Collection
Károli

Monograph

Marcus Kracht

Knowledge and Material Culture

How much knowledge can we afford



KRE | KÁROLI GÁSPÁR REFORMÁTUS EGYETEM

L'Harmattan

Ez a mű a Creative Commons Nevezd meg! - Ne add el! - Így add tovább! 4.0 Nemzetközi Licenc feltételeinek megfelelően felhasználható.
https://creativecommons.org/licenses/by-nc-sa/4.0/



Knowledge and Material Culture. How much knowledge can we afford?

Marcus Kracht

https://orcid.org/0009-0007-9633-9955

Cultural history, cultural memory / Kultúrtörténet, kulturális emlékezet (13056)

knowledge preservation, knowledge transmission, endangered knowledge

ISBN 978-963-646-546-9 (pdf)

DOI: https://doi.org/10.56037/978-963-646-546-9



https://openaccess.hu/

MARCUS KRACHT

KNOWLEDGE AND MATERIAL CULTURE

L'Harmattan Hongrie

Collection Károli

Collection dirigée par Enikő Sepsi

ISSN 2063-3297

MARCUS KRACHT

KNOWLEDGE AND MATERIAL CULTURE

HOW MUCH KNOWLEDGE CAN WE AFFORD?

Károli Gáspár University of the Reformed Church in Hungary
L'Harmattan Publishing • Éditions L'Harmattan

Budapest • Paris
2020

Publishing Director: Géza Horváth, Ádám Gyenes, Xavier Pryen

Series Editor: Enikő Sepsi

Károli Gáspár University of the Reformed Church in Hungary
Kálvin tér 9.
H-1091 Budapest, Hungary
T: (+36-1) 455-9060

L'Harmattan Kiadó
Kossuth Lajos utca 14–16.
H-1053 Budapest, Hungary

L'Harmattan France
5-7 rue de l'Ecole Polytechnique
75005 Paris

© Marcus Kracht, 2020
© Károli Gáspár University of the Reformed Church in Hungary, 2020
© L'Harmattan Publishing, 2020
© Éditions L'Harmattan, 2020

ISBN 978-2-343-20655-4
ISBN 978-963-646-546-9 (pdf)
DOI: https://doi.org/10.56037/978-963-646-546-9

Volumes may be ordered, at a discount, from

L'Harmattan Könyvesbolt
1053 Budapest, Kossuth L. u. 14–16.
Tel.: +36-1-267-5979
harmattan@harmattan.hu
webshop.harmattan.hu

www.amazon.fr

For Johanna
For the benefit of all

CONTENTS

Foreword to the English Edition 9

Preface 11

Chapter 1 • Introduction 13

Part I • Basic Concepts

Chapter 2 • Theory 23
What Is Knowledge? 23
A Definition 23
Knowledge a Priori—Knowledge a Posteriori 30
Place, Time and Subject 34
Information and Entropy 38
Traces 44

Chapter 3 • Sources of Knowledge 49
The Four Sources of Knowledge 49
Experience 51
Culture 58
Science 67
Thinking 74

Chapter 4 • Language 81
Language-Dependent Knowledge 81
Signs 83
Codes 93
Language and Worldview 96
Context and Interpretation 103

Part II • Material Culture of Knowledge

Chapter 5 • Conservation of Knowledge 109
Archives 109
Administration 113
Durability. 117
Knowledge as a Mass Product 124

Chapter 6 • Passing on Knowledge 133
Surpassing the Limitations of Space 133
Networks 139
Generations 146

Chapter 7 • The Use of Knowledge 153
Cost and Benefit 153
Big Data. 158
Technical Limits of Access 161
Dead and Living Knowledge 163
Relevance 166

Chapter 8 • The Future of Knowledge 171
Science and Religion. 171
Ignorance. 177
Blackout 181
The Education Catastrophe 185
Peak Knowledge 187
What Are We Losing, Really?. 190

Bibliography 195

FOREWORD TO THE ENGLISH EDITION

After publishing the German edition in 2017 I began to realize just how far I was from understanding the problems I was addressing. Since then I have toyed with the idea of rewriting it completely. But that proved to be a very difficult task, taking a few years at least. I feel that much of the scientific data that I would like to have is not available and would have to be collected. But the topic cannot wait. I feel that the message must go out now. I decided instead to retell the story in a translation. The idea to this was prompted by talks with Ivan Mikirtumov, who shares my concerns. So, I sat down and translated the text into English, but not without changing it at places where I now feel I know a better way to tell the story. Many thanks to Daniel Milne and Robert Denis, who edited the last version of this book.

Werther, 2018/2019
Marcus Kracht

PREFACE

This book is an attempt at gauging the challenges for our knowledge at the dawn of the great energy descent. That knowledge and energy are linked may come as a surprise. This is why I felt that this book had to be written. For if one tries to understand what the link between the two is, one realizes that no scientific invention of any sort will make it obsolete. For the roots of this connection lie in the fact that any immaterial phenomenon, in as much as it exists in this material world, must have a material manifestation without which it simply cannot be. In particular, this holds for knowledge. To make others appreciate the depth of the problem, I have found it expedient to begin with a multidisciplinary approach to the phenomenon. For if one begins to see that it isn't just a philosophical problem to determine the nature of knowledge, more and more disciplines come to mind that can contribute to an answer. This concern gave rise to a manuscript that is neither philosophical nor scientific, nor do I approach the subject exclusively from a cultural or sociological perspective. I had to transcend the disciplines to capture the problem as a whole. And yet, the disciplinary background is still visible, because they all contributed to the subject matter in one way or another.

Although the text is written in a linear form, it need not be read chronologically. Often I go into some detail which is not necessarily needed later on. Some sections can be skipped without losing track of the argument. Especially the first chapters will be somewhat theoretical. I did not shy away from certain discussions so as not to give the impression that the conclusions are based on vague conceptual development.

I thought it necessary to go into some depth with respect to theoretical questions because they will invariably appear once one tries to think the matter through. The matter at hand is too important to gloss over these issues. Even though I find it completely obvious that our modern way of handling knowledge is threatened by the energy crisis, arguing my case involves explaining the mechanics of our knowledge-based society to some degree. And it also requires clarifications as to what exactly I mean by that. Exactly what forms of knowledge and which organizational structures are most at risk is worth knowing. Our well being depends on it.

There are many people who have contributed in one way or another to this book. The most important discussions I have had with my wife Johanna, who listened with patience to my rants about energy and society. Other people who have given me valuable feedback were Franziska Bergner, Christoph Brandt, Wolfgang Braungart, Peter Finke, Peter beim Graben, Lutz Graner, Udo Klein, Stefan Kühl, Martin Kleinelanghorst, Kai Löser, Daniel Milne, Waldemar Nebe, Paul Nellen, Philipp Nürnberger, Jürgen Schnack, Jan und Bärbl Wirrer, and Werner Zittel. They all have inspired me and given me the patience to go on.

Werther, February 2017
Marcus Kracht

CHAPTER 1

INTRODUCTION

The computer and information technology has been so immensely successful that it created the impression that just by increasing the body of knowledge we can add value to our lives. Everything else, labour, energy, raw materials are just sitting there idle waiting to be called into service by the ever increasing information machine. To see how influential this way of thinking is you just have to read the Lisbon Declaration of 2000. Its aim was to turn the EU into a knowledge society within 10 years based on the conviction that the only way to secure competitiveness—and therefore growth—was to foster knowledge.[1]

This knowledge society did not come to be. One may argue whether this was because it was not properly put into practice or whether it was the big crisis in 2008 that put an end to these dreams. Or maybe the goal had just been too ambitious. I would certainly agree with the latter. It was to be expected that Europe would not be able to pull this off alone, especially if everyone else was probably trying to do just the same. But what it was exactly that thwarted this hope is hard to say. My intuition is that the reasons run deeper than what can be measured with statistical data.

Perhaps it is hard to imagine that one could motivate humans to learn and study hard simply by telling them that this would boost the sales figures or grow the GDP of their respective countries. This suggests that knowledge is not so important for the people themselves as for the companies or the society as a whole. The constant refrain of generating knowledge just for the sake of keeping up the race is itself damaging in the long run. The crisis of the Western world would thus indicate that its inhabitants are not knowledgeable enough or have acquired the wrong kind of knowledge. Although this is not actually off the mark, I do not find this too plausible either. A modern society does not need to distribute its knowledge widely, so few and select groups of experts are often enough, especially in high skilled areas.

In this book I shall propose a different answer. And that is that the knowledge society that was dreamt about did not become reality not because people did not try hard but because it simply was not feasible. Not now, and certainly

[1] For a definition of the term *knowledge society,* see Nico Stehr: *Arbeit, Eigentum und Wissen: Zur Theorie der Wissensgesellschaft,* Frankfurt am Main, Suhrkamp, 1994; Nico Stehr: *Die Zerbrechlichkeit der modernen Gesellschaft,* Weilerswist, Velbrück Wissenschaft, 2000.

not in the future. The problem is not only that it is hard to organize a society in such a way that humans will always strive to enlarge their knowledge and skills. The problem is more fundamental. It is the idea of the knowledge society as such that is flawed. I contend that the idea that knowledge of whatever sort can be enlarged to any degree is inherently wrong headed. The most important difficulty is not the quality of knowledge; it is the quantity. The project will fail because of the quantity of data that it needs to generate.

The problem is foremost of a physical nature. Knowledge acquisition and preservation cannot be separated from the physical world we live in. No matter how we turn it, it costs energy and raw materials to gain and maintain knowledge. This is not a new insight, but what is often lacking in public discussions is the rigorous application of this insight in practice. It will probably take some time until the effects of the energy crisis on the question of knowledge will become fully visible. Traditionally, in a society without books and information technology, the limit to the acquisition of individual knowledge has been the amount that one was able to absorb in a lifetime. This is a really hard limit. For a while it seemed that this limit no longer applies. Not only did books change the way we treat knowledge, the computer and its virtually limitless capacity for information processing seemed to change everything for good.

If we just briefly recall the recent developments, this enthusiasm is understandable. The computer became so fast that we do not really see it calculate, and big documents can be transferred from one corner of the world to another in a fraction of a second. And with the advent of deep learning it seems that even the days of theoretical research are being counted. It seems that in the near future we do not need to produce new theories about anything, we just need to produce reliable data in sufficient quantity and have a machine learn the rules of the game. If the projections are correct, then the computer will soon enter even into professional fields that were deemed to require too much skill not long ago.[2]

However, there is a new development on the horizon that may soon change all this again. It is the resource scarcity that is creeping into all facets of life. The fundamental premiss of this book is the following: "The era of expansion in terms of material and energy consumption is over." In other words, we are facing peak energy and peak materials.[3] I shall take this as my starting point, though I shall return to the question later when we need to be more specific. The aim of this book is to show what this means for our knowledge,

[2] See Carl Benedict Frey – Michael A. Osborne: *The Future of Employment: How Susceptible Are Jobs to Computerisation?*, Tech. Rep., Oxford University, 2013.

[3] See Richard Heinberg: *Peak Everything. Waking Up to the Century of Declines* (2nd ed.), British Columbia, New Society Publishers, 2010; Ugo Bardi – Leigh Yaxley: *A Lotka-Volterra Model of Worldwide Energy Consumption – ASPO-5*, 2006; and many more.

both everyday knowledge and scientific knowledge. I will argue that the end of energy growth also spells the end of the scientific era, by and large. To this end, a number of points need to be confirmed. The first is that there really is no way we can decouple knowledge acquisition and preservation from its energy and material base. Although mankind has struggled hard to keep the material footprint and the energy costs low, the overall costs have so far always increased. The second point is that quantity is not all that important. Given a choice between various fields of inquiry, there certainly are differences in importance. So, while the quantity of knowledge will not be able to grow, the overall impact of this fact to society is less than clear. Third, no matter what one believes about the energy problem, the foundations of the knowledge society are not as secure and stable as we like to believe. It is not and it never has been insignificant what kind of knowledge we wish to attain. Not as individuals (a fact that modern societies like to downplay except when it is related to employment) and certainly not as a society. Moreover, as we are facing a growing list of threats it is of utmost importance to ask what kind of knowledge really is helpful.

The process of civilization is often portrayed as if we have now left behind the era of industrialization and are now entering the phase of information.[4] What is overlooked is that the industries are still there and are producing an increasing quantity of goods. The apparent contradiction is simply being ideologically outsourced. We are not seeing it because it is not discussed. But also the concept of the knowledge society as such is not what we make it to be. It suggests that there are societies that are not of this kind. What are these societies like? If, as it is claimed, a knowledge society is one in which science is important in decision making and in creating progress, then a society that is not a knowledge society is either one where we base our decisions not on the knowledge that we have (instead of, say, religious beliefs) or one in which no scientific knowledge exists. I am yet to be convinced that there exists a society that does not apply science when it sees its advantages. Even in Islamic countries that are considered theocracies it is engineers and physicists that are in charge of machines and technology. So, I guess it is fair to say that the distinctive trait of knowledge societies is the possession of scientific knowledge at a large scale.

However, it should also be stressed that science runs only a small part of our lives. Most of what we do is based on experience and a kind of learning that has been with us for hundreds of thousands of years: listening and talking to others. It is hard to overstate the importance of this kind of learning as opposed to the scientific education commonplace in Western societies.

[4] Just compare the theory of the fifth Kondratiev wave of information technology, starting 1990.

In sum, I argue that basically all societies have the potential for being knowledge societies in our sense—if they so wish. At closer look however it does seem that other societies are not much different in how they use the knowledge they have; it is rather the question of what is the target of their curiosity to begin with. And it should be stressed that science isn't all there is of importance. Although religion, art and philosophy do not generate scientific knowledge they take up a large part of our life in Western societies as well. And deservedly so.[5]

This book is not primarily a philosophical book. I am not interested in defining exactly what knowledge is; also sociological discussions will be not as detailed as one might hope for. The reason is not that I think they are not worthy of attention; it is rather that discussing them often distracts from the main line of argument that I wish to pursue here. The main focus is the nexus between knowledge on the one hand and the available resources, namely energy and raw materials, on the other. The question is how mankind has dealt with this problem and how it may likely deal with it in the future. The message of this book is that there is no way to fully eliminate the energetic costs of knowledge acquisition and preservation. The illusion that knowledge is immaterial is a dangerous one and did curiously enough flourish in a society that is known for its energy waste. Maybe it is at heart the same dangerous philosophy that drives these developments. And it is the thought that there is no limit to anything whatsoever. I shall try my best to convince my readers that the contrary is the case. That there is a limit even for the seemingly limitless things such as knowledge.

To do this, we have to start with a definition of knowledge and then ask in what forms it exists. In fact, the idea of knowledge has been a controversial issue in philosophy since Plato. One might think that the main issues have been settled. That is however far from true. The idea of knowledge turns out to be a multifaceted one. For the question of how we come to acquire and pass on knowledge is not simply a philosophical question or one of the philosophy of science. For it does not only touch on the question of what constitutes knowledge but also how it is embedded in society and how humans or groups can effectively make use of it. This is the domain of the sociology of knowledge. Then there is the question of reliability and durability of archives. Or the question of the reliability of the education system as a whole. How a book or a hard disk can become a source (or storage) of knowledge is not only a physical problem. More is required than the ability to read in order to

[5] For a defense of mythology as a store of ancient wisdom see Jordan B. Peterson: *12 Rules for Life. An Antidote for Chaos*, Toronto, Random House Canada, 2018.

turn a book into a useful source. And it requires a whole culture surrounding the individuals to equip them with the mastery of turning the stored data into knowledge from which they can benefit.

The problems run deeper, however. In many areas of physics new results can often be obtained only by using extremely sophisticated machines. These do not only raise questions about their reliability. They also raise questions about finances. The more governments run into financial trouble the clearer it will become that such kind of research is costly in all respects. And it will also become painfully clear that answers to some questions—no matter how fundamental we think they are—might not be obtainable not because that is impossible as a matter of principle; but rather because obtaining them would be too costly.

In addition, the world is threatened by the increasing quest for experimentation. Our curiosity interferes in a dramatic way with the livelihood of all species on this planet. The dramatic scale that the interference between the observer and the world that he aims to understand needs to be taken seriously. It no longer is negligible. The experiment succeeds at the cost of introducing a change into the world. Ideally, of course, we would like to see it in another way. The observer is outside of the world and simply observes the change that is occurring. The world would anyway proceed in the same way even if there had been no observer. However, that is not the way it works in actual practice. For the observer is not idle. He is actually exerting force on the world in order to run his experiment. That is precisely why it is called an experiment and not a simple observation. The experiment is a way to force nature into doing something that it would otherwise not have done, and precisely the element of force is what is at stake. To understand the scale of the problem just recall how many atomic bombs have been exploded simply for experimental purposes, how many animals have been subjected to „tests" or how many experiments have been conducted with genetically modified plants. So, our quest for knowledge leaves visible traces on this earth because we do *not* simply observe, we experiment. And the curiosity risks to destroy that which it plans to study. In principle, that could be a reason not to perform the experiments at all. Because the experiment would yield knowledge about something that no longer exists. While there now exist guidelines about ethical experiments in practice they just alleviate the problems. What we can see is that they are incapable of changing the course of things. For the problem is really pervasive. The anthropologist, for example, that sets out to study an indigenous culture so far untouched by the outside world does so at the peril of destroying what he sets out to study. All he can really do is document the decay of a once existing culture. The experiment goes through at the expense of destroying its object of curiosity and rendering the knowledge so obtained useless.

One is confronted with exactly these questions if one sets out to study the relevance of knowledge. The situation has been different in antiquity. While the scientists of the Greek world were theoretically minded, their modern equivalents are not.[6] Today the attitude is that knowledge is useful and you obtain it by asking nature some questions and receiving answers. The use of knowledge is in the increased domination of some phenomena or—more abstractly—in the increased competitiveness of society.

More facets to knowledge have meanwhile been discovered. Knowledge has a place and a time. Knowledge once obtained is not everywhere available and does not persist through time. Passing it on, be it to other people, be it to other generations, requires language. Language is the *master code* of society. No wonder much of early school education consists in learning how to read and write properly. Language accompanies everything we do. No occupation, no hobby without a particular set of words or phrases that one needs to learn in order to communicate with people of like interest. In this way language is constantly changing, due to our changing interests.

Also mathematicians and logicians have been discussing the problem of knowledge for a long time. For mathematicians have been claiming that they are the only ones to discover timeless truths. Everything that empirical science can discover is only preliminary. Today's truth and conviction can be tomorrow's superstition or outdated belief. However, even those mathematical certainties are somewhat less that eternal. 150 years ago it was believed that mathematics rests on certain solid, eternal basis. And that anything that was recognised as true could actually be proved. However, it has since been realized that any theory that can at least deal with natural numbers is deficient in that not everything that is true is also provable.[7] At the beginning it was only quite exotic claims that could be shown to be true but not formally unprovable. That has since changed. But, the attentive reader will surely ask, what does it mean to say that these claims are true? How do we know if we cannot prove that? And if we *can* prove that, aren't we using a method, or calculus, different from the one before? Is there a way to choose which one is better?

