knowledge they can acquire, and in their behaviour patterns, by some innate capabilities. Now, no one can doubt that the basic human faculties *are* bounded by some low-level physiological processes occurring in the brain. The high-level phenomena of mind and society must then reflect these limitations, just as our physical characteristics and abilities reflect our genetic code.

But, to be consistent, the scientists ought to build their theories starting from those low-level physiological processes (from neurons, for instance). Instead, the starting elements in their structures are relatively high-level entities: for linguists they are phonemes and words; for behaviourists they are simple movements, reflexes, and the like; for structuralists they are the binary

opposites found in thought, or in stories and traditions. These are not the simplest entities from which human phenomena are made up, but merely the smallest entities that we notice or understand. The scientists believe that they can start their projects from any level of abstraction they like, so they choose high levels, which are more convenient. They base their theories on mechanistic processes assumed to occur at low physiological levels; but then they ignore the low levels, and the many intermediate levels, and build their models starting from some arbitrary, relatively high-level, entities.

Their mistake, thus, is to use as building blocks in their models, elements that are not independent. Elements like words, or limb movements, or pieces of mental logic, are related – both mutually and to other *types* of elements; so they give rise to complex, not simple, structures. They are related because they serve as elements in other structures too, besides the particular structure that each scientist is concerned with. All these structures are formed at the same time from the low-level physiological elements, so they reflect the countless interactions that take place at levels *lower* than the level where the scientists decided to start their projects.

When ignoring the interactions – when assuming that those starting elements are independent – the scientists separate the one structure that forms their particular project from the complex structure that is the human phenomenon. The top element in the complex structure represents all the alternatives that human beings can display; and the theory fails because it can account for only *some* of these alternatives – the ones that would occur if those starting elements were indeed independent. The theory may well account for an infinity of alternatives, but this is still a small fraction of the alternatives that constitute the actual phenomenon. When attempting to represent a complex phenomenon with a simple structure, and when starting from higher levels, the scientists are committing both fallacies, reification and abstraction. The dramatic reduction in alternatives is then the impoverishment that these fallacies inevitably cause.

We can understand now why these scientists are convinced that freedom and creativity are illusory, that what we perceive as free will is merely the freedom to select any act from a predetermined range of alternatives. They base this notion on the existence of low-level physiological elements, and on the assumption that, *in principle*, we can account for all the alternatives generated by these elements. At the same time, they admit that they *cannot* develop their models starting from these elements. So they continue to claim that freedom and creativity mean just a selection of alternatives, even while starting their models from much higher levels – as if the infinity of alternatives they *can* account for were the same as the infinity generated by the low-level elements. They are seeking mechanistic explanations, but they are *not* following the

mechanistic doctrine: a model cannot explain mechanistically a given phenomenon unless the starting elements are atomic and independent, and *their* starting elements are neither. (Even the low-level physiological elements, which they *fail* to reach, are only *assumed* to be atomic and independent.)

The scientists believe that it will soon be possible to explain all the alternatives and thereby account for the full range of possible human acts. But this optimism would be warranted only if their theories were completely mechanistic. They admit that they cannot find a true mechanistic explanation – a continuous series of reductions from the human phenomena to some atomic and independent entities – but they refuse to take this failure as evidence that human phenomena are indeterministic, and hence unexplainable with mechanistic theories.

Thus, no matter how we feel about mechanism and about mechanistic explanations of human phenomena, the conclusions drawn by these scientists are unjustified simply because their theories are fallacious even *within* the mechanistic doctrine. The indeterminism, the creativity and unpredictability of human minds, are precisely this impossibility of finding a set of atomic and independent entities for the starting elements.

The Software Theories

1

The evil concepts of mind and society engendered by the traditional mechanistic theories are being implemented already, without anyone noticing, through the mechanistic *software* theories. While our democratic system has safeguards to protect society from the political consequences of the traditional mechanistic delusions, no safeguards exist to protect us from the ideas promoted by the software elites. And few of us recognize the similarity between the social systems envisaged by the traditional mechanists, which we reject as totalitarian, and those promoted by software experts and by the software companies, which we perceive as scientific and progressive.