The longer one meditates on these questions the clearer it becomes that the vision of mathematics as an edifice built on some undoubtable foundations, creating one level after the other is seriously flawed. Knowledge remains

[6] It needs to be said, though, that the practical consequences of their theorizing where not unwelcome, either. In retrospect it seems that the essential difference was that the quest for knowledge arose in people because they could afford to think about it at all. Recall that people in antiquity were mostly required to earn a living with hard work. Philosophy was a hobby of the upper class, of those who had free time and who could afford to devote that time to whatever they liked. Today, scientists are the heirs of this aristocracy. They therefore insist that they want to pursue what they think is worthy of their interest. They are constantly reminded, though, that they are being paid to deliver.

[7] This is the famous incompleteness theorem by Gödel.

a mystery in as much it is not clear how new knowledge is acquired that is not implicit in the knowledge we already had. The correctness of formal methods can only be shown within their own boundaries. But these boundaries are being constantly redrawn.

The rational method has therefore above all convincingly shown one thing: How little is possible to achieve by rational thought alone. We have learnt a lot about logic, language, proofs and knowledge and so on. And yet, all this struggle has led to the somewhat meek conclusion that there is nothing that we know for sure. It seems that at least that we know for sure now. Socrates would have rejoiced.

What is the lesson of all this? One of them could be that we have a somewhat too euphoric view of science and the scientific method that voices an ever growing claim to power exactly in the moment when its deficits become patently clear. Another is that science and its institutions, like the church, are not theoretical beings but are in fact quite ordinary congregations, where people fight for recognition and all kinds of benefits. And they are sometimes ready to sacrifice truth for them.

We are often all too eager to forget that science is part of the culture and therefore suffers from most if not all of the illnesses that beset it. Many people say, for example, that the energy crisis is not a problem: scientists will discover new sources of energy. Or the oil crisis will be dealt with by drilling for new oil. And if that does not work, well, then we shall use our information technology to substitute for the need to use oil. Using the latest technologies we shall transition to a new fantastic phase of innovation and technology.

Somehow, the underlying thought is that science is exempt from the earthly needs. Yet, this idea is not only useless, it is dangerous. It is easy to see that it is false. Just consider the fact that all the knowledge we have needs to be passed on to the next generation. And so we need to put people in charge to educate our children. While they used to be parents, grandparents and other relatives, now it is teachers, professors and other professionals. Even if all the knowledge transfer were for free, we would still need to feed, clothe and house these people for the services that they provide. As soon as the financial crisis hits we can see how real that need is. Around the world we can see the effects of a growing education sector. It is costing more money, money that is increasingly becoming scarce, whatever the reasons. Of course, we may appeal to their idealism, but everyone needs to eat, live, and feed their children. It is misleading to think that we need to invest all it takes into education since that is the most important thing of all. Somehow this suggests that education is sacrosanct and somehow does not incur real costs. But it does. And therefore defending education will inevitably happen at the cost of something else.

Thus the question will not be whether or not we shall preserve our knowledge but rather which part of it. Heaping up knowledge as if costs did

not matter or even exist is not the answer. Increasingly, the question will be raised whether a certain project really does need funding. Whether what it purports to answer is really that important.

I am not claiming to have a list of not-so-important research projects ready, though there are some immediate candidates. When research projects depend on expensive equipment often in the range of several billion dollars, questions of their necessity are being regularly voiced. This is the case in astronomy, with space missions, and high energy physics. Clearly, what is not at stake is the quality of the research being done. It is quite feasible to continue high profile science and the pursuit of knowledge in areas where that costs a fraction, like sociology, linguistics, philosophy, or mathematics, to name just a few.

Not all parts of science and life can exist without this expensive research machinery. When the support dwindles, they will go extinct as well. This however might not be life threatening at all. I am pretty sure we can do without these costly endeavors. Billions of people on this planets give testimony to this every day. Of course, they all need to know an awful lot to be able to make it through the day. But their knowledge is not of the kind science likes to talk about.

One thing is certain beyond doubt: in the coming years we shall witness an unstoppable transformation of knowledge. In contrast to most gurus of the information age I see less knowledge on the horizon, not more. All we can do is direct this transformation in the most beneficial way.

PART I

BASIC CONCEPTS

CHAPTER 2

THEORY

What Is Knowledge?

A Definition

Before we can ask meaningfully about the material manifestations of knowledge, we need to inquire into what it is that we are looking for. In other words, we need a definition of the term "knowledge". There are, as always, different opinions. This ought to be no surprise. People mostly have different opinions about the meaning of words, especially when it comes to abstract words. Additionally, different people like to emphasize different aspects of the same things. So, if a philosopher is asked for a definition he will probably emphasize different aspects than a linguist or a computer scientist. While a philosopher is more inclined to look at the essential nature of knowledge, a linguist will remind us of the fact that all knowledge needs to be expressed in some way in a language, while a computer scientist will most likely highlight the connection with data, its transfer and storage. None of them is right or wrong. The philosophical definition is necessary in order to understand what the nature of the object is so we can identify it in its manifestations. The linguistic view is necessary so that we understand how the thing can at all manifest itself. And the computer scientists' view is important since we also need to be practical about knowledge transfer. This is why I need to discuss the basics at somewhat greater length than some of my readers would like me to. Please bear with me, it will pay off later.

The definitional method will be as follow. I shall propose a definition and then proceed to inquire into its exact understanding and the possible ramifications. I shall start with a somewhat restricted idea: I shall inquire into what it is for a person A to know some proposition P. Here, P will be a sentence of some language and not its content, that is, what it takes for P to be true. The latter will be referred to here by the word *proposition*. So, I am not asking what it is to know that snow is white but rather what it is to consent truthfully to the sentence "Snow is white." Many questions will ensue, for example, what it is for a group of people to consent (truthfully) to that sentence. That in turn will raise interesting questions such as: What does

it take for a group to consent to something? How are meanings negotiated in a group? For as we shall see, different people have different opinions about the meanings of words. So, the sentences they give consent to, actually have different content. Thus, they give consent to different claims about reality. This is a troubling aspect; it is however not so different from problems that individuals face with respect to their own individual judgments. We know by experience that we are answering the same questions differently depending on circumstances.

We shall also raise the issue whether or not animals or artefacts such as books or computers can be said to possess knowledge. Indeed, if knowledge is independent of language—since the original definition talks about propositions and not sentences—then it ought to be possible for animals to possess knowledge even though they do not talk. However, they do possess the ability to judge circumstances. With books and computers, it is the other way around. They do seem to use language to code certain statements, but they do not seem to possess the ability to judge for themselves. Yet they can be used by individuals to answer successfully questions the individuals would not be able to answer alone. Thus, books and computers are artefacts to extend our knowledge of the world. To be able to draw the boundary between things said to possess knowledge and those that do not is going to be difficult but worth the while. For it will shed light on the way in which the culture creates artefacts and uses them for its purposes.

Let's begin with a definition, though. I basically agree with Timothy Williamson who defines knowledge as a basic attitude towards a true proposition.[8] To say that A knows that P means that P is true and that A believes that this is so. Actually, his story is more complicated as he struggles to define the nature of the propositional attitude. We shall use a somewhat different definition that will highlight the linguistic nature of knowledge. For us the test whether Anton knows that Paris is the capital of France means—apart from the fact that it is true—that he will answer "yes" to our question "Is Paris the capital of France?" Again, as the previous definition requires, knowledge requires giving a correct answer. Hence, Anton cannot know that Paris is the capital of Belgium even if he correctly answers the question "Is Paris the capital of Belgium?" And this is because this statement is false. Likewise, if Anton answers "no" to the question "Is Paris the capital of France?" then he is also said to *not* know Paris is the capital of France. So, it is required that the question requires a positive answer *and* that the person gives exactly that answer.[9]

[8] See Timothy Williamson: *Knowledge and Its Limits*, Oxford, Oxford University Press, 2000.

[9] This is a case of presupposition. To know that P means that P is true; to not know that P also means that P is true.

In contrast to the first definition the second one is operational. It tells us how we can recognise knowledge. However, we need to be somewhat careful here. First of all, this test assumes that Anton is not lying to us. To get around this difficulty we imagine that it is actually Anton who asks himself that question. He will surely know whether he is telling himself the truth. This however means that our criterion is no longer a truly empirical one. It does no longer tell us whether or not Anton knows something; it tells us more about what knowledge actually is.

Let us next scrutinize the object of the attitude. What is this P exactly? The standard answer is that P is a proposition. And a proposition in turn is something that is either true or false by nature of how the world is. A proposition is independent of the way in which it is expressed. So, we say that the proposition expressed by the German sentence "Paris ist die Hauptstadt von Frankreich." is the same as that expressed by the English sentence "Paris is the capital of France." We therefore expect that the answer to the first question by a German speaker to be the same as that of an English speaker to the second—provided they both know that Paris is the capital of France. Yet the ability to answer in that way depends on the ability to represent relevant properties of the world in our heads as well as mastering the respective languages. For what does it mean to ask ourselves a question? The independence of propositions from language does not mean that they are not linguistic objects. We can for example consider them to be objects of some mental language that we humans all share and that the ability to speak some other language means simply to be able to correctly translate from the language into that mental language.[10]

The next question is a practical one: how do we recognize if something is knowledge? To return to our definition, if A claims that P is true, how do we ascertain that P is true? Intuitively, since the definition presupposes that P is true this seems to be a given. But it is not. Belief that P and knowledge that P are separated by the difference whether or not P can be false. So, while we can *believe* P without it being true, we cannot *know* P if it does not hold. This, however, is merely a linguistic distinction. It only tells us when it is appropriate for us to say that A knows or simply believes that P. So, in order to ascertain whether A actually *knows* that P, we need to establish that P is true. How do we do that? We could ask someone else, for example. But then the same question arises. How do we know that that other person knows that P? We could go on to ask a third person but that would simply defer the problem. The search is circular. Something else must come to the rescue.

[10] In a sense the mental representations could be said to be expressed in that language. This is effectively how many computer implementations work. They code states by means of a language, typically first-order predicate logic, and then translate human language sentences into this language.

Thus, we turn to a different definition of knowledge, going back to Plato. It says that knowledge is true justified belief. Of course, this definition is at first blush as circular as the previous one.[11] Williamson argues that knowledge is a mental state that is directed towards a true proposition—and nothing more. The notions of belief and justification are according to Williamson misplaced.

We shall nevertheless scrutinize this definition and see how it can be useful for us. For a proposition P, it says that "A knows that P" is true if and only if (i) P *is true,* (ii) A *believes* that P is the case, and (iii) P has justified grounds to believe that P is the case. Take the case of the proposition that Paris is the capital of Belgium. A cannot know that since it is false (this constitutes a violation of clause (i)). Now take the proposition that Paris is the capital of France. It is true therefore (i) is satisfied. By (ii), A must believe that it is the case. But according to the definition this is not enough. This is crucial. Consider the case where A has read in a novel that Paris is the capital of France.[12] Then we would say that this is not an adequate justification for knowing that. Not because it is not true but because novels are unreliable sources in this respect. Indeed, it is part of general knowledge that what is claimed to have happened in a novel does not have be true. In fact, we rather expect it to be pure fiction. By contrast, an encyclopaedia is supposed to be a source of facts. So, if instead of claiming to have read about Paris in a novel, A claims to have read about it in an encyclopaedia, we would be inclined to say that A knows that Paris is the capital of France.

Of course, all this is less than guaranteed. There are popular science books which claim that our civilization has been created several thousand years ago by aliens. Is the source reliable? How about our friends, parents, teachers, Nobel laureates? How about Wikipedia? We all know from experience that the reliability of a source is not a given. Even though we might expect a Nobel prize winner to be an expert in his field, it is far from clear that his claims on other subjects are more trustworthy than that of other scientists. Similarly, we may admire our teachers, but they might sometimes say things that are beyond their expertise. And very often it is not even clear that we know where the boundaries of our expertise really are. How well do we know what we know after all?

[11] It has been criticized among others by Edmund Gettier: Is Justified True Belief Knowledge?, *Analysis* 23 (1963), 121–123; and Williamson: *Knowledge*. I agree that it is problematic. But it will nevertheless be useful as a heuristic.

[12] We shall return to this problem in subsequent chapters. Basically, we can have true beliefs that are unjustified. And that is a big problem in science. Consider the attitude of our modern medicine towards traditional medicine, e.g. Chinese medicine. It is considered flawed not so much because we know that it does not work but because it has not been obtained using rigorous scientific methods. Based on the definition above it can indeed not be counted as knowledge.

It seems therefore that even our second definition is deficient. The reason is not so much that it operates with terms like *belief* or *justification* that are no clearer than what they seek to define. Rather it is that they use terms that once again talk about mental states, and those are almost impossible to ascertain. This holds not only for the concept of "belief". It is also true for "justification". Everything that we know must by definition have a justification. That justification in as much as it too is something we claim to know must be justified as well. And so this leads us down the same path again. Anything we claim to hold must be justified. Thus, we need to go on forever or we stop at some point and declare that an axiom. Furthermore, it is not always clear what makes one fact a justification for another.[13]

And how can we verify whether someone believes something? Is it useful to use that concept in the definition? Williamson does not seem to need it so maybe it is superfluous?

I shall propose a revision of this definition, one that can also be extended to cases where the subject is not a human. When talking about machines we are not likely to say that a computer believes the answer, but we are inclined nevertheless to say that it knows it. In a culture that treats computers and other machines as the extension of the body and mind it may be useful to understand in what ways that is (or is not) a justified attitude. Thus, we shall formulate a version of this definition that does not require the term "belief". That will be useful also in human interaction. There often are situations where we would like to say that we know something is the case, but we do not believe it.

The new definition is not meant to define knowledge. It is meant to be a practical tool for determining what may or may not be regarded as knowledge. We start by replacing propositions with sentences of a language. There are different kinds of them, questions, commands, and declaratives. Declaratives are sentences, that express propositions. Given a declarative sentence S the question is whether the proposition P that it expresses is true. This is called a *problem*.[14] The solution of the problem is constituted by the correct answer. So, if P obtains we should say that S is true and if P does not obtain we should say that S is false. Equivalently, we can turn S into a question ("Is Paris the capital of France?"). And then if P obtains the answer to that question should be "yes" and if it does not obtain it should be "no". The idea is now this. To know whether or not P is to reliably answer correctly to the associated question. Or—more cautiously—to reliably judge S correctly. There is some terminological problem in talking about declarative sentences

[13] See John Leslie Mackie: *The Cement of the Universe: A Study of Causation*, Cambridge (Mass.), MIT Press, 1980, for a similar discussion on causation.

[14] This terminology derives from complexity theory. Problems are functions from some set F to {0, 1}.

rather than propositions. This will turn out to be important. I will maintain that the primary target of our attitudes are the declaratives sentences, but that ultimately, they are about the propositions thereby expressed. To make this distinction visible, I shall talk about judging sentences to be true, while claiming or believing propositions. This distinction is quite crucial. A judgment is a momentary act. It happens in time and can be repeated. Thus, one way to establish whether or not A knows that P will be to repeatedly ask him to judge the declarative sentence S. Thus, while it may be that A answers a question correctly once this maybe just by chance. To take an extreme example, if we were to "ask" a coin whether Paris is the capital of France, stipulating that heads means "yes", and the coin shows heads after the toss—would we say that it knows the answer? Hardly. But if we asked a computer whether 2 plus 2 is 4 and the answer is "true" we would be so inclined. The difference is that the coin is not consistent in its answer. We cannot repeat the experiment. While the computer will always give the same answer the coin will not. This is why the definition contains the word "reliable". This requires that the act of judgement is repeatable and yields the same result all the time.

This definition turns knowledge into an action disposition. We have thus dispensed with mental states such that doubting, believing, and so on. Instead the criterion is that the person acts in a certain way. This however requires encoding the proposition in a language. Instead of being able to ask directly about P we need to cast P into some question Q and ask A that question. Of course, A can be lying. But that is, as Williamson rightly observes, beside the point. We are actually asking what his judgement is. Whether or not he truthfully responds, is another matter—though it is one that needs careful scrutiny in practice.

Since the judgement concerns a sentence this brings in another variable, namely the language in which that sentence is expressed. We need to ask whether the sentence "Is Paris the capital of France?" really *is* expressing the question whether Paris is the capital of France.[15] That sounds trivial but becomes very complex once we look at foreign languages or abstract concepts. And we need to ask to what degree a person can be said to understand the meaning of a sentence.

I touch briefly on an important philosophical problem. It is clear that we can never be one hundred percent sure whether something is true or not. We can be convinced to the utmost degree about the truth of a claim and yet it may turn out to be false. This happens even to mathematical claims. There is a gap between our beliefs and the world out there that cannot be bridged. Anton can be convinced that Paris is the capital of Belgium. And David may disagree

[15] Recall the famous truth condition by Alfred Tarski: The Concept of Truth in Formalized Languages, in J. Corcoran (ed.): *Logic, Semantics, Metamathematics*, Indianapolis, Hackett Publishing, 1983, 152–178, that "Snow is white." is true if snow is white.

strongly. Each of their beliefs may be as strong as the other's even though one of them must be wrong. There will however be no end to their debate as long as none of them produces some evidence (other than just continuing their respective claims) that the other one accepts. And it may be David who ends up being convinced by Anton that Paris is the capital of Belgium.

While this discussion seems absurd—after all we all know that Paris is the capital of France—it is far from obvious when we turn to other questions. For example, is the climate warming an effect of burning fossil fuels or not? Have we reached peak oil or not? Is poverty rising in Western Europe or not? None of these questions are easy to answer, for various reasons. It is not only a question of facts; it is also a question of words and definitions. Not to speak of the fact that the answers to each of these questions touch on the lives of millions of people and so we cannot expect everyone to be answering to their best knowledge.

Finally, I shall touch on a point that is becoming more and more important as we move into the age of internet. The definition of knowledge talks about a subject answering a question. In short, knowledge is not some object, it requires some subject to judge it. Data is turned into knowledge by a mind. A book is an array of letters, just like a hard drive is an array of bits. What makes them into knowledge is the mind that synthesizes them into a judgement. This is a difficult point, but it explains why knowledge, information and data are not simply synonymous. Data is the raw stuff, the letters, the stored bits, the graphs on a sheet of paper. Information is what they can be computed as saying. Knowledge is the ability to judge whether what they are saying is true. Information becomes knowledge only by passing through some mind. It is part of what Searle is after when he says that language is *social reality*.[16] Knowledge likewise is a social reality since it is being symbolically coded; so, what constitutes the social reality of knowledge is both the fact that it is coded in language but also the fact that the way it is turned into an aspect of reality is again socially determined. We shall discuss later how that is to be understood.

The line between the subjective and the objective is however not really easy to draw. Information, too, has a subjective element. Without a code that lets us understand what a particular string of letters or bits really is saying, data simply does not become information. And that code also is a social reality in the sense of Searle. The ink marks on paper do not become letters, the letters do not become parts of words if there was no subject that is in possession of the code that makes that happen. We could of course write down the code itself—as is done in many cases—but again it remains to be decided which

[16] John Searle: *The Construction of Social Reality*, Free Press, 1997.

data is to be seen as information under what code. No matter how we turn it, the sense that makes all this hang together is not a given. It is us humans that decide for ourselves how the traces are to be read.

I have resisted to use a restricted interpretation of the term knowledge, though, for the reason that it does not reflect the current usage. I will however discuss later the dangers of equating data or information with knowledge itself.

Knowledge a Priori—Knowledge a Posteriori

After having settled on a definition of knowledge, let us start exploring its consequences. The first problem to look at is the infinite regress built into explanations. Since claims to the fact that we know such and such is the case because there is a reliable source that says so, for example, lead to asking the same question about the source, and additionally about how we come to know that it is reliable. These are then things to be investigated in turn. This can end only if there are claims that need no backing because they can be known in and of themselves. It would be interesting if they could be used to explain all the other things that are not self-evident. To make progress on this issue Kant introduced a distinction between *analytic* and *synthetic* judgements and between a judgement *a priori* and a judgement *a posteriori*.[17] These distinctions also make sense in connection with knowledge since as we have seen knowledge is closely connected to judgement. So, we can distinguish between knowledge a priori and knowledge a posteriori. Judgements a priori are judgements that do not require any knowledge about the external world and can therefore be evaluated without any experience. An example of such a judgements are mathematical claims such as "2+2=4". They surely are true a priori. Anything that needs empirical confirmation is by definition a posteriori.

The distinction between analytic and synthetic according to Kant is not the same as in modern days. It will not be of much relevance in our investigation and therefore I will touch on it just briefly. A statement is analytic according to Kant if it's true in virtue of the concepts that are being used. A clear case is the claim that all bachelors are unmarried since being unmarried is a condition for being a bachelor. Another case that Kant explicitly mentions is "5+7=12". In his view the concepts of the numbers are not made in such a way that the concept of the number 12 literally contains the concepts of 5 and 7, so it is not stated explicitly in the definition that it is the sum of 5 and 7. Hence, according to Kant, although we can say that "5+7=12" is certain a priori since we do not need any empirical confirmation,

[17] Immanuel Kant: *Kritik der reinen Vernunft*, Frankfurt am Main, Suhrkamp, 1974.

it is not analytic. It is therefore synthetic. Synthetic means that the mind needs to perform extra steps (of synthesis) to reach its conclusion. And that is certainly the case, since children *do* need to learn that it is true even if they know what the number 12 is.

Let's return to the distinction between a priori and a posteriori. As we said a judgement is a priori if its truth can be established by the mind alone. Now, judgements are couched in terms of declarative sentences. The question then is what is it that makes S true a priori. A first approximation would be to say that S is true in virtue of what it means; that it to say, S is true in virtue of expressing a certain proposition P that is necessarily true. Since sentences are couched in a language, and a language assigns meanings to sentences in an arbitrary way, we must therefore be a bit more careful and say that S is *true in a language L*. It is the language that turns a sentence into a carrier of meaning. Just as an aside here: von Weizsäcker discusses the idea that the language of physics is exact because it is expressed in terms of equations such as $E = mc^2$. He notes that this is not self- evident. For it is the education of physicists that makes them read the equations in a certain way. As such, the equation is a string of symbols none of which has any meaning in English. One must be trained to understand what they stand for and how to interpret their particular arrangement.[18] In a similar vein it may be noted that logic classes begin by instructing students to take expressions like "every" as allowing for no exceptions. That is necessary since in colloquial usage we mostly do allow for exceptions.

Similarly, to understand what the equation "2+2=4" means one must know the meaning of the individual symbols, "2", "+", "=" and "4". In addition, one must know what the concatenation of these symbols is signifying. Once that is clear, though, one can find out that this is a true statement. It is similar to the game of chess. To become a chess player all you have to know is the rules of the game. There is no luck involved in winning the game. Of course, to become a good chess player you need to have some talent; likewise, to become a good mathematician you need talent. Even being able to perform complex calculations is not without challenge. However, one thing is certain: in principle everyone can pick up these arts, they need no extra knowledge or equipment. You can learn them in your room.

Any other statements need verification, though. There is no way to become a historian, for example, without learning about actual historical events. It is perfectly clear what the sentences "Caesar crossed the Rubicon.", "Cicero was a great orator." and "A slave of Cicero invented a shorthand." mean.

[18] Carl-Friedrich von Weizsäcker: *Die Einheit der Natur*, München, Carl Hanser Verlag, 1971.

Yet it does not follow that we can know whether or not they are true. This requires knowledge of something else. The same applies to many other fields of inquiry.

The case of physics and chemistry is less clear, though. Many physicists believe that the natural world is what it is by necessity. And yet it is not clear whether even the most basic facts of physics are clearly non-empirical or a posteriori. Take the fact that we are living in a four-dimensional world. Does it have to be like that? Many physicists think that there are more possibilities and they are actively exploring them.[19] Thus the claim "Our universe is four dimensional." may actually be false. And I have not started to raise doubts about whether we know the exact meaning of the words "universe" or "dimension".

Clearly it is not simply a matter of whether our words possess exact meanings. For all intents and purposes the meaning of physical statements is precise. And yet they are not a priori. The reason is that we are perfectly capable of at least in principle investigating what it takes for them to be false. In mathematics one often does the same. When it seems impossible to prove a statement one might actually try to come up with counterexamples. However, the difference is that if a statement is true, counterexamples cannot exist so we cannot hope to find one. In physics, though, there seem to be alternative theories of space-time. They are just not theories about *our* space-time.

However, the situation is more difficult when it comes to verification. In principle, no amount of empirical evidence is enough to make sure we know some physical law with absolute certainty. In mathematics, however, we seem to be certain that there is knowledge of universal claims such as "For no n greater than 2 there are rational solutions of $x^n + y^n = z^n$".[20] Obviously, even though the meaning of the sentence is clear, it has taken centuries to establish its truth. And the method for doing that is to provide a demonstration, a proof. The proof is our evidence that no counterexample can ever be produced.

Evidently, then, it is not enough to know exactly what a particular sentence means to know whether or not it is true. There are claims in mathematics that were believed to be true and then turned out to be false, just like there were claims that were believed to be true and resisted proof for decades, even centuries (or still resist proof). And then there even are those that have been proved and still there are people that try to disprove them. A particular example is the problem of trisecting an angle with ruler and compass, or of constructing a heptagon. And then there are those claims that were believed

[19] String theory is based on the idea that we are living in a ten-dimensional world, see Peter Voit: *Not Even Wrong. The Failure of String Theory and the Search for Unity in Physical Law*, New York, Basic Books, 2006.

[20] This is known as Fermat's conjecture which has been proved by Andrew Wiles.

to be true but the methods for obtaining a proof were shown to be deficient. A particularly instructive case is Euler's polyhedron theorem.[21] It says that in a polyhedron, $F - E + V = 2$, where F is the number of faces, E the number of edges and V the number of vertices.[22] Although it seemed at first perfectly clear what it means, it had become apparent that no proof could be given partly because the concepts have turned out to be less clear than originally thought. This has led to a wonderful branch of mathematics called algebraic topology (and homological algebra), but it does raise an important point. Obviously, Euler was right, and he did seem to possess all the evidence. And yet in retrospect it does seem that he *intuited* the truth more than that he actually *knew* it. Similarly, the four color theorem, which says that any map can be colored with four colors so that adjacent countries would receive different colors, was posed in 1852 and was proved more than a century later—with the help of a computer! That sparked a heated debate about whether or not we can accept a proof that was delivered by a machine.[23]

The most important complaint was not that the proof had not been checked by a human but by a machine. After all people showed that the program was bound to compute correctly, so that there is no need at least in principle to verify the proof as well. And it could be argued that a mathematician is more reliable for proof checking than a machine. After all, if one wanted one could check the proof delivered by the machine.[24] The real reason was that computer proofs give us the answer but do not allow to advance mathematics itself. A proof is not simply a means to establish an answer to a question. It is a means to understand why the claim is true. The more it makes us understand, the more transparent the proof is, the more fruitful it becomes. And it will allow us to go beyond the original theorem proved. For example, when algebraic topology and homological algebra were developed, it was possible not only to show that Euler's formula was correct. It was also possible to derive many formulae of that kind: for a torus, a pretzel, even higher dimensional objects.

Whatever was ultimately proved to be the case, the intuitive idea that the four-color theorem just had to be true, was clearly wrong. Although one had the right feeling, there is nothing that lets us see why our intuitions must be correct. In this connection Frege insisted that mathematicians did not have a concept of "coloring" or "polyhedron" because it turned out they could not say what exactly these objects were. On the other hand mathematicians do need

[21] Imre Lakatos: *Proofs and Refutations. The Logic of Mathematical Discovery*, Cambridge, Cambridge University Press, 1976.

[22] For example, a cube has 6 faces, 12 edges and 8 vertices, and $6 - 12 + 8 = 2$.

[23] See Haim Gaifman: On Ontology and Realism in Mathematics, *The Review of Symbolic Logic* 3 (2012), 480–512, for a discussion.

[24] The reason why the computer was used in the first place was that the proof required hundreds of configurations and so producing it was a rather tedious affair.

intuition even if that intuition turns out to be fallible (or—who knows—maybe its fallibility is a precondition for being able to make progress at all). In ordinary language we are inclined to say that Euler knew that he was correct.

Knowledge a priori is a small fraction of knowledge that we possess. Only mathematics and logic deliver knowledge of that kind. However, there is a third kind. The knowledge about a language is actually of neither kind. On the one hand it is an empirical fact that "cat" denotes a kind of animal. On the other hand we can decide at any moment that "cat" can denote something else. This is why one says that language is an *arbitrary* relation between expressions and meaning. So, it is not appropriate to say that linguistic knowledge is a priori or a posteriori. It is merely knowledge that is required to be able to decode the meanings of expressions.

Place, Time and Subject

Socrates claimed that learning is remembering; that all we come to know consists in rediscovering the knowledge that was once had.[25] He likened the philosopher to a midwife, simply helping to give birth to what was already there: the knowledge. Of course, Socrates was aware that people made divergent claims and so only one of them could be true. So, only one of them could be said to possess knowledge while the others only had opinions. He would have to say that it was important to remember in the correct way.

Now even if we think this was problematic, let us put that worry aside. There is some aspect of his theory that is worth discussing in detail. It implies that we can have knowledge of future events. For if all we can do is recollect the knowledge we had before, then that which we will come to know in the future is again something that had been known before. Surely, Socrates would have had trouble saying whether or not Paris is the capital of France. It is not clear how to understand the notion of rediscovering the knowledge in this case.[26] In everyday life the only things we really claim to know are those that we have witnessed with our own eyes.[27] Everything else we know is because other people have told us so or because we correctly deduced them from something else. It is less controversial that we can know about things that happened in the past, even though we cannot have witnessed them. But it is nevertheless also possible to know about events that will happen in the future. One such type of event is the movement of comets and

[25] See the dialogue *Menon*.

[26] Interestingly, Socrates demonstrates his theory by having a slave discover a mathematical truth, namely that the surface of a square whose one side is the diagonal of another square is double the surface of the latter square. So, this is a truth a priori according to Kant. In that sense this truth is timeless.

[27] The Greek word for knowing, *oida*, means literally *I have seen*. Similarly, German *wissen* is related to Latin *videre*, to see.

an upcoming solstice. Apparently, even the Greeks had machines to calculate solstices.[28] Seen from this angle there is nothing mysterious about knowledge of future events. In principle, it could have been explained to Socrates what this sentence is claiming. However, he would have had no means to find out whether or not it will turn out to be true.

The fact that it claims however is timeless. Of course, there was a moment when Caesar crossed the Rubicon however seen from an external point of view, God's perspective if you will, this event exists in this universe, so the fact that Caesar crossed the Rubicon on the 10th of January 49 BC is timelessly true. And so Socrates would in principle have been able to know about it. The fact that it occurred in the future is of no particular concern. We also do not know from direct experience that it happened. There are no witnesses alive, no traces except some manuscripts that are not even originals! So if there is any merit in saying that we *can know about events in the past* we will have to allow for knowledge of future events as well. The same is true for the place. We can know of events that happen at other locations. The mere fact that we are not a witness is not much of a problem. So, while the event has a time and a place, this is not so for propositions. Thus, we are led to conclude that propositions are timeless and placeless. To make this clear we must replace the sentence "Paris is the capital of France" by a variant that includes a time point (and some location, if necessary): "In 2017, Paris is the capital of France".[29] Viewed this way, propositions are extra-worldly entities without time and place. What makes them true are facts. And facts are located in space and time.

Even though propositions are timeless, this is not the same with knowledge. Obviously, we do not know all of the things that are true. The reason is that knowing something requires some mental representation. The facts alone are not enough. For if they were, we would instantly know about everything that happens. Having to have a representation means having to build and maintain it. Therefore, knowledge can not only grow it can also disappear. We can forget, individually or collective, that something is the case. We shall return to this issue in a later chapter. And in as much as knowledge requires representations there is no final answer to a question. For the answer obviously depends on what is represented in our minds and how. And this means two things: knowledge possesses a subject that can answer to the question, and a time point. Therefore it would be more appropriate to start with a definition of *A knows P at t* where *t* is a time point. However, the way we defined that throws up a problem. We said

[28] The mechanism of Antikythera is said to be such a device. When it was found in Aegean Sea it was so badly damaged that it was unclear what its precise purpose had been. Since then it was reconstructed and seems to have been a portable device the size of a book that was able to compute the movement of the planets.

[29] Jon Barwise – John Perry: *Situations and Attitudes* (4th ed.), Cambridge (Mass.), MIT Press, 1986, were proponents of the idea that facts are located.

that it requires that A answers reliably (and correctly) to the question whether P is the case. Reliability can be ascertained by repeatedly asking the same question. That however is not possible if knowledge is time bound. This is why the earlier definition talked about justification. A knows that P if P is correct and A has justified grounds for believing P. So, if we know what grounds A has for believing P that would constitute evidence for A's knowing that P, evidence that does not require asking A repeatedly. Notice however that there is a difference in applying the definition. Reasons for *us* to know that A knows that P need not be A's own reasons. For us the reasons may be in terms of A's behavior, his upbringing, other knowledge and so on. These are surrogates for asking A multiple times the same question. Just a simple example might illustrate this. If A is considered an expert in a particular field, then we are inclined to say that he knows what he is claiming. And knowing that P we would even predict that A knows it as well even if we had not previously asked him.

Knowledge has its time. Knowledge also has a space. To see this, consider a group of people and ask yourself: Who among these people knows the answer to the question? And what does it take for a group G to know that P? Consider this case. High ranked officials of the German ministry of economic affairs were negotiating with other parties a treaty (the so-called MAI, multilateral agreement on investments) without the knowledge of government. But in a sense, it was part of the government. Therefore, did the government know of these negotiations or not? Does not the fact that the officials of the ministry know imply that everyone else knows? Obviously not. And that is because to know something means that it is represented. And it means that the representation can be accessed. If A has knowledge of something, B does not need to know it as well. Because even though that fact is represented in A's head, B has no direct access to that representation. So, the government does not know about the negotiations just because some of its members do.

It is a difficult and very often legally relevant question of what constitutes knowledge by a group of something. It is likewise difficult to establish a consensus concerning a matter. High courts have several members despite the fact that the documents on the basis of which the verdict will be issued are known to everyone. That is often because the application of the law is not straightforward and there can be different opinions on the matter. Yet, while the individual judges have their own opinion the verdict is issued as a verdict of the court. In that sense, the individual members may have dissenting views. But since the group verdict is based on their viewpoint whether it is dissenting or not, they can't technically claim to see the matter differently as a court. That is because the *subject* of the verdict is different. How and when a consensus has to be formed is a matter of definition and is in many cases established beforehand. Often a simple majority vote is enough. But often a group will establish a committee and then follow its recommendations even if it does not agree.

Of course, it seems strange to say that there is a majority vote on whether such and such occurred. But this is what happens in reality. Juries at one point will have to say whether the accused actually committed the murder. So, they vote on whether or not that is the case. Likewise, it is often said that there is overwhelming consensus among scientists that climate change is manmade. Of course, everybody knows that this is not meant to make it so; rather, it is a way to strengthen our conviction that this is what actually is the case.

With this in mind let us return to the place of knowledge. When we type in a query for Google and it responds with a list of sites, we like to say that Google "knows" the answer. However, Google is not a person nor is it a particular machine. It is a distributed network of thousands of servers each of which permanently communicate by sending data back and forth. The cloud is—despite its somewhat ephemeral terminology—a huge collection of servers, stacked in tens of thousands in large hangars. They are orchestrated by certain other machines, often enough however they perform their actions autonomously. In that they are no different from a typical bureaucrat who is tasked with reviewing applications for social welfare, handing out food stamps, or deciding on whether or not a building may be built on a particular piece of land. From the outside it looks as if there is a single unit, Google, the cloud service or some communal government, but in reality, it is a network of machines, people or even institutions, that collaborate according to some predetermined schedule.

Thus, it seems clear from this discussion that since knowledge needs representation, and that representations exists in space and time; therefore, knowledge also has a space. It is the space where the representation resides that effectively constitutes the knowledge. I am being somewhat vague here since I am have earlier claimed that knowledge is not reducible to data. The answer to the question "How many dimensions does the universe have?" requires not only some representation of whatever sort but also a subject that puts together (synthesizes) all the facts recorded in various places into an answer to the question. We do not exactly know where to place that mind. Luckily, for the purposes of this book we don't have to. It is enough to say that it is somewhere inside the human body.[30]

In what sense can we say that a machine knows something? I already said that we need to distinguish two things: the representation that collects the facts necessary for the decision and the final arbiter that synthesizes the judgement on the basis of the facts recorded. It is customary to ignore the arbiter in the determination (or to subsume it under the representation). For example, colloquially we will say that the book knows the answer to the question which

[30] Depending on what one believes or what the content of the question is, our belief is that the answer is given by the brain or the heart. However, we do not need to be that precise, even though that last answer is still somewhat vague.

city is the capital of France. Of course, it is not literally the case that the book can answer any of our questions. However, it might be that the person we are asking does not know it by himself; but if he were to consult that book, he could give us the answer. Systemically, it is the unit of human agent as arbiter and the book as representation that possesses the knowledge. But we are more inclined to attribute the knowledge to the book. So, technically we are saying the place of the knowledge is the place where the corresponding data is being stored. In that sense it does not matter whether we talk about ink on paper, magnetic particles on a tape or connections between nerve cells.