Recall the distinguishing characteristic of the mechanistic delusions we studied in this chapter: the claim that, in a complex structure, it is possible to derive with precision the value of the top element from the values of the starting elements. The complex structure represents a complex human or social phenomenon, so what the mechanists are claiming is that a simple structure – a deterministic model – can account for all possible manifestations of the complex phenomenon. They are claiming, in other words, that the other structures that are part of the phenomenon are unimportant. But the other

structures are the knowledge structures present in human minds, or the social structures formed by the interaction of many minds. So, by making this claim, the mechanists are depreciating the role played by each mind in the complex phenomenon. The reason these theories are failing is that the assumption is wrong: the contribution made by individual minds is in reality an important part of the phenomenon.

In the case of software theories, the complex structures comprise the various activities performed by people when developing and using software applications. And what the software theories claim is that it is possible to account for the top element of these structures from a knowledge of their starting elements. Accordingly, we should be able to replace the knowledge involved in those activities with formal methods and, ultimately, with software devices based on these methods.

In programming activities, the starting elements are details like the definitions and statements used in an application, the fields in data files or display screens, and the software-related acts in the human environment where the application is used. The top element consists of the combinations of operations performed by these applications, and the business, social, and personal needs they serve. The knowledge and experience of programmers and users provide the structures that, together with the structures formed by the elements just mentioned, give rise to the complex phenomena observed when a social or business environment depends on software.

The software theories, though, claim that these are not complex, but mechanistic, phenomena: the top element can be described, precisely and completely, as a function of the starting elements. Whether they invoke mathematical principles or the similarity of programming to manufacturing, the appeal of these theories is understandable. Let us design our applications as hierarchical structures, we are told, as neat structures of independent modules. We should then be able to develop applications of any size and complexity simply by combining these software parts, one level at a time: starting with some atomic entities whose validity is established, we will create larger and larger software subassemblies – each one guaranteed to be correct because built from proven parts – until we reach the top element, the complete application.

All theories, and all methodologies and development environments, are grounded on this principle. But the principle is invalid, because even the smallest software elements share attributes, and are therefore interrelated. Thus, like the mechanistic theories of mind and society, the software theories praise reductionism and atomism, and at the same time they violate these principles by using starting elements that are not atomic and independent. Also like the traditional theories, the software theories are failing because

these elements give rise to multiple, interacting structures. Applications, and software-related activities generally, are not the simple hierarchical structures the mechanists assume them to be. (We will study the various types of software structures in chapter 4.)

2

Psychologists, sociologists, and linguists think that mechanistic theories can account for all the alternatives displayed by individuals and societies: all possible manifestations of knowledge, behaviour, customs, language, and so forth. Their theories account for only *some* of the alternatives, but despite these failures, the scientists conclude that what human minds do is simply select various acts from a predetermined range of alternatives: what we perceive as free will or creativity is an illusion caused by the large number of alternatives that minds can select from. Let us see how the mechanistic *software* theories lead to the same conclusion.

Like the mind mechanists, the software mechanists are searching for a deterministic model that can account for indeterministic phenomena – software-related phenomena, in this case. Specifically, they hope to find a mechanistic model that can account for all the alternatives displayed by human minds when engaged in software development and use. And they think that if a model can account for an infinite number of alternatives, this means that it can account for *all* the alternatives.

But we saw how easy it is for a mechanistic model to display an infinity of values while accounting, in fact, for only a small subset of the possible values. The alternatives in software-related phenomena are all the applications that can be implemented with software, and all the business, social, and personal aspects of programming and software use. These alternatives constitute an infinite number, of course. And, as is the case with the other theories, the infinity of alternatives that make up the real software phenomena is infinitely greater than the infinity of alternatives that a mechanistic model can account for. The difference is seen in the perpetual need for new programming theories – a need arising, obviously, from the failure of the previous ones.

While the mind mechanists attempt to represent with exact models *mental* phenomena, the software experts attempt to represent with exact models *software-related* phenomena. This must be possible, they argue, because the totality of software-related human acts can be described with precision as a function of some low-level elements. We should be able to build software devices, therefore, which incorporate these elements. These devices will then allow us to generate any software-related structure – that is, any alternative of

the phenomena of programming and software use – without ever again having to start from the low levels.

True to the mechanistic doctrine, the software experts attempt to explain complex software phenomena by repeatedly reducing them to simpler ones. And in so doing they commit both fallacies, reification and abstraction: they take into account only one structure, failing to see that software entities are related in several ways at the same time; and they stop the reduction long before reaching the lowest levels.

The experts notice that in processes like manufacturing we benefit from the mechanistic principles while starting our structures, nevertheless, from high-level elements (prefabricated subassemblies), and they conclude that software-related processes, too, can start from high levels. They also notice that physical entities can function as starting elements only if independent, and they conclude that software entities, too, can be independent: just as, in a physical structure, the internal properties of one subassembly are unrelated to those of the others (so that we can build them independently of one another), the internal operation of each software module can be unrelated to those of the others.

Both conclusions, however, are unwarranted. While it is true that we can build software structures starting from independent, high-level elements, if we limited ourselves to such structures we could represent with software only *a fraction* of our business, social, and personal affairs.

The experts treat software development as a manufacturing process because they don't appreciate how much richer are human phenomena than physical ones. The number of alternatives lost when we start from high levels and when we separate structures in *physical* phenomena is relatively small, and we gain important benefits in return; when we ignore the low levels and the links between structures in *human* phenomena, we lose an infinity of important alternatives.

Simple structures and high-level starting elements – that is, the use of standard parts and subassemblies – are acceptable in activities like manufacturing because there are only a few kinds of tools, appliances, vehicles, etc., that are useful, or convenient, or economical. The same is not true, however, in mental and social phenomena. There is an infinity of sentences, ideas, customs, cultures, forms of knowledge, and types of behaviour that are correct, or practical, or suitable. Unlike physical processes, therefore, our models of human phenomena cannot be restricted to simple structures, and cannot start from high levels. For, the loss of alternatives is then so severe that the benefits of simple structures and high starting levels become irrelevant: the number of human acts the model can explain, relative to those it cannot explain, is too small for it to be useful.

Software-related phenomena, in particular, despite their dependence on physical structures like computers, are largely *human* phenomena, because they entail various intellectual and social processes. Consequently, simple software structures and high-level starting elements can account for only *some* of the alternatives, just as the mechanistic models of mind and society can explain only *some* aspects of human phenomena.

3

The immediate benefit of software mechanism is thought to lie in explaining the phenomenon of programming itself. If we view as hierarchical structures not just the applications but also the activities and the mental acts involved in developing applications, we should be able to account for all the values displayed by the top element; that is, all the combinations of applications, requirements, software concepts, etc. We will then search for some *high-level* entities that can be used as starting elements in these structures. And, once we discover these entities, we will no longer have to develop applications starting from low levels.

Everyone agrees that it is possible to develop all conceivable software applications if starting with *low-level* elements. The promise of software mechanism, thus, is not to enable us to perform tasks that we could not perform otherwise, but to perform them more easily: less work and lower skills are needed to reach the top element – the application, the business system – when starting with higher-level elements. So what software mechanism promises us in effect is devices that would permit us, not only to attain everything that can be attained through software, but to attain it sooner, and with less knowledge. This is true, we are told, because the infinity of applications possible when starting from the higher levels is about the same as the infinity we could create by starting from the low levels.

It is not surprising that the scientists make such fantastic claims; after all, these ideas are nothing but the software counterpart of the theories of mind that mechanists have been proposing for centuries. What *is* surprising is that we all accept the claims now, and that we continue to accept them even as we see them refuted in practice.

A little thought will reveal the true nature of these claims. We are told that software devices will permit us to create *any* software applications, or to address *any* software-related matters, *without* the need to develop all the knowledge that human minds *can* develop. So, by stating this, the experts admit in effect that we *can* attain greater knowledge; we simply *do not need* that extra knowledge. But if we *are* capable of greater knowledge in software-

related matters, we necessarily could, through that knowledge, perform certain tasks, or develop certain ideas, or create certain applications, or make certain discoveries, that we cannot *without* it. To say that those additional alternatives are unimportant, that they represent an insignificant part of the potential achievements of human minds, and that they can be forsaken, is to make an astonishing statement about the future of humankind. For, we can only judge how important or unimportant those alternatives are by *having* them first; we cannot know in advance how human knowledge will evolve, or what our future needs and capabilities will be.