Once this is accepted, we can go on to attribute a place to particular knowledge. We can then say that since the answer to the question whether Paris is the capital of France is written on Page 95, the place of the knowledge that this is so is exactly that page. Curiously, this leads to unexpected indeterminacy. For let P be the proposition that Paris is the capital of France; and let Q be the proposition that the declaration of independence of the US was signed in 1776. And assume that the latter is found on Page 163 in that same book. Then what is the location of the conjunction of the two, P and Q? It is definitely in that book; but it is neither on Page 95 nor on Page 163 alone. Instead it is on Page 95 and Page 163 *together*. Similarly, if a team of experts working on a particular area of science is located partly in Toronto, London and Marseille then their joint competence is shared between these places. Their knowledge resides in a location that includes Toronto, London and Marseille (or whatever parts thereof they happen to reside in). Knowledge, it appears, is not confined to an arbitrarily small location. It is distributed. This is reminiscent of the fact that the brain does not actually reserve precise locations for contents (unlike a computer). The fact that we know certain things is not due to a particular cell being connected with another cell in a particular way. It is a property of the network of connections between a whole region of nervous cells.

Information and Entropy

Knowledge is tightly connected with information but is nevertheless not the same thing. One major difference is that information is independent of meaning, and it is not necessarily connected with truth. The foundational paper by Claude Shannon does not mention the word *meaning*. It simply asks how accurately we can predict the next symbol in a stream. Nevertheless, we need to look at this concept since information turns out to be the bridge between the concept of knowledge and the material aspects of it.

Unfortunately, information turns out to be hard to measure.[31] Some definitions stress the novelty character of information and hence say that it is the propositional content that remains after we subtract what we already know. Hence, if I know that Paris is the capital of France then the sentence "Paris is the capital of France." contains no information. If I only know that either Paris or Bordeaux is the capital of France, then that sentence contains one bit of information. If I know that Paris, Bordeaux, Marseille or Lyon is the capital of France, the sentence is said to provide two bits of information. For the four choices can be recorded using two bits of information: 00 for Paris, 01 for Bordeaux, 10 for Marseille and 11 for Lyon. Such a definition of information is useful if we want to measure how much one needs to memorize in order to capture exactly what has been said. This is relevant in many areas, for example, computer storage.

Other definitions do not relativize information. They do not ask what is known beforehand. This is the concept of information that I shall use here. It measures how many abstract possibilities a proposition eliminates. In contrast to knowledge it is a *quantitative measure*. For example, if we ask which of the countries Paris is the capital of, then there are 193 possibilities. (This is the number of officially recognized countries in 2017.) So, the proposition eliminates 192 of them. Information is a logarithmic measure. The logarithm to the base 2 of 193, written $log_2 193$, is 7,592. We usually round up to the next integer. In this way, the information content is 8. The reason is that we are asking how many information bits are needed to store the exact answer. Obviously, 8 bits are enough since they allow to store $2^8 = 256$ choices.

This approach obviously has its weaknesses. It depends on the question being asked. It might have been asked for example which French city is the capital of France. Since there are several thousand cities the number of bits needed to distinguish them increases and so now, we require roughly 13 bits of information ($2^{13} = 8{,}196$). By choosing an appropriate question we can obviously claim this sentence to have any information content whatsoever. The original definition of entropy avoids this dependency entirely. Let's see how that goes.

[31] See Holger Lyre: *Informationstheorie. Eine philosophisch-naturwissenschaftliche Einführung*, München, Wilhelm Fink, 2002, for the multiple uses of the term. It will emerge in the course of this discussion that all of these different uses of the term are of relevance. It is however not useful to list them here. I also mention Lienhard Pagel: *Information ist Energie. Definition eines physikalisch begründeten Informationsbegriffs*, Wiesbaden, Springer Vieweg, 2013, about the difference between information and entropy. Suffice it to say here that information is an abstract concept whose physical realization is entropy. Analogously, the proposition is abstract but knowledge about it is not.

We are assuming that we have a source Q that is sending out letters from an alphabet A at random. The probability that the letter a is used is $p(a)$. It is clear that the sum of the $p(a)$ over all a in A equals 1. Now define the *Shannon-entropy* of Q to be

$$(2.1) \qquad H(Q) := -\sum_{a \in A} p(a) log_2 p(a)$$

The reason for the minus sign is that since $p(a) \leq 1$ we have $log_2 p(a) \leq 0$, so the sum is less than or equal to zero. However, we want this to be a positive number. It represents the optimal average length of encodings of texts given that the letter probabilities are as indicated. To see why this is so consider a text T of length n, in fact, a sequence

$$(2.2) \qquad t_1 t_2 t_3 \cdots t_n$$

of length n of symbols that have been issued by Q. It is assumed that the probability of any symbol a occurring is independent of any previous and subsequent symbols. Thus, we can expect that on average the symbol a occurs n times $p(a)$ many times. Second, there are methods to code the letters of A such that the symbol length $c(a)$ of the code of a gets arbitrarily close to $log_2 p(a)$.[32] In practice, however, we only have $c(a) \geq -log_2 p(a)$. Given this, the length of the binary string coding T is

$$(2.3) \qquad \sum_{i=1}^{n} c(t_i) = \sum_{a \in A} np(a)c(a) \geq -\sum_{a \in A} np(a) log_2 p(a) = nH(Q)$$

This shows that $H(Q)$ is the lower bound of the average symbol length in a binary code. This is, of course, only a theoretical lower bound. If the code is allowed to code only single letters, then the length of a code is a natural number while $-log_2 p(a)$ generally is not. So, equality will not hold. But that is not our biggest concern. The biggest concern is the assumption that Q has no memory, that is, that the probability of a symbol occurring at a given moment does not depend on the previous symbols. This is evidently not the case in natural languages. The probability of the letter "u" occurring right after "q" is far higher than its unconditional probability. The first is nearly 1, while the second is 0.00095, close to 1 : 1000. The probability that a capital letter occurs after a period (and the blank following it) is far higher than its probability

[32] To the initiated, this is the Huffman-code. If one codes only single letters of A then it is not yet optimal. One can improve the length of the code by coding 2-letter combinations. And even that can be improved by taking 3-letter codes. And so on. In the limit the code length of a reaches $log_2 p(a)$.

in general.[33] And so on. And then we have not even touched on grammar or morphology, which also regulate how letter sequences have to look like not to speak of the content of a string. Thus, the Shannon-entropy is a quantitative measure that is useful in many circumstances, scientific and technical, and its decisive advantage is that it allows to quantify the informational content.

However, its disadvantage is that it treats symbol occurrences numerically with no regard to content. So, we would ideally like to have a qualitative measure of content. But how could this look like? Informally, we say that *P* has more content than *P'* if *P'* follows from *P*. In other words, every time that *P* is true then *P'* is true as well. Could we then have a quantitative measure of information that is qualitative as well? The answer is no. To see this, I take a little detour. In the literature on philosophical logic, it is customary to view a proposition *P* as a set of *possible worlds*. Possible worlds define ways things could be. There is a possible world in which Caesar crossed the Rubicon in 49 BC, there is one in which he crossed it in 48 BC, and one in which did not cross it at all. Let *W* be the set of all possible worlds. Denote by [*P*] the set of all *w* from *W* in which *P* is the case. Then *P'* follows from *P* if and only if [*P*] is a subset of [*P'*]. For every world in which *P* is one in which *P'* is true. So, the qualitative measure of information reduces to set inclusion.

This seems like a satisfactory measure of content. If it was also a quantitative measure of information we need it to be linear. This means that any two entities can be compared. Weight is a case in point. Any two objects can be compared in terms of their weight; if *O* and *O'* are objects then either *O* is heavier than *O'* or *O'* is heavier than *O* or they have the same weight. The same goes for length, duration, speed and so on. Always one is more than the other, or they are equal in quantity. So, we need to have that for any two statements *P* and *P'* either *P* follows from *P'* or *P'* follows from *P*.[34] Translating this into sets we need to have for any two propositions *P* and *P'* that either $[P] \subseteq [P']$ or $[P'] \subseteq [P]$. As Mancosu has shown, this is impossible.[35]

So there is no hope to go significantly beyond the Shannon-entropy. However, the dependence on language remains an issue. For notice that propositions must be coded using some language after all. So, we can approach the question of information also from a totally different angle. The alphabet of English has 26 letters, which come in two shapes (upper and lower case), to which we add the digits and a few punctuation marks. The number of symbols needed to code English sentences is somewhere between 64 and 128. Indeed,

[33] András Kornai: *Mathematical Linguistics. Advanced Information and Knowledge Processing*, Berlin, Springer, 2007.

[34] The third case is contained in the other two. If P is the same as P', then P follows from P' and P' follows from P.

[35] Paolo Mancosu: Measuring the Size of Infinite Collections of Natural Numbers: Was Cantor's Theory of Infinite Numbers Inevitable?, *Review of Symbolic Logic* 2 (2009), 612–646.

the ASCII code, which is the most widely used code in computers, uses 8 bits per letter, where one bit is reserved for error detection. Conveniently, 8 is a power of two, and the bits used to be grouped into so-called bytes, which are bit groups of size 8. Now let us look again at the sentence below.

(2.4) "Paris is the capital of France."

Including the period but excluding the delimiting quotation marks, this string contains 31 symbols. Thus, stored in ASCII this would require 31 bytes, or 248 bits of storage. This is a lot more than we have computed earlier. Of course, we did need only 7 bits, so we actually need only $7 \cdot 31 = 217$ bits.

If we use more than the necessary length one speaks of redundancy. An encoding is redundant if it is longer than it could be under an optimal encoding. For example, the 8 bit encoding contains one bit more than necessary (it is not used to encode part of the message) so it is redundant. If the question to be answered is

(2.5) "Is Paris the capital of France?"

then it is enough to answer "Yes." or "No."—or "I don't know." if one does not know the answer, a case which I shall ignore for the moment. If one answers instead

(2.6) "Yes, Paris is the capital of France."

then the answer is much longer than necessary and therefore redundant. As we shall see later, there are reasons to welcome redundancy, not the least because humans do not only talk to exchange information.

There is an important relation between information and energy, which was first discovered in thermodynamics. Consider a box that is separated into two equal parts by a wall. Place a single particle into one of the chambers. Then the particle will push the wall in the direction of the empty chamber to expand its space. This push can be used, it is equivalent to some energy. However, it can only be used if we know in which direction the wall is being pushed. So, given one bit of information (the particle is in the left hand / right hand chamber) we can extract a certain amount of energy which turns out to be

(2.7) $\Delta E = -k_B T ln 2$

Here, T is the temperature (measured in degrees Kelvin), ln 2 the natural logarithm of 2 (a number roughly equal to 0.693) and $k_B = 1.38 \cdot 10^{-23}\, J/K$ the so-called Boltzmann-constant. At a temperature of 20°C, which equals 293.15 degrees Kelvin we get

(2.8) $\Delta = -280 \cdot 10^{-23} J$

The dimension is that of energy, Joules, since multiplying degrees Kelvin with Joules per Kelvin gives Joules as the dimension. The sign is negative because that energy is needed to pay for one bit of information. However, it is important to keep in mind that the experiment suggests energy gain per bit, which is different. It has been estimated that in 2007 the entire internet contained 290 Exabytes of information.[36] This is $280 \cdot 10^{18} \cdot 8$ bits, which is $2.32 \cdot 10^{21}$ bit. Multiplying with the number above gives

(2.9) $280 \cdot 10^{-23} J \cdot 2.32 \cdot 10^{21} = 650 \cdot 10^{-2} J = 6.5\, J$

Now, since $W = J/s$, we have that a Joule equals Ws, one Watt applied for one second. So, the energy equivalent to the internet storage in 2007 is roughly 1/554 of a Watt hour. (Recall that $1Wh = 3600Ws$.) So this really is a minute fraction but nevertheless a demonstration in principle that we need energy to store this information.[37] Hilbert and López claim growth rates of 58 percent per year, which means that the consumption is fourfold in 3 years and a hundred times that amount within 10 years. It is clear that even from these modest beginnings there will arise sizable amounts of energy.[38]

Physicists have estimated the size of the universe at 10^{80} particles whose energy content is equivalent to 10^{40} bits of information. Taken together this means that the informational content of the universe is around 10^{120} bits. At growth rates of 58 percent per year we need a mere 500 years to exhaust that. It is evident that we cannot grow at that rate.

However, notice in addition that a mere bit of information is pretty useless if we cannot make sure that it can be recruited for the intended purpose.

[36] See Martin Hilbert – Priscila López: The World's Technological Capacity to Store, Communicate, and Compute Information, *Science* 332 (2011), 60–65. The prefix *Exa* stands for 10^{18}.

[37] I have tried to establish more realistic bounds. In the absence of tangible data, I took the power consumption of a hard drive (approximately 3 W) and compared that to the data throughput (10 Gb per second, which is roughly 100 Megabytes per second). This gives $3\ 10^{-9}\ J$ per bit. Comparing that to the above $280 \cdot 10^{-23}$ yields a factor of around 10^{12}. That means, in practice that one bit of information costs 10^{12} times (1 Million Million times) the thermodynamically minimal price. Applying that to the internet of 2007 would yield $1{,}800MWh$, an amount of energy that a single wind turbine could deliver.

[38] Hilbert – López: The World's Technological Capacity to Store, Communicate, and Compute Information.

For example, if we want to store the answer to a given question in a certain bit then that bit must assume a material shape (say, a register in an array of flip-flops), and it must remain stable over time so that we can retrieve its content when we need to. In fact, this is how it is done in a computer; the data chips give an address to each individual register, which contains one byte of information (typically a byte nowadays is much longer than 8 bits, mostly 32 or 64 bits). The register keeps that information for a pretty long time, so we can be sure it does not change. And we can retrieve its content as long as we know what its address is. Thus, the chip contains an ambient infrastructure around its registers which ensures that information can be stored and accessed. This ambient infrastructure comes at a cost much higher than that of its storage alone. This is why quantum computing will solve only part of our problems. While it can compute at zero costs the results of the computation nevertheless need to be recorded, that is, stored. Storing means in first instance to make them resistant to distortion, and that incurs energy costs.

Traces

The concept of a trace is the most important one for the purposes of this book if we want to really understand what role knowledge can play in our society. In his book *Der Aufbau der Physik*, Carl-Friedrich von Weizsäcker emphasized the role of traces for physical theory.[39] At the heart of his theory of physics is the way in which past events are able to influence present or future events. Buried inside that lies the question of whether or not it is possible in principle to have access to the past. For if the past is not able to physically influence the present then since knowing something requires being of a certain physical makeup we cannot know about this past. For if an event e taking place at time t has no bearing on the present then whether or not the content of any register in my computer is 0 or 1 is independent of it. So, the fact that it is 0 (or 1) cannot tell me about whether e has occurred. In this way, knowledge is embedded in a physical theory of causation, which, in turn, according to von Weizsäcker, is predicated on a theory of traces.

Intuitively, this idea makes a lot of sense even without its physical underpinnings. Anything that happens leaves a certain trace by which it can later be recognized. We can have knowledge of past volcano eruptions by the lava that has been dispersed, and by the climate change it has given rise to. These in turn are documented in the ring structure of trees. We know about the rhythm of ice ages, we know what humans ate in a particular settlement, and so on. All this is known because scientists have very carefully

[39] Carl-Friedrich von Weizsäcker: *Der Aufbau der Physik*, München, Carl Hanser Verlag, 1985.

investigated the effects of activities and have done very careful research in all possible regions of this earth including excavations and drilling down deep holes in stone and ice. Of course, the smaller the event in terms of impact and the more distant it is the less clear is its trace and the less certain we can be that it actually happened.

Reading traces is a matter of interpretation. It is one thing to say that an event *e* causes a certain event *e'* and quite another to say that *e'* happened *because e* has happened. Logically, the difference is whether *e* is necessary or sufficient for *e'*. Our mind is geared towards reconstructing entire histories as explanations. Some of the processes of interpretation are quite automatic. Leyton discusses cases of visual perception that lead to expectations of causation. If we see a flat can of some soda drink on a street, we immediately conclude that it has been flattened by someone, most likely by trampling on it.[40] We do not consider the possibility that the can has been artificially made to look the way it does. So, while one event (the trampling on the can) produces another (it being flat), it is not necessarily the case that its being flat has been caused by someone trampling on it; but we conclude that nevertheless. Leyton claims that symmetry has a say in it—it is more probable that symmetry gets broken than that it arises randomly. I point out that it also makes sense to consider the likelihood of something taking place. While the case at hand is pretty clear cut, it is not so easy to apply a rigorous logic to events that happened some billion years ago.

So, to discover the past means to read what is around us. Just as we know from Caesar's actions through certain books—written mainly by himself and copied in medieval monasteries (we do not have the original, only a copy from the 9th century AD)—so we know about the Neolithic past by reading the remains of settlements because years of scientific research have shown us how to do that. Whether our claims are accurate or not is one thing. Quite another is that unlike the book by Caesar most of the traces have not been fabricated with the intention to tell us something. It is us who read them and therefore they speak our language.[41] They become a mirror of ourselves. For all that reading the past is predicated of our scientific discoveries, that is, our knowledge.

But as a judgement is a momentary affair, how can we be certain to know something? Recall that knowledge requires reliability in judgement. One way of knowing I am reliable is to repeat the question. However, once we have issued a judgement it is history. Only our memory allows us to say that we have judged this or that claim in a particular way. Of course, when I answer

[40] Michael Leyton: *Symmetry, Causality, Mind*, Cambridge (Mass.), MIT Press, 1992.

[41] Adherents of a theory called "intelligent design" according to which the earth has been created 4,000 BC do say that the evidence showing that the earth is much older has been manufactured by God himself. Why he did so is however not so clear.

a question I am mostly certain that I am consistent in my judgements.[42] Whatever the justification for this trust in my intuitions, the question is why this is at all a reliable indicator. How could there be a connection with a judgement that I made yesterday and one that I am making today? What can be the connection between the person I was yesterday and the one I am today? What guarantees that the computer program will produce the same numbers again? In essence it is actually the same as the connection between the constellation of planets some years (or moments) ago and the one right now. The connection is that nature operates on immutable physical laws. They determine the state the universe is in here and now depends on what it was some time ago.[43] And these laws not only mean that traces of events remain to a certain degree to allow us to know about the past (and the future, in fact) but they also allow us to trust our own memory so that we can know who we are and what history we have ourselves. This continuity is necessary for there to be any knowledge whatsoever.