By making such statements, the software elites are claiming in effect the right to decide what knowledge and mental capabilities we, and future generations, are permitted to acquire. They see themselves as an enlightened vanguard: they are the select few who can appreciate the future of software, so it is their duty to guide us.

It is in the interest of the elites to maintain a general state of ignorance in all software-related matters, and the mechanistic software ideology is an essential part of this plan: we perceive the theories, methodologies, and devices as expressions of software science, even as they are preventing us from using our minds. By controlling the way we create and use software, the elites are restricting those aspects of our life that depend on software to a certain range of alternatives. But few of us realize just how narrow this range is, how many alternatives are lost when we are limited to mechanistic software thinking.

We would have no difficulty recognizing a similar limitation in areas in which we are already knowledgeable – other professions, or naturally acquired skills. Programming, however, was taken over by incompetents and charlatans before a body of responsible professionals could emerge; and as a result, we believe that what these impostors are doing is the utmost that a society can attain in the domain of programming. Since the programming we see is the only kind we have ever had, we cannot know how limited and inefficient our software-related affairs are. We cannot even imagine a society where programmers are as competent in their work as other professionals are now in theirs; that is, a society where programmers are permitted to attain the highest skill levels attainable by human minds.

The mechanistic software ideology, thus, fosters incompetence among programmers – but also among software *users*, because their performance is impaired by the inadequate applications. These conditions, then, impoverish our life by limiting our expectations in all software-related matters to only a fraction of the possible alternatives: those that can be accounted for through mechanistic concepts. Although our minds are capable of non-mechanistic knowledge, and we could therefore have infinitely more alternatives in our software-related affairs, as long as we limit ourselves to mechanistic thinking

we cannot know what these alternatives are. We trust the software devices provided by the elites, while the real purpose of these devices is to induce a state of ignorance and dependence. We like the many alternatives that we gain so easily through devices, and we are convinced that they constitute *all* the alternatives. We are unaware of the *missing* alternatives, because the only way to attain them is by starting our software-related structures from lower levels; that is, by *avoiding* the devices, and developing instead programming expertise.

Again, we can easily understand this for other types of knowledge, in areas in which we have had the time to become proficient. We all use *language*, for example, starting with low-level elements – with words. We also know that, when expressing wishes or ideas, we do more than just build linguistic structures; what we do is combine our linguistic knowledge with other types of knowledge. Hence, if we have difficulty expressing ourselves, we know that language devices would not solve the problem.

Imagine now a society where an elite suggested that we use language by starting with higher-level elements – with ready-made sentences and ideas. Thus, instead of *creating* sentences and ideas, we would have to express ourselves by combining the ready-made ones produced by certain devices; and instead of increasing our knowledge and our linguistic capability, we would learn only how to operate these devices. Even without recalling the mechanistic fallacies, or the difference between simple and complex structures, we sense intuitively that this could not work, that we would lose something. We sense that, even if we could still create an infinity of ideas, it would be impossible to create the same variety of ideas as we do now, starting with words. And we realize that, by restricting the use of language, the elite would impoverish all knowledge and all aspects of our life.

We accept in the domain of software, thus, mechanistic concepts and theories that we would immediately dismiss in the case of language. We accept them because we fail to see that they serve the same purpose: to restrict our knowledge, our values, and our expectations. (We will return to this subject in chapters 5 and 6.)

4

The mechanistic theories of mind, we saw, claim that we can account for all the alternatives in human knowledge, behaviour, language, customs, and so forth; and they conclude that no creativity is involved in mental acts, that what we perceive as creativity is merely the selection and combination of bits of knowledge leading to a particular alternative. In the mechanistic *software* theories, the counterpart of this belief is the belief that we can account for all

possible alternatives in software-related phenomena. If we can account for all the alternatives, the conclusion must be that we can account for all possible applications, and also for the processes that take place in the mind when creating applications. So what we perceive as creativity and originality in the work of an experienced programmer is only an illusion: what the programmer really does in his mind at every stage of software development is akin to selecting and combining bits of programming knowledge in order to generate one of the applications.