Now let us take a step back and ask: Is there a way to understand in purely physical terms what knowledge consists in? Neurobiologists say that the nervous tissue is composed in such a way that cells are connected with other cells by connections operating in a particular way. Once the network is in place it is the weights that control the signal propagation that are constantly being adapted. This is what constitutes learning. This network of connections and weights could then in principle be modeled mathematically and subjected to an analysis in just the same way as one could model a computer.[44] A computer consists in some hardware, an array of logical units wired together, and some software, a program that is stored in the memory of the computer for execution. The nervous connections are causing the humans to behave in a certain way. Had they been connected differently, the human would behave differently, just like the computer would compute differently had its hardware or the program been different. It is not hard to program a computer to answer "false" to the question whether or not 25 times 25 is 625. But that software once implemented takes on a physical appearance. It is a file, which in turn

[42] Williamson: *Knowledge*, denies that introspection is trustworthy in this matter. It is customary that once we know something then we also know that we know it. This is what Williamson calls luminosity. And he rejects it because knowing means requiring a safety margin for error, and that margin must be increased at every step of iteration.

[43] Notice that this formulation does not mean a commitment to determinism. It does not say that the state of the universe at $t < t'$uniquely determines the state at time t'. It just says that it determines that state to some degree, and that therefore certain reliable predictions can be made.

[44] Indeed, this is what is currently being done. The network structures are already being used as models of adaptive learning (for example Hidden Markov Models). Of course, for studying the brain these models would have to contain enormous numbers of connections.

is a sequence of registers containing a specific sequence of zeros and ones. Similarly, anything we know is coded as some architectural property of our brain. In sum, knowledge by necessity has a physical manifestation.

This means that it is in principle possible to predict what the next state of a computer will be. This can be done on an abstract level by looking at the registers and their contents and then calculating the next states; and it can be calculated concretely by looking at the physical shape of the connections as printed on a silicon wafer and then calculating the propagation of currents in that structure. The energetic and physical behavior of a computer are known in every imaginable detail. And they allow us to predict quite accurately, for example, how much a given device can store, how many pictures of a certain size, how many videos, and so on. If one wants to store more than that amount there is going to be a problem.

Similarly, one can estimate how much information we can put on a page of a given size. Once the size of letter is fixed, we can count the numbers of rows and the number of letters in each row and get an upper bound, at least for text. And this in turn gives us a good estimate on the storage capacity needed on a computer. So, we can either write text on paper or we can type it into a computer. And when we have it in a computer, we can even transfer it to a memory stick, put in into the cloud, or write it onto a CD. Whether we are changing the medium or not, each of the processes consist in copying by means of rearranging the physical layout of certain objects.[45] On paper the reconfiguration consists in putting ink spots on paper. In a computer the technique is more complex, but it still consists in rearranging certain parts so that the structure is a different one. How else could we read out the information?

Thus knowledge is coded by a particular form into which matter has been put. Though it is neither of these forms in particular as it can be transformed into different forms but cannot exist without being at least one of them. This material trace will be called the *trace* of the knowledge. To be able to understand what is expressed by means of the trace one needs to possess a certain *code*.

[45] Note that the word *information* contains *forma*, form.

CHAPTER 3

SOURCES OF KNOWLEDGE

The Four Sources of Knowledge

After having established the nature of knowledge we can now turn to the question where it comes from. In this chapter I will review four sources that each play a different role in human life and society. These are:

- experience
- culture
- science
- thinking

By culture I mean everything that involves transmission of knowledge from other people, be it in the form of talking, reading, writing, or other forms of interaction. Science is the form of scientific inquiry that we typically associate with natural sciences, though it is also used in psychology, sociology and many other disciplines. And finally, thinking is a process of synthesizing new knowledge from existing knowledge.

First of all, knowledge can be obtained by means of experience. This generally leads to what is sometimes called anecdotal or episodic knowledge. We experience a certain incident with animals, plants, people and they are first and foremost individual, one-time experiences. They can be generalized, of course. We can and do distil general laws from our experiences. The process of distilling such laws is not dissimilar from science, but the two need to be kept apart for reasons that will be discussed at length below.

One and the same content can be known through different ways. We can experience it directly, we can be told by other people, we can read it in science books, or we can deduce it from something else that we already know. The method of obtaining knowledge is unrelated to the content. However, the way we can come to know at all depends on many factors. During the first months a newborn human can only experience events, nothing more. There is very little transmission through culture. This only sets in when the child has learned to speak. Even then the scientific method is new, and so it has little

use for that method as a means of discovery. Thinking however seems to be involved at all ages, even though I should stress that the method of thought changes over time as well.

During childhood humans learn to speak a language. This is a powerful tool in human development whose role can hardly be overestimated. Consider how much a human learns throughout his lifetime. It is what the brain can digest the entire time that this human is awake. This puts a tight limit on human knowledge in general. Because knowledge cannot be passed on to the next generation every generation must learn the same things all over. All that can be transmitted is what can be shown here and now. Language changes everything. It allows to preserve knowledge beyond generations. Suddenly the body of knowledge can grow since knowledge can also be distributed over several people, and it can be recalled when necessary. However, this also has some downsides. The claim of cultural transmission on our knowledge adds to the claim of our own experience. What if what we are being told conflicts with what we are witnessing? Who is to be trusted? This problem is exacerbated with the advent of science. For scientists like to contradict common wisdom (it makes a scientist proud if he can do so), so what weight do we have to attribute to his claim that our traditional thinking is mistaken? With every step we are taking the number of voices is increasing. Just take the internet and the availability of thousands of different opinions on a given topic. Now pick a controversial topic such as nuclear power, or climate change. How many different people are there debating these issues and who is to be trusted? They can't possibly all be right.

The central tenet of modern societies is that the scientific method is the most reliable method for obtaining truth. For its body of knowledge is obtained through controlled experiments, and its data is (or at least ought to be) publicly available so that one can in principle check how the conclusions were being reached. A particular experiment is in principle also anecdotal evidence. However, it is conducted within a theory that makes predictions. Hence the experiment can be used to establish whether or not the theory makes accurate predictions; if so, the theory is considered to be more justified than before. And if not, maybe it should be changed or even rejected.

The possibility to perform all this is predicated on the ability to think rationally, however. It requires being able to use knowledge in the form of the affirmation of certain propositions to conclude that we need to affirm other propositions as well, or else risk inconsistency. To this end, mathematics and logic have been developed. But other forms of thinking are also important. They do not always lead from truth to truth but have a more pragmatic character. Such is the case with abduction, which is the reasoning from an effect to its cause. The cause need not be the only explanation, but often enough it might be that with overwhelming certainty. Logic has been

part of the arts in medieval times, in fact it has been part of the so-called *trivium*, the three-fold way. Logic was also used to reason about the nature of God. Nowadays logic is not taught as widely. In fact, in many corners of science it seems more important to memorize many studies than to think through a single one of them.

Experience

Every organism separates the world into two spheres: the inner world and the outer world. It aims at keeping the conditions inside stable (homeostasis). The surroundings are needed to make this possible. The environment is necessary because it supplies energy and materials. On the other hand, to keep the inside stable, also the surrounding conditions must be kept within certain bounds. Thus, organisms have sensors that tell them about the surrounding conditions. They can respond by moving elsewhere or by actively engaging to change the conditions. Humans are very skilful at doing the latter.[46]

To this end, humans must have an internal representation of the external world. It is an internal replica of the world, a mental map of some sort that allows us to navigate the surroundings and document the changes. These representations are propositional. They tell us what things there are and where we can find them.[47]

The sensory input on the other hand is not propositional.[48] The eyes receive light pulses on the retina. These constitute in the words of Gibson a 2-dimensional color theater. At the same time the ear is perceiving sounds, the skin is sensing some pressure and certain places, and so on. From these individual inputs the brain constructs the outside world, or better still, its own representation thereof. It works on several layers. First it organizes the input of the various sense organs. From the 2-dimensional color theater it creates a 3-dimensional picture in which there are certain things in the background from which others in the foreground are being separated. This is called spatial cognition or Gestalt cognition. On the other hand, the inputs of the various organs need to be compared with others and connected to a multidimensional representation. The image of someone moving his mouth and the talk that the ear is hearing will then be interpreted as belonging to one and the same event: someone is talking. This process

[46] Jakob von Uexküll: *Kompositionslehre der Natur. Biologie als undogmatische Naturwissenschaft*, Berlin, Propyläen, 1980.

[47] Ray Jackendoff: *Semantics and Cognition*, Cambridge (Mass.), MIT Press, 1983, claims that the brain has two separate systems for storing this information. There is a so-called *what-system* that stores things and a *where-system* that stores spatial information about their location.

[48] James J. Gibson: *The Senses Considered as Perceptual Systems*, Boston, Houghton Mifflin, 1966.

leads to the buildup of a representation in which there are things and events, and the senses are delivering the raw input to this process.[49] The underlying problem has been sketched already by Kant. In his *Critique of Pure Reason,* Kant argues that what he calls the transcendental object is a result of synthesis of sensual impressions. In modern terminology the synthesis is the process that collects individual sensory data and constructs the representation of an object (the transcendental object) out of those.[50] Importantly, Kant says that the mind hypothesizes the unity of the images as emanating from a single object. This is a hypothesis not warranted by the data alone. Thus, he calls a collection of data "objectual" if it can be construed consistently as partial sensations of a single object. Thus, the question remains for the mind as to which data must collectively be thought as belonging to one and the same object and which ones do not. This is the puzzle the mind has to solve.

Now, this is only the static aspect. There is a dynamic aspect as well. The world is constantly changing, and consequently our representation of the world needs updating. The mind needs cognitive routines to anticipate the way in which the world changes. For example, it can rotate an image of an object without the object itself (or its actual sense impression) rotating. It can anticipate where a car will be in a second from now and whether it is safe to cross the road. So, it can create representations of something that is not actually there.

In addition, the representation not only contains replicas of everything that exists around us. It also needs to contain an indication of where the subject itself happens to be. If I want to know where I am it is not enough to have a map of the city, I also need a pointer that says "you are here". It is often found on maps of a city that have a fixed place. Notice that it would mostly be possible to find out where one is situated on the map by taking a good look at the surroundings. This is why portable maps still are very useful. It is obvious that such a pointer exists in the mental representations inside our heads as well. For even if we could deduce where we are by comparing what we see with what is on the map, this knowledge is still there when we close our eyes. Most notably in known surroundings we will answer without hesitation to the question where we are. So, we are constantly updating the pointer as well.

These two dynamics are intertwined. The world is changing, and we are changing our place within it. And our sensual impressions of the world are changing accordingly. Certain objects are visible at this moment and invisible

[49] See Martha Fara: *Visual Agnosia. Disorders of Object Recognition and What They Tell Us*, Oxford, Wiley-Blackwell, 1990, for an account of how complex this process is and the many ways in which this can go wrong.

[50] See Theodora Achourioti – Michiel van Lambalgen: A Formalisation of Kant's Transcendental Logic, *The Review of Symbolic Logic* 4 (2011), 254–289, for a modern account of transcendental logic.

somewhat later. And this is either because they moved or because we did. And it is important to be able to distinguish the two. (Recall the irritation when you see the train next to you move slowly and you conclude that your train is leaving the station, only to find out that it is the *other* train that is leaving.) The cognitive mechanism is trying hard to make the sensual impressions congruent with what it thinks ought to be going on. Therefore, it does not conclude that the world has turned upside down when we are doing a head stand. Neither does it conclude we have become heavier when we just jumped into the air. The mechanism for this has been discovered in the 1950s, though it has been postulated as early as 1811 by Johann Georg Steinbuch. Basically, the so-called efferent copy is a copy of the signal that is sent to the motor system. It allows to compare what is sensed with what should be sensed anyway if that movement takes effect. And it allows to effectively control movement by comparing its effect with what was intended.[51]

The problems with these strategies can be experienced as well. Think about lifting a container with milk that you have incorrectly considered half full but that is in fact empty. What happens next is that we almost throw this container into the air. Our motor control system has been given the wrong instructions on the basis of this estimation. It applied too much force. However, as soon as that happens, the system tries to correct this.

These details are not mere curiosities. At present, the computer and information technology tries to replicate virtually all human cognitive abilities. It started with calculations, writing, playing board games, and expanded to all kinds of conscious activities. At present it has expanded this to include unconscious activities where humans are incapable of explaining what it is, they are actually doing. This includes driving, walking and playing soccer. When at first it was thought that the human mind is something of a pretty inefficient search engine, slowly it emerged that not only are the everyday routines pretty efficient. It also turned out that the subconscious tasks are horrendously difficult to perform. What is easy for a human is not necessarily easy to program. Nowadays this realization has led to some kind of coevolution between brain research and computer science.[52]

This however is just a small list of capabilities that humans and animals alike acquire very early on in life or possess them already at birth. Without them they could not survive a single day. Thus, the amount of knowledge that infants and children are acquiring is already enormous; that knowledge seems to be entirely trivial, but that is only because it shielded from our introspection. Knowledge of their own body and its inner workings belongs to that.

[51] Norbert Bischof: *Struktur und Bedeutung. Eine Einführung in die Systemtheorie* (2nd ed.), Bern, Hans Huber, 1998.

[52] See Arturo Carsetti: *Causality, Meaningful Complexity and Embodied Cognition*, Heidelberg, Springer, 2010, for an overview.

Humans as opposed to animals are born with a minimum of preprogramed mental structure. The first nine months are crucial in developing the basic "hardware", the basic nervous connections, which then form the basis of all future knowledge it will encode. How this is at all possible can be understood best if we look at the various human pathologies.[53]

This will make us appreciate the task for the human mind to construct order out of this chaos of sensory input. And like ourselves, the brain at a lower level needs to know how it exists in this world and how it needs to interact with it. This knowledge is not linguistic. It is prelinguistic. It is built in at such a deep level that we do not remember having had to acquire it at all, that it even is knowledge in the proper sense of the word. However, the progress of developmental psychology has brought to light the enormous complexity of the task of understanding reality from sensory input.

Once representations have been established, that is, once the external world has been mirrored inside the mind we can actually speak of knowledge in the propositional sense, the one that is the topic of the first chapter. In contrast to much of knowledge that I shall be discussing later on the kind of knowledge that is learned through experience is not coded in language. Even if we can talk about or write down what we have experienced, in actual fact we are not able to say exactly what it is that we have learned.[54] For example, we cannot really say what the rules of our language are and what meaning a given word has. Not even linguists do a really good job at that. It may happen that they declare a construction ungrammatical only to find themselves using it half an hour later. And the same holds for nearly everything that we do rather skilfully in the end. Take riding a bicycle. Although at one level it can be explained what one needs to do (where are the pedals, the brakes, how to change gears or switch on the lights), but at the motor level it is impossible to describe. The body needs to get used to the delicate balancing act while propelling itself not by means of walking but by pedalling. How exactly this is done cannot be explained, it must be experienced, and the body will learn it by itself.

There are therefore two separate reasons why knowledge that is acquired only through experience is limited. The first is that we can learn only what we experience directly, and so there is a gap due to the fact that our radius

[53] Oliver Sacks: *The Man Who Mistook His Wife for a Hat and Other Clinical Tales*, London, Duckworth, 1985. In science it is somewhat customary to consider the normal state of affairs as not in need of explanation. This often means, however, that we are not able to produce an account of what is going on. Fara, *Visual Agnosia*, comments on this in connection with vision: "One counter to this strange view is, broadly speaking, computational: Make me a formal model of how the mind-brain transforms peripheral transduction [...] into the full perceptual richness of the three dimensional objects in space."

[54] Part of this problem is the fact that we may not know what it is that we know. This is the failure of introspection or lack of luminosity mentioned in Footnote 42.

of experience both in space and time is bounded. This means that it is not only the amount of knowledge that is limited; also, its content is. Nothing that happened in the past, and nothing that happens far away can be subject to our experience. A second limitation is that this way of learning results in prelinguistic knowledge: it cannot be verbalized easily and therefore not shared. Thus, the addition of language enormously expands the volume and the coverage of human knowledge.

Experience is of course not only responsible for collecting data on the world. It will also lead to uncovering regularities about the immediate surroundings. This includes elementary physical laws (gravity for example) and facts about the behavior of people. This is not a trivial matter. The behavior of other people, especially adults, is a great mystery to children. Often, we tend to dismiss the problem of maneuvering the adult world as trivial, but it is not. We only know about our own feelings and we need to learn to read people's minds, as it were. In the literature, the technical term "theory of mind" is being used. It refers to a theory that each of us has that helps us explain why other people behave the way they do in absence of the possibility to experience their feelings or having their beliefs. To be able to rationalize about other people and come up with a theory is a very complex ability.[55] At an early age children do not have such a theory, they need to develop it. What helps this is an endless curiosity about other people. This interest actually never stops. The newsstands tell us a story about how much people are obsessed with knowing about others, neighbours, politicians, celebrities and so on. The social media are another indicator how important it is to be socially engaged.

Learning never stops. This is almost trivial for one part of it is remembering all the details of our everyday life. The plasticity of the brain may diminish over time, but it actually never stops to learn.[56] Of course the kind of knowledge that we mostly pick up is neither culturally nor scientifically of great importance. Nor is it morally prefigured. We more or less record what we see and remember it more or less against our will. The idea is that one never knows what will become useful later on. The same is true for the internet. Initially, when storage space was expensive only very important information was being stored. Now that the price of storage is basically zero,

[55] D. G. Premack – D. Woodruff: Does the Chimpanzee have a Theory of Mind?, *Behaviour and the Brain Sciences* 1 (1978), 551–526, actually claim that what distinguishes the humans from the chimpanzees is that they possess a theory of mind. One seemingly small difference between humans and other primates is the white in the eye. It allows others to detect what we are looking at and so they can guess what we are interested in or up to. This detection of human gaze is actually enormously accurate.

[56] There is a special region of the brain responsible for recording experiences, the so-called episodic memory. It is possible that only this region is damaged, in which case remain "frozen" in a state of their life. They can experience more but never remember more. See Sacks: *The Man.*

the internet is filled with all kind of trivial, useless, sometimes bizarre, even dangerous and illegal information. For some time the idea of an internet that never forgets anything was being praised as a virtue. Nowadays people are not so sure anymore. The situation is quite similar to that of people with perfect memory. They remember absolutely everything they experience with machine-like precision. The downside is that their brain is unable to select which of these facts are actually important and which ones are not. Typically, when something is important it remains more vivid and can be accessed faster. So, it is not necessarily the case that we forget certain things—although that too is the case—but that the brain is constantly trying to filter out the irrelevant things and puts them aside. The criteria for "importance" may have nothing to do with our culture (taboo stories are somehow more easily remembered even though they should not be). The reason for that is clear. Access to knowledge becomes difficult the more we know, so there must be some sort of ranking between all the things we know. The more important the faster it needs to be accessible.

Experience is not the only source of knowledge, but it is the most fundamental of them. Even scientific knowledge is ultimately based on basic experience. The scientist performs an experiment and witnesses how it plays out. It would be ridiculous to run an experiment when no one is there to make notes. Of course, nowadays much information is gathered with the help of machines, but at the end of the day someone takes a look at the machine to see what data it has collected. Thus, the divide between scientific knowledge and experience is not whether or not we should trust what we see. It is rather what it is that this tells us. Is our everyday experience a good guide or is it deceiving us?

Scientists are quick to dismiss our everyday experiences as merely accidental. They are quick to point out that we may be biased and do not understand the nature of circumstances we are in. So, we may fail to witness global warming because we live in a particular location where the temperature did not rise at all for some years. If we therefore conclude that global warming is a hoax, we have indeed fallen into a trap. On the other hand, a scientist often enough just sits in a room behind a computer screen and is prepared to think that his data is superior to what people witness around them. But take a mountaineer in the Alps, or a native of Greenland. These people are out in the nature every day for decades. And their parents and grandparents were probably doing just the same. Thus, they might actually know how the climate was one hundred years ago. And so, they can actually tell us something significant about the climate as well. To claim that their knowledge is somehow accidental would be to say that they were unable to understand the significance of what they were seeing. And it would be to declare that even though the climate is a very important factor in their lives they actually do not understand it. This is hardly believable.

Another problem is that scientific knowledge needs to be statistical. But it is applied to specific cases. This is the case in medicine. Whether or not some drug is effective against some disease is established through tests. And they never yield unanimous results. Rather, they establish a statistical likelihood that a given drug is effective. On the other hand, a patient while knowing little about medicine actually knows more than anyone about his own body. If a patient after taking a drug declares that it did not help him that presumably the doctor would believe him and search for some alternative. He would not insist that the patient is not a trained doctor and therefore unable to tell whether or not this drug can cure him. While we cannot exclude such behavior, it is illogical. Why would the patient be seeing the doctor in the first place if it were not for certain symptoms? Why did the doctor trust those? And what is the method for gathering data on the efficiency of a drug? Surely, we also want to know whether people feel well. Or at least the patients themselves will insist that that is important.

While these are clear cases, there are many more that are not clear or not understood to be unclear. One problem is that even though many disciplines call themselves "sciences"[57] they often do not apply the scientific method and more or less openly use anecdotal evidence. For example, in economics it is said that deflation has devastating effects on the economy and should therefore be avoided at all costs. The biggest piece of evidence is the big recession in 1930. But didn't Japan experience deflation in the 1990ies? And what would be a strategy to avoid it? Should we trust marxist theories or rather neoclassical theories? Is there a way to know, for example by modelling the economy and then seeing how it develops? Exactly this is the criticism of Keen.[58] It is not that economics could not be a science. It is that economists often rely on their intuitions when they should be doing something else. For if all they do is just intuit then how can they insist that their intuitions are to be trusted more than those of a layperson?

This is a very important discussion. For every time laws are passed or projects are initiated the question arises whether that is legitimate. Laws for protecting the environment are a case in point. Everyone likes to protect the wildlife but as soon their own lives get compromised over this they become hesitant or simply change sides. So, they start asking for scientific evidence. On the other hand, locals often know better what needs to be done. Regulations may be better on average, but the effectiveness depends on local circumstances.

[57] Even more so in German than in English, I should add. What is called "Wissenschaft" ("science") in German is often just labelled "studies" in English.

[58] Steven Keen: *Debunking Economics. The Naked Emperor Dethroned*, London, Zed Books, 2011.

Culture

To share knowledge, you have to be able to talk. Otherwise, all you can do is point at something here and now. It is possible to teach an instrument simply by playing it and expecting the student to learn by watching you. This is effectively how much of education still works today (even master classes in music). And this is by the way how animals pick up knowledge that is not instinctive.[59] And yet that method is severely limited, since it requires physical presence. So, knowledge transfer (which is what this section is all about) is almost impossible without language.[60] While transfer is possible without language, the ability to use language changes the cultural transfer beyond recognition. We can talk about events that happened in the past and add that history to our knowledge. We can talk about places that are far away. We can talk about situations that are merely conceivable and speculate and argue in public about them.

One advantage is that it eliminates the need to make certain experiences. For example, in order to learn about the danger that certain animals pose, like snakes or bears, it is no longer necessary to wait until one of them shows up. We can also talk about what needs to be done beforehand. This is very useful since a present danger leaves very little room for mistakes.

Knowledge collected from episodic data tends to be unreliable as well. If on various encounters bears have shown to be harmless for me how reliable is that? Again, it would be foolish to try and put the hypothesis that they are harmless to a real test. A much better way would be to collect stories throughout a longer time and more varied sources until a more general picture emerges. The ability to share knowledge by using language has the advantage that we do not have to rely on particular experiences, and that we can even out irrelevant factors. Knowledge can be passed on through generations. Jared Diamond reports that a tribe had knowledge of what plants were edible (and how to prepare them) that they would normally not touch, only in severe circumstances. Such knowledge might turn out to be needless for a long time and so it would disappear in case circumstances are good for a few decades. Likewise, these people claimed not to know how to make certain artifacts until they could no longer get them, a point at which they started making them themselves.[61]

[59] Many skills that animals possess that were once believed to be instinctive turned out to be learned by direct teaching. One example is bears. Bears that have been brought up in a zoo do not necessarily know that a river is a source of food. So, they would not go fishing.

[60] Language is another difference between humans and other animals. For Jared M. Diamond: *The Third Chimpanzee: The Evolution and Future of the Human Animal*, London, Hutchinson Radius, 1991, it is the most important difference.

[61] Jared M. Diamond: *The World Until Yesterday: What Can We Learn from Traditional Societies*, New York, Viking Press, 2012. That raises the obvious question whether and how they preserved that knowledge. Perhaps they did learn to make these kind of things but declared that knowledge secret.

Evidently, even the transmission of knowledge by means of language has its limits. There is a limit on what a single person can memorize. And, if in the case of traditional societies people more or less stayed put in their village for all their lives then what that mode of cultural transmission offered was a distribution of knowledge within a small tribe or a village at best; and it offered the possibility of preserving knowledge indefinitely.[62] The kind of society that this produces is pretty egalitarian since there is a limited way to arrange for private knowledge. Such a model can be seen still today, for example with the Amish in the USA.

The middle ages actually produced a new development. Some knowledge was now in the hands of guilds. These were organizations consisting of particular craftsmen. They did exist also in Roman times (known as collegium or corpus). These organizations not only served as a trade union, or an organization that furthered the interests of their own people, but they were instrumental in the transmission and development of knowledge concerning their craft, for example by organizing the details of apprenticeship. In this way the guilds were able to exert quality control both over what knowledge people commanded and over what they produced. While division of labour was almost nonexistent within the guild, it was obviously practiced between various guilds.

Guilds were also instrumental in establishing the first universities in Europe, Bologna (1088), Oxford (1096), Paris (1150) and others. Universities offered courses in law and medicine (for those who wanted to make money) and general arts (logic, philosophy, theology). However, the biggest share of education was offered by the church and the monasteries. They took over the education from childhood on, in much the same way as schools do today, although they took in only a fraction of the children at that time. Most sons were supposed to become farmers anyway, not to mention girls, who could at that time practice only a very limited number of professions.

Schools did exist already in antiquity but were restricted to the rich who could afford to pay for tuition. Schools in the modern sense, which teach every child from a certain age on, became the rule in Europe quite late. In Germany, certain regions already had general schools in the 16th century, but schools only became mandatory in 1919.[63] Schools, like universities, were founded with the mandate to educate children, which effectively meant passing on the knowledge to the next generation. At first they were in the hands of the church and the monasteries. The reformation put an end to that, since monasteries

[62] It is by now known that written records do not necessarily last longer than oral tradition, which can pass on memory of events that have happened more than a thousand years before.

[63] Franz-Michael Konrad: *Von der Geschichte der Schule. Von der Antike bis zur Gegenwart*, München, C.H. Beck, 2007.

were being deserted. In this situation the local authorities (counts, kings) intervened and started to run the schools themselves, often with the approval of the clergy.[64]

Schools like many other institutions replaced the somewhat informal learning in the family or village. It did to some degree also lead to a situation where families expected the school to do the teaching for them. The authorities favour schools for the reason that they can use them to control the educational standards. It is far easier to implement a standardized teacher training and a school curriculum than to mandate uniform skills among the population and let local communities choose how to meet these standards. This is even more relevant in a situation where there are more and more types of occupations and very little hope that children will find someone in the neighborhood to teach them the necessary skills. This can now be left to higher education. However, higher education can only be effective when it can rely on a standardized school education. This is one example of what Beniger calls *preprocessing*.[65] Preprocessing is a way to reduce complexity by increasing the complexity and standardization of products for further processing. In education, higher education is the actual end. It prepares students for their profession so that they can be useful for society. The more difficult and diverse this higher education becomes, the more important it is to start at already higher educational level. And the more important it is that the high schools deliver uniform standards.

To become a high school teacher, one must by now have a university degree. In order to secure standards in schools it is in theory enough to implement uniform standards in the teacher education. Indeed, the introduction of new methods such as set theory ("New Math") or the Whole Language Method started at universities and were included in the course program so that they could then be taught as part of the school curriculum. This top down approach has its downsides, however. The biggest problem is that many people feel that the government simply mandates what teachers have to teach. And not all that they are mandated to teach is popular. New Math was highly unpopular also for the reason that it introduced abstract methods that many felt were superfluous. Evolution theory is unpopular with certain religious groups, and some sex education is seen as teaching downright immoral values by some parents.

Schools and universities are not the only institutions to pass on knowledge. However, let us leave this subject and turn to another problem, namely that of culture as an organizational framework to articulate knowledge. Notice that what we possess in the form of propositional knowledge that can be

[64] Martin Luther was explictly asking the nobility to help run the schools.

[65] James R. Beniger: *The Control Revolution. Technological and Economic Origins of the Information Society*, Cambridge (Mass.), Harvard University Press, 1986.

elicited is only the tip of the iceberg. Most of what we know is not of this form. And much of it, even though it is of propositional form, is not available for elicitation. We cannot determine all that we actually know. In this connection, the cultural tradition acts as a framework that allows us to compare and evaluate what we see and determine whether or not we are actually seeing it the right way or whether we are living under an illusion. Often enough culture incites our curiosity in asking questions or proposing the existence of things that it deems important to ask or propose. Examples are the ideas in science (theory of evolution, relativity theory) or objects (aether, atoms, molecules). Not of all them are equally successful. But they have all inspired innumerably many people to think about certain ideas to see how they structure the way we think about the world and many other areas of life as well. To see what I mean by that you just have to recall how much the ideas of evolution biology have influenced our thinking (for example in economic theory) or how much people have been inspired by the ideas of relativity or quantum theory.

It does not only happen because some of these ideas and theories are really hard to understand in detail. It is also because there is an inherent need to make sense of everything that surrounds us. Many people believe for example that nothing happens by chance. And that everything that happens to some person has actually been caused mentally to some degree by that very person. It follows that people who have been robbed are to some degree responsible for having been robbed. A variant of that theory is the theory of karma, which says—in very simple terms—that what you experience in this life is determined by what you have done in your previous life. Whether or not we believe this particular theory does not matter much for the point I am making. For it is true that regardless of what we believe, we are always trying to make sense of the world around us. We are carrying with us certain ideas and theories about what should happen and then compare that with what actually happens. This is really something of a triviality. If we did not do that then there would be no way our experiences could give rise to a useful theory of the world. If we did not compare the behavior of this bear in front of our eyes with bears that were previously encountered, we could never formulate (and test) a reasonable theory of what that bear will do next. Scientists sometimes propagate the idea that laypeople should enlighten themselves by reading the scientific literature instead of trying things out themselves.[66] What matters really is whether the belief system is useful. Of course, we want

[66] I'd like to give an example here, namely the discussion on efficiency of solar panels, see the blog post by Alice Friedeman (http://energyskeptic.com/2017/tilting-at-windmills-spains-disastrous-attempt-to-replace-fossil-fuels-with-solar-pv-part-2/). As usual in such discussions there exist numerous apparent contradictions, and it is pretty impossible to discern the truth. Many scientific indicators are either insufficient at closer look or the numbers can't be obtained in practice.

it to be true. That people are not always telling the truth, also in science, is an additional factor to be aware of. It is never good advice that people should simply trust rather than think for themselves. However, it is useless advice to tell them not to believe at all. We shall return to that problem in the last chapter of this book.

A large part of culture is (or has been at least) telling stories. Most of them are not openly educational. Many of them are not told because they are true but because they are beautiful. But even those stories do have some content, and it is important to understand what that content is and how it relates to the world, our culture and its values. And in the modern age people do mostly tell stories that are supposed to inform others more or less directly.[67] Such is the case with newspapers and journals. It is part of the ethos of journalists that they report claims only after having verified them using independent sources. How they do that is part of their training. Not all stories support the same idea; they are directly or indirectly in conflict with each other. Not all that is being told is of equal relevance. The more information is available the more it becomes vital to distinguish what is important and what is not.

Behind the question of truth however lurks another one that is far harder to address. It is that there are different ways of looking at reality that are incommensurable and yet equally valid. Ludwik Fleck coined the word "thought style" *(Denkstil).*[68] His claim is that the thought style is a way to make sense of the world and to systematize experiences. There are many thought styles, different societies have different ways of thinking and these ways are also subject to change. To illustrate this difference, think of what counts as an explanation. Modern science consider an explanation that which reduces an object to smaller parts and thereby explains its behavior by an interaction of its parts. In other words, we are atomists all the way through. The converse is almost unthinkable. An example of a non-atomistic explanation that comes to mind is the so-called "Gaia hypothesis" by James Lovelock.[69] According to Lovelock, earth is a system that is self-balancing (homeostatic). It therefore reacts to disturbances whether induced by nature or humans with a counteracting force. I would venture a bet that the immediate reaction is this: OK, but we can explain that homeostatic behavior with the use of some differential equations that govern the behavior of the atmosphere. So, that would again mean going to the level of particles. So, we do not accept this as an explanation but rather as a first step towards an explanation. Now the question that arises is the following. Is there any

[67] And most stories get told to make money. This is not a good state of affairs, since what sells is often not what is educating us. The film "industry" is a good example.

[68] Ludwik Fleck: *Entstehung und Entwicklung einer wissenschaftlichen Tatsache: Einführung in die Lehre vom Denkstil und Denkkollektiv*, Frankfurt am Main, Suhrkamp, 1980.

[69] James Lovelock: *Gaia: A New Look at Life on Earth*, Oxford, Oxford University Press, 2000.

intelligent system larger than humans? Is there an intelligent system of which humans are a part? If we believe the modern atomism then the answer is negative, *by definition*. Atomism of this sort is incompatible with a conception of God as an immanent being, for example; again, this is almost by definition the case.[70]

To prove his point Fleck looked at the history of the concept of syphilis.[71] Before it was possible to isolate pathogens, it was more common to classify diseases using the symptoms. After that it was "clear" that they must be classified according to their pathogens. So, what is now classed as syphilis has been seen as several diseases two centuries ago, and vice versa. But there is no way in which we can say that one of the classifications is correct. Similarly, there are many ways to look at a forest. For some it is a place for recreation, for others a source of money, for some it is a dangerous place. And this is their belief regardless of any previous experience. Until recently, robbers and other illegal people liked to hide in forests, which were therefore not without justification regarded as dangerous places. Traveling, as much as it is enjoyed today, has been an expensive, tiring, and dangerous activity. Evidently, there is no correct way to look at the notion of traveling. It is however clear that we look at it in a particular way, regardless of our personal experiences. This is the way the society at large looks at it. So, even if a particular trip turns out to be not so pleasant, we would say that we were unlucky and traveling is enjoyable after all.

Or look at the way society treats genders. Some decades ago, women were considered not intelligent enough for certain subjects, now it is commonplace to claim that they are as intelligent as men. Some still claim, though, that the variation in intelligence is higher in men so that we should expect more male geniuses, which is contested by many others. Society has not shifted to the viewpoint that an egalitarian treatment must result in a fifty-fifty balance in all relevant aspects. While the first claims can be tested (and have been) the last one is an article of faith that is used to judge which development in society is good and which one is not. Also, while many societies stress the difference and complementarity of the sexes the prevalent belief is that they are actually identical in virtually all respects.

Cultural knowledge is not restricted to facts. First, there is art and music, about which we also possess knowledge. The so-called cultural heritage is not only a collection of buildings and pieces of art. It also consists in stories, folk songs, traditional cures, moral values, and so on. To know these was

[70] Notice, however, that naturalism cannot explain the emergence of consciousness, see Holm Tetens: *Gott denken. Ein Versuch über rationale Theologie*, Stuttgart, Reclam Verlag, 2015. So it isn't simply the case that our version of naturalism is a better way to explain the phenomena, because there is none that we possess.

[71] Fleck: *Denkstil*.

an integral part of belonging to that society. In times when there was no radio and no sheet music it was an integral part of education that one knew the songs by heart. There was no other way they could be preserved through the generations.[72] Today that knowledge is in principle dispensable. However, what happened was in fact something different. As more and more music gets written, the body of music is increasing. While up until the 19th century music was written to be performed at some occasion and then pretty much forgotten, nowadays societies try to preserve any musical piece whatsoever. Since it is impossible for any individual to remember all of that, music has become diversified. There are any number of different styles, and they all have their canonical music which one must know in order to be a respectable connoisseur of that style. Take classical music. There not only is classical music as a whole, with its popular pieces by Mozart, Beethoven, Schubert and so on. There is renaissance music, modern music, romantic and late romantic music, electronic music and more. A visit to a larger music store gives testimony to this diversification.

This phenomenon is not restricted to music. It exists in literature, visual arts, and, curiously enough, also in religion and in the moral sphere. For centuries the church has tried to have everybody adhere to one and only one belief (Catholicism). While that led to internal differentiation the reformation has effectively put an end to that dream. The enlightenment then proceeded to dismantle religion as a foundation of society even further. Rulers no longer forced a unique religion on their subjects, some even explicitly stated that everybody was entitled to his own religious belief.[73] Today, it is hotly debated whether schools should teach religion, and if so, whether they shouldn't just give an overview of them without any judgement. This had an effect on the philosophy of law as well. When moral judgements are no longer uniform it becomes harder to justify any particular legal code. Instead one has turned to positive law: there are rules by which laws get sanctioned, and any proposed law that successfully underwent this process is valid and needs no more moral underpinning.

Finally, I shall turn to a subject matter that seems to be as far away from knowledge as one can think of: mythology. There is two kinds of myths. The first are the foundational stories, the myths that explain how the world came to be, and why it is the way it is. This is what Claude-Lévy Strauss was dealing with. Some think that these kinds of myths do contain important

[72] A recent lecture by Richard Heinberg on music in the Post-Carbon era makes the point that complexity in music is dependent on a lot of material artifacts that could soon disappear (https://www.youtube.com/watch?v=m_kAumEqwmY).