More accurately, what the software theorists say is that programmers *could* create applications in this manner, so it is wrong to depend on such inexact resources as personal knowledge. Professional programming entails nothing but rules, standards, principles, and methods. Only old-fashioned practitioners rely on unscientific notions like skill, experience, or intuition. In other words, the theorists say, programming phenomena are deterministic: practically all the work involved in developing a particular application, or in modifying it later, can be specified precisely and completely from a knowledge of the requirements – just as the steps needed to build a car or an appliance can be specified in advance from a knowledge of their physical characteristics.

So, if we can account for all software applications that human minds can create – if there is nothing indeterministic, or unpredictable, in what programmers do – we should be able to replace programming knowledge with devices. The devices known as development environments, for example, materialize this idea by providing high-level starting elements. They simplify programming by minimizing the number of levels between the starting elements and the final application. Programming expertise, clearly, is becoming redundant, since even novices can now create applications. It is also interesting to note the trend in these devices to reduce the programmer's involvement to a series of selections, and selections within selections. The process of programming – the definition of files and variables, the specification of operations and conditions – has given way to a process of *selections*: instead of a blank screen where we can enter our definitions and specifications freely, we find colourful screens replete with menus, lists, icons, buttons, and the like. This structure of selections within selections attempts to emulate, obviously, the hierarchical structure believed to constitute the mental part of programming: the selection and combination of bits of knowledge leading to one of the alternatives, to a particular application.

The devices, thus, are an embodiment of the software theories. And, since the software theories are the counterpart of the traditional mechanistic theories of mind, we are witnessing the actual implementation of the idea that human creativity is an illusion, that the mind works simply by selecting and combining bits of knowledge within a predetermined range of alternatives. The

substitutes for human intelligence – the models of mind that the mechanists were proposing all along in psychology, sociology, and linguistics – have finally found their practical expression in the world of programming. While the traditional mechanists were content with speculations and theories, the software mechanists are actually replacing human intelligence with devices.



We have already discussed the consequences of mind mechanism: by claiming that all human acts can be explained – by claiming, therefore, that human beings are incapable of truly creative, or unpredictable, acts – the traditional mechanistic theories can lead to a society where freedom, expertise, and creativity are redefined to mean the selection of acts from the range of alternatives sanctioned by an elite. In the world of programming, clearly, the shift has *already* occurred: the software theories assume that programmers are capable of nothing more creative than using programming substitutes, and this has given rise to the belief that it is these substitutes that are important, not the individual minds. This, in turn, has led to the belief that we must depend on the companies which provide the substitutes, and that the only way to improve our programming practices is by constantly adopting new versions of the substitutes.

Both the traditional theories and the software theories, thus, lead to the belief that human intelligence can be replaced with deterministic models. But, whereas traditional mechanism has had so far no serious consequences outside academia, *software* mechanism is causing great harm – by distorting our conception of knowledge and skills. Like the traditional ones, the theories behind the programming substitutes have been repeatedly refuted. Yet, while the theories themselves are failing, our belief in substitutes has become a self-fulfilling idea: Because we assume that programmers cannot advance beyond their current level, we encourage them to depend on substitutes. Consequently, no matter how many years of experience they have, their programming skills remain at the same low level – the level needed to use substitutes. Given this state of affairs, the adoption of the next substitute always appears to be the only way to improve their performance. So they waste their time assimilating yet another theory, or methodology, or language, instead of simply programming and improving their skills. And so the fallacy feeds on itself.

Human beings, as a matter of fact, *can* attain higher skill levels. It is only because programmers are forced to stay at novice levels that their capabilities do not exceed those of the substitutes. Recall the process of skill acquisition we studied in “Tacit Knowledge” in chapter 2: the skills and knowledge of a novice programmer constitute *simple* structures, just like the structures on which the

mechanistic software theories are based; only when attaining expertise can the programmer's mind process the *complex* structures that make up software phenomena. Expert programmers outperform the substitutes because their minds can cope with whole software phenomena, while the substitutes can only cope with isolated aspects of them.

So it is the dependence on programming substitutes that is preventing the emergence of a true programming profession. Application programming has been so degraded that it now entails nothing more difficult than the performance of acts which almost anyone can master in a year or two. We are underrating the potential of our minds, and we have forgotten that there exist levels of knowledge and skills higher than those reached after a year or two of practice. In effect, we are degrading our conception of human intelligence to the level attainable by the mechanistic substitutes for intelligence.