[73] Friedrich II. ("Frederick the Great") is quoted as having said concerning religion: "Jeder soll nach seiner Façon selig werden." ("Everybody should attain happiness according to his own way.").

messages.[74] And there is a second type of myth, which puts it closer to the thought styles of Fleck. They are a way to communicate or reflect. This is how Ernst Cassirer looked at them.[75] An example is the myth in the works of Plato. It is not a characterization of content but a characterization of the form in which it is put. Plato uses the myth as an alternative way of explaining.

Plato did not think highly of folk tales. His position was much like the modern one. We mostly think of myths as scientifically unjustified stories or beliefs. This is the usage of the word that is quite commonplace today. You call some claim a myth, implying that it is both false, and without rational basis. I will call them mythologies after Roland Barthes.[76] Basically, for Barthes anything can be called upon as a sign and then imbued with new meaning, be it red wine (which has very special relevance in France), the Romans in films, or washing powder. In his view, these signifiers become focal points which allow societies (or rather in his view the upper class) to produce new images and values. Thus, speaking with Barthes, it is a way to naturalize beliefs by associating them to common objects.

One does not need to buy into the idea that mythologies are created by the upper class, as Barthes has it. What is more important is the aspect of focussing on certain explanations, things or actions and changing their overall significance. To give some examples, the notion of a Murphy-type law (what can go wrong, does go wrong) or of a Peter-principle (people get promoted because of their performance until they are no longer overqualified, i.e. typically under-qualified for their job) are labels which we put on happenstances or stories. They allow us in this case to give an explanation where none could otherwise be given.[77]

As much as one would want to eliminate such mythologies, it is impossible. There are various reasons for that. First, even though something is classified by someone as a myth that does not mean that it is actually false, just like claiming something to be true doesn't make it so. Earlier I have spoken about the impossibility of knowing whether something is in fact true. The mirror image is that it is impossible to know whether something is false. So, whether or not a given claim is a myth or whether or not something is knowledge is open to debate. We can see this clearly in science. There are many claims that were once established scientific facts that we now know or at least believe to be definitely false. Nobody knows what the future will bring in this regard.

[74] See Jordan B. Peterson: *Maps of Meaning. The Architecture of Belief*, London, Routledge, 1999. He quotes among others Carl-Gustav Jung and Mircea Eliade who have studied the archetypes of these stories very carefully.

[75] Ernst Cassirer: *Philosophie der symbolischen Formen Erster Teil. Die Sprache* (2nd ed.), Darmstadt, Wissenschaftliche Buchgesellschaft, 1953.

[76] Roland Barthes: *Mythologies*, Paris, Editions du Seuil, 1957.

[77] Similarly the phrase of exceptions that prove the validity of a rule, which is semantic nonsense if taken literally.

Their case could be reopened at any time. But if that is the case where do the alternative hypotheses come from? How can it be possible to think that an established scientific fact is actually false? It could only be because we allow ourselves to talk about ideas irrespective of whether they are true or false; and that in addition myths persist even if current scientific evidence has it that they are false. Because, as funny as that sounds, the mythologies are necessary precisely because science is fallible. To use an analogy here: to think that we should aim for a society where people's beliefs are based on scientific truths is tantamount to building society with only adults in it. This is impossible. Myths provide the unconscious pool from which we draw our inspiration. Somehow, by whatever reason, we establish a belief. Before we have done any explicit testing all we are left with is our instinct. Something tells us that we are onto something real. That provides a powerful motivation for doing scientific research. There seems to be no other way for humans to generate hypotheses. If we want to dismiss mythologies, we must ask: how is it that we can come to believe what we believe if not through experience and hypothesis testing? And what is better to our advancement than constantly challenging the established consensus?

A similar point can be made with respect to prejudices and stereotypes. The idea that societies are full of biases has attracted a lot of research in social psychology. Curiously, as Jussim notes, one has very often failed to address the question whether the stereotypes were actually justified or even true.[78] The consensus view, now even more accepted, is that stereotypes are false. Moreover, he notes the conundrum people face. If a stereotype is a belief about a group that is actually false, then labelling any belief a stereotype requires some argument. Instead if stereotypes are just any kind of belief about a group of people then they may be true. But shouldn't we welcome true beliefs? And if so, which ones are they?

Now in order to be able to establish correct views it is necessary to put out hypotheses. They may turn out to be false, but if we do not have any faith in them or at least a commitment to their relevance, we will not make any advance in knowing what they are worth. So, having a hypothesis involves a risk, that of a misplaced bet. But it is risk that we have to take. As strange as that may sound: if we ban mythologies, stereotypes or prejudices, we will pay a heavy price. Society would become totally irrational. It will not be able to properly sort out its beliefs.[79]

[78] Lee Jussim: *Social Perception and Social Reality. Why Accuracy Dominates Bias and Self-Perception*, Oxford, Oxford University Press, 2012.

[79] Two remarks are in order. While modern societies have many outlets for irrational behavior, they still more or less adhere to the idea of a rational society which formed itself in the 19th century. This can be exemplified with religious belief. On the one hand, religious beliefs are deemed irrational and should therefore not be used to justify moral values even less scientific theories. On the other hand, religious beliefs do not disappear just because we

Science

Although many disciplines are called sciences, I shall restrict the usage of that term here to the empirical sciences. The ideal science is physics. Physics deals with the laws of nature. These laws are mostly immutable and universal, they are beyond time and space, like the laws of Newtonian motion, quantum theory or of general relativity. They tell us how the world is (or can be). There are no exceptions.[80] But not all of the physical laws are like that. A different type of physical law is Hooke's Law. It states that the length of a spring is proportional to the force that is applied to it. This length can be computed by multiplying the force with a constant that depends on the spring, and nothing else. This law is not universal: it applies within a range of values for the force. It also requires certain values for the surrounding temperature, for example. Additionally, it needs to be clarified what counts as a spring and what doesn't. But nevertheless, even Hooke's Law is universal within these bounds. It applies to *all* forces within the range.

To establish or verify this law, we need to perform an experiment. Ideally, we choose a variety of springs, and apply various forces to them, measuring the extent to which the springs get longer depending of the force. While we cannot even hope to check the law for all values, there is no bound on the values for which we can check it.

However, all that is not enough. Science is a social endeavor. In order for the results of our experimentations to be accepted by other people we have to convince them. For that it is not enough to claim to have performed the experiments, we should have proof of that. One way is to perform the experiment in public (like Otto von Guericke with his Magdeburg hemispheres). Another is to document them and allow others to perform them for themselves. We should also keep in mind that we might have performed an invalid experiment or drawn the wrong conclusions from it.

The scientific method has been invented to avoid all these shortcomings. One important ingredient is that every step in the experiment as well as the reasoning behind it must be recorded. This is important to avoid the haphazard nature of uncontrolled learning. This of course needs justification

---

do not welcome them. They are in most instances more fundamental than any particular scientific theory. To think that science can be used as a basis for moral judgement is to invert the nature of our human psyche. It is *because* we attach value to something that we think it worthy of study. And the second remark pertains to what I call the infantilization of science. In the absence of a clear notion of what is important and what isn't we are without guidance as to wherein to invest our efforts. Since any activity is performed under external constraints, whatever we are investigating goes at the expense of everything else that we are not doing at this moment. We cannot and should not pretend that there is no value system over and beyond science.

[80] Not even those that prove the validity of the laws...

in the first place. Say we are planning to put a hypothesis H to test, for example Hooke's Law. Typically, there is evidence in favour of H as well as against it. A trivial counterexample against Hooke's Law may consist in producing a spring that is not of good quality (the thickness of the metal varies) or whose shape is irregular. In virtue of the fact that there may be any number of reasons for its failure we must be prepared that the law is not universally valid and that we shall find counter-evidence. What will convince us or any other person that H is valid? Will it be enough to just count the number of times the outcome supported H against the number of times that it didn't?

Philosophers of science have tried to answer that question. The central idea is that of inductive verification.[81] Suppose that H is true with probability p; it follows that it is false with probability $1 - p$. This means that the likelihood for it being true is p. In general, we do not know what the probability for a given theory is. To be sure, we are *not* asking about our confidence, as confidence is subjective. We are asking about some objective measure using which we can identify how far we are off the mark with our hypothesis H. (Recall that most claims are not universally true, they are just likely true. An example is: *This drug is a cure for disease X*.) Suppose that we run n independent trials of which k support the law, while $n - k$ do not. Then on the basis of this experiment we should put $p := k/n$. This is because given that experiment, the most likely value for p that makes us observe k out of n confirming trials is k/n. If we perform 28 trials and 23 supported H then we put $p := 23/28$. If 28 of them support H, we put $p := 1$. This is explained as follows. Probability theory tells us what the likelihood of a value for an outcome parameter $\mathcal{O}$ is given a certain value for an input parameter $\mathcal{I}$. We perform an experiment and observe a value v for $\mathcal{O}$. Then we ask: *What value should we assume for $\mathcal{I}$ so that the observed value v for $\mathcal{O}$ has highest probability?* This turns out to be same as asking: which model is the most likely given that our experiment yielded this result?

The theory of inductive confirmation goes beyond that in the following way. It associates a probability p for a hypothesis, which initially is 1/2 if nothing is known. We perform an experiment and ask the following question: *Given that our initial probability of likelihood for H was p and we performed our experiment, what probability p' of H should we assume now? p* is called the *prior probability*, *p'* the *posterior probability*. So, a single experiment can make us confident to some degree in our hypothesis. Another experiment can increase or decrease that confidence, depending on outcome. A cornerstone of this theory is the Limit Theorem of probability theory: If we run an experiment infinitely often then the estimated probabilities will coincide with the real probabilities. In other words: if we repeat the experiment often enough,

[81] John Earman: *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*, Cambridge (Mass.), MIT Press, 1992.

we can approximate the probability to any degree—and with any degree of confidence. That is, we can increase the reliability of our estimates and increase their precision in the long run.[82]

But how do we actually run an experiment? We can either wait until the circumstances are right and we can make an observation that counts towards our hypothesis. If we want to know whether cats like to eat rabbits we may simply wait and see if we observe a cat that is chasing a rabbit. The biggest problem with that approach is not that we have to be patient. The problem is that we do not know how long we should wait until we may say that the hypothesis is confirmed or disconfirmed. Maybe we have picked an unfortunate spot and the are no rabbits to chase. No wonder then that we did not observe any rabbit chasing cats. But did that tell us whether or not cats like to eat rabbits? Hardly.

So maybe a better way would be to take a cat and a rabbit, put them into a cage and wait until something happens. That is to say, rather than wait for circumstances to be of the right sort we actively make them so. The advantage is that we can plan ahead when we want to do the experiment and that we have control over the conditions under which it takes place.

This is why experiments are so popular. Nearly all disciplines routinely run experiments: physics, biology, medicine, psychology, linguistics, even philosophy. The body of results that are obtained experimentally is growing fast. The more we rely on experimentation the more important it is that it is done right. This is where the documentation becomes important. As I said above, a repetition of an experiment increases the probability of the hypothesis for which it was designed. Now suppose we want to repeat the experiment, either to check whether it really was performed as claimed or just to corroborate its findings. Then in both cases it is important to know exactly how it was conducted. For if there was no documentation we cannot tell whether we are running a repetition of the previous experiment. Or if we just want to see whether or not it was performed properly we need to see the documentation as well. A recent case might illustrate this. Some years ago, Kenneth Rogoff and Carmen Reinhart had published a study *Growth in a time of debt* where they argued that if the national debt exceeds 90 percent of the GDP (gross domestic product) then growth is severely compromised. The EU commissioner Olli Rehn based his policy recommendations on this study. Many people scrutinized it. Thankfully, not only the raw data but

[82] Probabilities and confidence are actually two different things. Probabilities are internal to the model itself. The model outputs values with given probabilities. Confidence has to do with the (nevertheless objective) likelihood that our estimated values for the parameters are correct. Typically, as we want to estimate some real number, we can determine that number only with some error. The error margin is inversely correlated with the confidence: the more error we allow the more confident we can be that the value is within the error margin.

also the excel sheet which they used for their calculations were publicly available. It was found that the calculations were incorrect. This could not have been discovered if the authors had not made all their information public. For all that could have been known was that either the data was wrong or the calculations—provided one could find some data on the matter.

Experiments should in principle be repeatable, and it is customary to try and reproduce important experimental findings. But this can fail for various reasons. One reason is that experiments may be expensive. This is the case with drugs. A clinical trial is a very expensive, typically running in the tens of millions of dollars.[83] This is one of the reasons that governments are not performing them themselves, they let the companies perform them and instead focus on setting up the rules for conducting them. It is clear that if one does not do the experiment then one simply has to believe that it was conducted properly. This however is not always guaranteed.[84] In other sciences fraud does also happen from time to time. This is why proper documentation is essential. If a new run of the experiment does not result in similar findings this is indication that the previous experiment has not been conducted properly.[85] In this connection it is interesting to recall the situation of the early miner's schools in Prussia. The authorities were highly interested in finding mineral deposits for exploitation. Initially, they employed dowsers for that purpose. However, they were keen to eliminate them not because they were dissatisfied with their results but because the method could not be recorded and done independently. Thus, in order to gain control over this process the authorities mandated trying out other means that did not involve dowsing.

Some experiments cannot really be done, however. I give two examples. The first is from economy. If we want to find out how an economy reacts to certain policies we can surely try to find out by implementing them. But that is too high a price to pay if it fails. And the experiment certainly cannot be repeated or controlled. There is no reasonable way to control the conditions under which the experiment can be run. Another case is astronomy. Stars are so far away and so big that there is no way we can control their behavior. All we can do is observe them. In both cases we are not left without tools, though. One is to create a *model* and then run the experiment on the model. This model can be physical. Ship researchers perform experiments mostly not

[83] See https://aspe.hhs.gov/report/examination-clinical-trial-costs-and-barriers-drug-development for a breakdown of these costs.

[84] This is why the Food and Drug Association (FDA) has recently announced stricter rules for data compliance, see http://www.pharmtech.com/fda-and-manufacturers-intensify-concern-about-data-integrity.

[85] Equally well could it indicate a problem with the second experiment. In practice it is assumed that the second experiment is conducted with the first in mind, and that it is being planned more carefully.

on big ships (because their experiments are done precisely to find out whether they will be successful) but on a much smaller model. Scaling the results from model to the actual ship is an art in itself but can be done.

Another kind of model is a computer model. Here there is no physical scale involved, so we can actually produce an electronic copy of the object in question. The modelling experiment consists in studying the behavior of the copy by replicating the conditions and doing forward calculations. This has become a very popular and successful method especially since computing power has become rather cheap. The basic idea is that the laws of nature (in the case of stars) or of economic development are not put into question, but rather are being used in doing the forward computation. This already points to a drawback. The experiment can fail not only because the data is sketchy or wrong, or because the computation itself is flawed but because the fundamental assumptions on the laws which govern the forward behavior are incorrect. And this means in turn that modelling cannot be used to prove or disprove a theory. For too much hinges on the initial parameters and the basic mechanism that we choose. For example, to show that the current physical theory is flawed by modelling the life of a star or a galaxy is a rather pointless exercise. There are many basic initial parameters that need to be estimated and there typically is no independent method to get them. Similarly, modelling an entire economy is no trivial affair. We cannot model millions of people buying and selling goods and services, there simply is no data available on that scale. We only have estimates. And then it addition to all that the model could be wrong in its fundamental assumptions. A rather prominent case is the climate model of the IPCC. To get at the data thousands of stations distributed on earth, including satellites, are being recruited. Additionally, historical records are being investigated. All this is fed into a large computer simulation model that then calculates the development of the climate depending on a few economic and environmental parameters. It is hotly debated how reliable this in fact is. And the debate is not only between those that deny climate change and those that do not but even within those groups.

Another problem underneath these experiments is that they are performed using increasingly complicated machines: computers, telescopes, colliders, and so no. While the working of an optical telescope is intuitively at least easy to understand, most modern telescopes are not optical. Similarly, with microscopes. The theory behind the workings of these machines is itself pretty difficult. And the more complex the machines the more daunting the task becomes to make sure that they work properly. Again, this is solved by invoking a theory that explains the workings of the machines. This does not only *sound* circular, it actually is. When a new machine is invented, we first check it thoroughly to see if it does what we expect it to do. When we are

certain of its functionality we can (and do) use it as a measuring instrument to discover new things or control other machines. In this way the advances in technology can help in developing and testing our theories. But also the other way around: new theory allows for the development of new machines.[86]

Finally, an experiment is a willful disturbance of the natural course of events. Thus, in principle no experiment can really be repeated. But most changes are deemed to be completely irrelevant. It seems not to matter how often we let something drop, we can still observe the gravitational pull. However, in many experiments we really do not know what the influencing factors are and how much of the circumstances is relevant. Sometimes we even know that the experiment cannot be repeated. The anthropologist who visits a tribe in the Amazon hitherto completely isolated from the rest of the world will be the last to do so. Every other anthropologist will be one visiting a tribe that had some contact with Western civilization. Evidently, the contact with the anthropologist changed the life of the people in the Amazon forever. For he will arrive there not only with stories of his home but a backpack full of equipment and presumably some more people around him. The experiment if you will succeeds but the price is that the object of study disappears. I add here that there are seven thousand languages on earth, and most of them will be extinct by the end of the century. Languages like cultures are being destroyed at about the same rate as species. The records that the linguists and anthropologists collect will at that point simply be of historical value. The corpus they collect will be the last, no more questions will be answered. Already when the numbers are low the evidence is getting questionable; people are losing the grip on their own language.[87]

The scientists are faced with a problem. If an experiment cannot be repeated then its findings will have to be taken with care. But what about experiments that will cause unknown damage? There are such experiments. Nuclear scientists were not sure at the beginning how much destruction will be caused by the detonation of a nuclear bomb. For example, the first hydrogen bomb turned out to be twice as strong as anticipated.[88] Unfortunately, even

[86] See Peter Janich: *Handwerk und Mundwerk. Über das Herstellen von Wissen*, München, C. H. Beck, 2015, for a beautiful illustration of the codependency of craftsmanship and theory.

[87] See Nicholas D. Evans: *Dying Words. Endangered Languages and What They Have To Tell Us*, Oxford, Wiley-Blackwell, 2010. Alan Dench told me that several years back he had received a phone call from an aboriginal tribe in Australia whose language he had studied. They wanted him to perform a burial ceremony. Because the number of speakers had dropped significantly, they were not sure they knew the proper words to say.

[88] According to John Horgan: *The End of Science. Facing the Limits of Knowledge in the Twilight of the Scientific Age*, New York, Broadway Books, 1996, 78, Teller had made calculations to the effect that the explosion would create a fire in the atmosphere that would not stop. It was Hanse Bethe, the head of the theoretical department at the Los Alamos project, who put

the anticipated risks have not deterred planners from going ahead with the testings. There are more dangerous experiments (genetic doping is of the them) where it seems a matter of time that they will be tried out.