5

We saw how the mechanistic ideology has affected the programming profession. The most serious consequence of software mechanism, however, is not the destruction of *programming* knowledge, but the destruction of the *other* kinds of knowledge – knowledge that took centuries to develop. The belief that intelligence can be replaced with devices is already spreading into other domains, especially in our business-related pursuits.

In one activity after another, we see the claim that it is possible to account for all the alternatives in that activity, discover the starting elements that give rise to these alternatives, and then incorporate the elements in a software device. In all activities, we are told, what the mind really does is combine some elementary bits of knowledge into knowledge structures, one level at a time, just as we build cars and appliances. By providing directly some high-level starting elements – prefabricated knowledge subassemblies, as it were – the devices eliminate the need for each one of us to develop in our minds the low-level elements and the combinations leading to the high levels. The lengthy and arduous period of learning and practice is a thing of the past: all we need to know now is how to operate software devices. In one occupation after another, we are told that it is unnecessary, even wrong, to rely on personal experience. Anyone can perform the same tasks as the most experienced worker simply by selecting and combining the ready-made, high-level elements available through these devices.

The devices known as office productivity systems, for example, which address office workers, are the counterpart of the development systems used by programmers. Instead of claiming that we can account for all the alternatives

in programming activities, the software companies claim now that we can account for all the alternatives in *business-related* activities. Instead of software concepts that programmers can use as substitutes for programming expertise, the office systems promise software concepts that replace the skills of office workers.

And here too we see the trend to provide these concepts in the form of selections, and selections within selections, rather than permitting the user to combine freely some low-level elements. Here too, therefore, the system attempts to emulate what are believed to be the mental acts of an experienced person: selecting and combining hierarchically, one level at a time, the bits of knowledge leading to one of the alternatives.

As in the case of programming systems, the fact that an office system can generate an infinity of alternatives in business-related matters is mistaken as evidence that these devices can generate *all* possible alternatives. But this infinity is only a fraction of the infinity that an experienced worker can implement by starting from low levels, through the complex knowledge structures developed in his mind. Software devices cannot replace business-related knowledge any more than they can programming knowledge.

The fallacy, again, lies in providing starting elements that are not atomic and independent. We are lured by these devices only because we forget that the knowledge required to perform a difficult task constitutes, not an isolated structure, but a system of interacting structures: it includes many knowledge structures besides the neat hierarchical structure of selections that can be embodied in a software device. The different values of the top element – values representing the various acts performed by an experienced mind – are determined largely by the *interactions* between structures. The most important interactions occur at levels lower than the starting elements provided by devices, because it is low-level elements like variables and individual operations that are shared by the software structures, business structures, and knowledge structures which together make up the phenomena of office work. These interactions are lost when replacing minds with devices, and this is why a device can display only *some* of the alternatives displayed by a mind – those alternatives that can be represented with isolated simple structures.



To conclude, I want to stress again the link between the mechanistic theories of mind and the delusion of software devices. If we believe that it is possible to account for all the alternatives displayed by human beings in their mental acts, we will necessarily conclude that it is possible to describe human intelligence as a function of some starting mental elements; that the creativity of human

minds is an illusion; and that everything the mind does can be explained as we explain the working of a machine. It should be possible to represent a person's mind, therefore, with a simple structure where the elements and levels correspond to the knowledge developed by the mind, and the values of the top element correspond to the various acts performed by that person.

What is left is to implement this structure by means of a software device. Then, any person operating the device will be able to perform the same acts as the person whose mind the device is emulating. All that human beings do, in reality, is operate devices. So, why depend on a device like the mind, which is inefficient and unreliable and, moreover, can only develop knowledge structures through painstaking learning and practice, when we can purchase modern software devices that already contain these structures?

In the domain of programming, we have already replaced minds with devices. Now, as our reliance on computers is growing, the software elites are degrading all activities where human knowledge and skills play a part, in the same way they have degraded the activity of programming. They are modifying our conception of knowledge and skills to mean simply a dependence on software devices. They are instilling in us the belief that our intelligence, our work, our initiative, our experience, can be reduced to the process of selecting and combining operations within the range of alternatives provided by these devices. They are shifting our definition of expertise, creativity, and responsibility from their traditional meaning – to do a good job, to solve an important problem, to make a real contribution – to merely knowing how to use the latest software devices.