Finally, despite all advantages of controlled experiments, their artificial nature is also a drawback. There is a controversy for example how useful speech data is that has been collected in a laboratory. Invariably, the intonation is artificial. But even when we exclude experiments done on human subjects, their value is not always what it is made to be. When a ship has finally been built, its performance is not actually rigorously tested. Rather, the ship is being checked once for how it performs on a particular stretch and then it is sold. No data of actual performance is available and therefore it is hard to tell whether or not the single test is of any value. However, extensive testing consumes time and money, and the future shipowner is not interested in spending more than necessary on it. Likewise, drug companies are not actually interested in too much testing; just enough to convince the authorities that their drug is safe.

A laboratory experiment allows to control the circumstances. We can control for air pressure, temperature, and so on. But typically, there is an endless number of combinations of these and we need to make a selection of what to test. Furthermore, in reality the conditions may vary randomly. It is not clear whether the results under constant conditions can be extrapolated to variable conditions, and so on. Much of science is stuck with discriminating more and more factors and combinations thereof without a clear understanding of how the results can later by synthesized into a theory.[89] It should also be clear that the controllability of test conditions is an illusion. In theory we can control as many factors as we like in practice there are all kinds of effective constraints on doing so. Money is just one of them.

This means that the relevance of controlled testing is limited. The size of groups on which a test is performed is a crucial factor, and there are any number of health conditions that need to be considered when asking whether or not a drug is effective. In the end, the clinical trials will only give a very broad picture of the possible effects and side-effects. There is a development underway that considers using all available data including actual patient data (after the drug has been approved) to evaluate it. However, apart from the fact that it is difficult to anonymize the data, it will never be able to make perfect predictions.

---

these worries to rest with his own calculations. See https://en.wikipedia.org/wikiList_of_military_nuclear_accidents. I am speaking of the test on the Bikini Atoll on March 1, 1954. Take a look also at the 1970 Banberry disaster for another sort of unintended consequences.

[89] This is the criticism raised by Masanobu Fukuoka: *The Natural Way of Farming: Theory and Practice of Green Philosophy*, Madras, bookventure, 1993. He says that the deductive method is superior to the inductive method because the inductive method has difficulties transcending the study of individual cases. Fukuoka had worked in customs as a plant pathologist before deciding to return to the rice fields of his father. The book summarizes his criticism of science—not just agrarian science.

Thinking

The fourth source of knowledge is quite simply thought, or maybe, echoing Kant, pure reason. Thinking is distinct from the previous methods in that we do not use experience or external evidence to move forward. But how can we achieve new knowledge? There are basically two ways. One, to be discussed at length here, does not actually extend our knowledge, it only brings to the fore the many conclusions that can be drawn from it. It is to use rigorous methods to explicate what the necessary consequences are of our already existing knowledge. The other is a more imaginative method. It actually expands our knowledge. It uses the experience so far to formulate a claim about something that has not been experienced. All empirical science is of the latter kind. It does experiments and then conjectures some general law. The latter can be done only with some additional margin of error. And that is precisely because it claims things that do not follow by necessity from what we have seen so far.

The are more modes of thinking. Humans proceed by all sorts of ways to generate expectations about the world. We see smoke and conclude that there must be a fire. This is reasonable, but it cannot be said to generate knowledge. We just provide some relevant explanation as to why we are seeing smoke. Similarly, we like to extrapolate events into the future to predict what is ahead of us. We see a red light and expect it to change to green very soon. We have a talented child who is good at mathematics and so we expect him to become a mathematician. All this is reasonable, but the conclusions are not certain. We know that and mostly we would therefore not speak of knowledge.

Let us therefore discuss the first method, that which proceeds by rigorous methods and from knowledge only generates knowledge. Although at first it is deemed to be a trivial affair, we shall see that it raises problems of its own.

Let me give an example. When Newton first proposed his laws of motion, they were so simple and yet so powerful that sooner or later everybody came to accept them as true. So, they became knowledge, and are by now regularly taught in school. This seems to settle the question as to how bodies move in space. But this is an illusion. The question really is: how much do we really know now? Here is therefore a real problem that has been studied extensively, namely the motion of the planets, for example the motion of the moon. The question is: how fast does the moon move around the earth and at what distance? It turns out that the motion is not regular (it is not strictly elliptical, so even Kepler's laws are just a simplification) and it depends on the influence of the sun and the other planets, and even the form of the earth (which is not a regular ball). In the 19th century Charles-Eugène Delaunay published a monumental work with the title "La Théorie de mouvement de la lune" (The theory of the motion of the moon). From a notice of the Royal Astronomical Society (Vol. 33, 1873) on the occasion of his death we learn that Delaunay starting 1846 devoted all

his time to this project. He developed an equation, which contains 42 terms, and then proceeded to the calculations, which he completed all by himself, twice, with sufficient time in between, so that no calculation was done with a memory of the previous one still present. The second volume appeared in 1867, a third one was planned but Delaunay died before he could finish that work. Newton himself showed how his theory could explain certain anomalies of the lunar cycle thus producing evidence that supported it, but he offered no complete calculation for the lunar motion. We can see why that was so. It took enormous effort. Indeed, the theory of planetary motion gave rise to the so-called *n*-body problem: given *n* bodies at certain locations and with particular initial velocities (the technical term is momentum), is there a complete way of expressing its development in time, say in the form of an analytic function? If so, direct methods for calculations were known. Sadly, while for 2 bodies this is the case, Heinrich Bruns and Henri Poincaré showed that no such expression exists if *n* is at least 3. The problem is not that we do not know the basic laws of motion, the problem is to say how the system looks like at a given moment in time in terms of where the particles are and what momentum they have.

The area of celestial mechanics is highly instructive in this respect. For notice that all astronomers had until the 19th century was a sky full of dots, perhaps with the exception of the planets that showed some structure. How could one know the position of these dots several years from now? Or, alternatively, how can we ascertain that a particular dot we see in the sky is the same star (or planet) that we saw elsewhere some hours or years ago? Recall the notion of synthesis. Kant says that our knowledge is objectual if we can consistently put the images together to make them aspects of the same thing. Carl-Friedrich Gauß famously calculated in 1801 where we should expect Ceres to be in the sky after astronomers had discovered it and subsequently lost track. This was no trivial affair. But it was a very important contribution because only by performing the calculations one could be certain which of the dots in the sky where known objects and which ones were new.[90]

This highlights another aspect. While mathematicians have done many complicated calculations, it was one thing to establish a complex formula and another to calculate the numbers that it produced. Of course, mathematical progress consisted mainly in doing both, but calculations were tiring (Gauß needed 100 hours for doing his calculations on Ceres). It can thus hardly be overestimated how much more progress the computer meant. It took away the tiring effort of manual calculations and made them almost effortless. However, it is also true that the theories that we are dealing with

[90] At that time astronomers did not even know all the planets we know today let alone their moons. Uranus had been discovered in 1781. Neptune was discovered in 1846 because of perturbations in the movement of Uranus.

are becoming much more complicated to the extent that it takes not only lengthy calculations but also highly abstract mathematics to even understand the basic consequences. Occasionally it has become possible to perform experiments to validate these theories. General relativity theory for example predicts that the space becomes "curved" in the presence of a body so that stars could effectively force the light rays to go around them ("gravitational lens"). This was indeed observed in a famous experiment in 1919. Another effect, the gravitational time dilation, was measured using atomic clocks in 2010. In the latter instance, the equipment to measure the effect itself is already very sophisticated. The quest for the Higgs-Boson has kept thousands of employees of CERN in Switzerland busy for years. In the case of string theory, however, the mathematical methods to understand its theoretical consequences are already so complex that it is not easy to understand how it could effectively be tested.[91]

Of equally fundamental importance is logic. The topic of logic is the study of inference. Mostly inference is to step from something that is true to something that is also true by necessity. Historically, logic was done in the Aristotelian way, called syllogistic logic.[92] One of the main features was that it was performed on sentences of natural language (mainly Latin). This slowly changed, partly because many were trying to mechanize not only calculations but reasoning. Leibniz was one of them. He was actively trying to find out laws of logic. His Leibniz 2000 was a forerunner of Boolean logic, named after George Boole.[93] The most radical shift was initiated by Gottlob Frege with his "Begriffsschrift", published in 1879. This led to the development of predicate logic, and in the beginning of the 20th century logicians tried to embed mathematics into logic by means of formalization. The most influential work was "Principia mathematicia" by Bertrand Russel and Alfred N. Whitehead from 1910. Two questions arose in this connection. One was whether or not the system of Principia Mathematica was consistent (that is, that it does not allow to derive a statement and its negation) and that it was complete (in that everything that was true was also derivable formally). It was Kurt Gödel who showed that there are statements that were true but could not formally be derived; moreover, he showed that no consistent system extending it that was useful for formal derivations[94] could show its own consistency.

[91] Voit: *Not Even Wrong*. See Edward Frenkel: *Love and Math. The Heart of Hidden Reality*, New York, Basic Books, 2014, for a personal account of a mathematician.

[92] See Terence Parsons: *Articulating Medieval Logic*, Oxford, Oxford University Press, 2014, for a modern account of the syllogistic.

[93] Gottfried Wilhelm Leibniz: *Die Grundlagen des logischen Kalküls* (Lateinisch – Deutsch), *The Foundations of the Logical Calculus* (Latin – German), Philosophische Bibliothek 525, Hamburg, Meiner Verlag, 2000.

[94] Formally, system that has a recursively enumerable set of axiom.

This has been the end of the so-called Hilbert program, the dream that mathematics would be able to formally derive everything that is true in virtue of the axioms. Clearly, this demonstrated not only the limits of formalization. It anticipated also the limits of computability. To show that in effect computers were limited in this way, Gödel and many logicians introduced many concepts which are still being used in computer science. Effectively, when computers actually appeared in 1940 the theory of computability was already quite developed and provided insights into what could be computed and how. The concept of a recursive function was developed, and Alonzo Church formulated the thesis that any function that is at all mechanically computable must be recursive. (The converse claim, that every recursive function is computable, is easy to see.) At the same time, Alan Turing produced a machine model named after him (the Turing machine) which was at the same time extremely simple and yet so powerful that it can compute all recursive functions. It is important to stress that all this was known before the computers were being produced. The first functioning computer was the Z3 built by Konrad Zuse. So, the theory of computers and computability existed, and it has been established by mathematicians and logician through pure thought.[95] Similarly, Einstein did not have to invent the mathematics he used for the relativity theory. The idea of a non-Euclidean space had been studied long ago, most prominently by Bernhard Riemann, who credited Kant for introducing him to the idea that there could be such a thing in the first place. Kant actually did not say that space is anything but Euclidean but he claimed that spatial cognition is a priori: we are born to conceive of it in this way. We don't find out in life. That surfaces could be nonplanar is however part of everyday experience. Nevertheless, it is quite a stretch to imagine our space to be curved in such a way that it is both finite (in volume) and infinite in the sense that it has no borders.

Mathematics and logics are therefore instrumental in understanding reality in relation to theory. Of course, a theory has to be true, but a theory that explicitly states what is anyway necessary is useless. It fails to be empirical. Karl Popper is credited with the principle of falsifiability. A theory is empirical only if it can be said what must be the case for it to be false so that it can effectively be put to test. So, mathematics and logic are methods to separate the empirical content of a theory from the nonempirical content. For example, mathematicians found through centuries of hard work that the fifth axiom of

[95] It is interesting to understand this properly. Beniger: *Control Revolution*, points out that we need to understand not only how computers came about but also why. For the latter he claims that there was a genuine demand for the task that they were supposed to perform, for example control of material and information flows. On the other hand once people understood what it is they wanted to have they first set their minds to it to see how such a thing could be built and how it could achieve these tasks.

Euclid does not follow from the others, so it can't be dropped from the list. On the other hand, it has become not only apparent that theories cannot be complete as Gödel has shown but also that even if they were we would not always be able to say given some proposition whether it was derivable or not. Predicate logic is a case in point (this is another of Gödel's famous results). It means that there are problems (many of pretty basic nature) for which no algorithm exists that can even in principle compute a solution. In computer science one simply has to accept that programming languages are strong enough to encode such problems. So, if a program computes answers to an unsolvable problem then in some cases it must simply run forever (if it calculates correctly). What is worse, we do not even have a method to decide which of the problems are decidable and which ones aren't. Indeed, it can be *proved* that no such decision method exists.

Even if a problem is decidable, algorithms may take a long time. Curiously enough, the development of computers has also given rise to several new disciplines concerned with practical matters of computability, notably complexity theory and numerical analysis. Complexity asks among other things how long it takes to perform a computation. Here is a very simple problem. We are given a million entries in a phone book and have to sort them according to name and/or phone number. Done in a naive way this would require 500 billion comparisons. First, we run through the data set and find the first entry, remove it, and then repeat with the remainder. Astonishingly, one can show that 20 million comparisons are enough if a more clever method is applied. The bound is in general $n \log_2 n$, and since $\log_2$ 1,000,000 is slightly less than 20, this establishes the claim. This is 25,000 times faster. It has in fact been estimated that the progress in computation since the 1970ies is an increase of 1 million. That means, that algorithms run 1 million times faster than forty five years ago. This is because computing power has increased by a factor of one thousand, and algorithms have become a thousand times more efficient. However, given that the amount of information has also increased rapidly, still most of what we would like to perform on the available data will be too expensive. (Algorithms that crawl the entire internet must be next to linear in time to be feasible.)

Next consider numerical analysis. This discipline deals with the problem of computing a problem in the real numbers with a computer. The computer, being a finite device, can only store and compute fractions m/n with a given range of integers m, n. Hence, every computation involving a dynamic system will use approximations. There are many problems with this. One is that every operation proceeds with a certain error (rounding error), and these errors accumulate. How can we be certain that the error is not larger than is acceptable? Second, there are problems that are "badly conditioned". This means that small errors in the input lead to large or extreme variations in

the output.[96] Problems that are badly conditioned do not necessary get easier to handle when computation is exact since we also have to take into account errors in the initial data (due e.g. to error in measurement).

Last but not least, mathematics and logic provide the foundations for empirical sciences through probability theory. Essentially, probability theory is a branch of mathematics and hence needs nothing but rigorous logic. It intends to provide an answer as to why the scientific method is at all successful. Rudolf Carnap was one of the first to try to prove the success using his inductive logic.[97] Inductive logic uses probability theory to confirm a hypothesis. It cannot do so absolutely. In effect, as Karl Popper has stressed, a law being universal cannot strictly speaking be proved, it can only be disproved. In principle one counterexample would be enough. But this would mean that no empirical science is possible. Hence, inductive logic was meant to assure us that even though we are never sure of our hypotheses we can nevertheless reach any level of confidence in them. The problem with that is that we do not know actually at what point we can be confident enough, since that depends on our initial assumption (the prior probabilities). Assigning initial probabilities is as much a problem as quantifying how much progress we have made through an experiment, no matter how good the numbers seem to be. One part of the problem is that we often do not know on what variables the outcome depend and whether we have included them all. Philosophy of science and analytic philosophy have made important contributions to the problem of knowledge and knowability and so would need to be discussed at greater length. However, our focus is more on the practical aspect of knowability, that is, on knowability under real life conditions. This changes the rules of the game rather dramatically.

[96] This is known as the "butterfly effect", which was talked about extensively in chaos theory, known also somewhat less drastically under the name theory of nonlinear differential equations.

[97] See https://plato.stanford.edu/entries/logic-inductive/ for an exposition and assessment.

CHAPTER 4

LANGUAGE

Language-Dependent Knowledge

The definition of knowledge contained one element whose role I have so far not elucidated: language. Philosophers like to define knowledge as an attitude towards a proposition. A proposition is thought of as an abstract content: it specifies ways the world needs to be like in order for the proposition to be true. But what actually is a proposition and, more importantly, how do we develop an attitude towards it? How do we know of other people what attitudes they have towards some proposition?

At first this seems to be a question of detail. However, when we attempt an answer we shall see that the way it has been defined actually always requires a language in which it is couched. This language boundedness is actually a virtue because it allows to acquire knowledge of things that are not present or events that happened long ago. But it also is a weakness because language acts as a filter.

I'll start by contrasting two kinds of processes by which we can acquire knowledge. One is direct experience and the other is instruction. Growing up as a child involves among other things classifying animals and objects into various groups and learning what they have in common and what distinguishes them. A child may thus learn from direct experience that dogs like to chase cats. It starts with observing that the dog next door likes to chase a particular cat and then extends to all kinds of other observations, namely that there is a group of animals quite like the dog of the neighbor in appearance, and another group of animals quite like that cat. And that animals of the first group like to chase animals of the second group. Not long after the child will have mastered the concept of a dog, the concept of a cat, and will have understood that dogs do not like cats and therefore keep chasing them. And it will probably know that there are some exceptions to this rule. All this can be learned quite quickly. In fact, it is known that mid range concepts such as "cat", "dog", "apple", "chair" are easy while most others take longer. Typically, when someone points at a dog and tells the child that this is called "dog", very few repetitions are needed until the child knows the kind of animals called "dog" in English as opposed to "doberman", or "cat". And this may be

explained by the fact that the concept of dogness has already been formed by the child itself before it had given a name to it. Although in practice concept formation goes hand in hand with learning words, it is not out of the ordinary to assume that some concepts are there before we have a word for it. And likewise, there will be words for which we have no corresponding concepts.[98] The concepts they have formed without words lend themselves to generalizations. They will expect then that when they see a dog and a cat, that that dog will start chasing that cat.

When children start to talk, we can also ask them what they think. We can ask them whether or not dogs like to chase cats and will get an answer. And that answer was probably there even before they were able to put it into language. Now contrast this with the question "Is Pluto a planet?" It is not realistic to assume that children know the answer to that question. For that they would have to know what the word "Pluto" refers to; and they would also have to know what it is to be a planet as opposed to a star, a moon or a comet. It would be difficult to show them directly what the difference is. All we can do is show them the sky and point at some dots. How are they to determine which of these dots is a planet and why? And how could they have found out themselves? The example itself is a curious one since astronomers have decided in 2006 that Pluto is not to be called a planet. So, they have changed the meaning of the word "planet". Before 2006 astronomers should have answered the question affirmatively, now they should answer it negatively. If you haven't heard of this fact, though, you would still consider Pluto a planet if you went to school before 2006.

This kind of fluidity is not only a problem of public languages. It also concerns the language of thought if you will. For if as I explained above our own concepts are subject to experience then we have to expect the same flexibility when it comes to our own concepts. When we see that the neighbor's dog chases our cat, then that is an instance of what is true of these two animals, say Fido and Garfield. The categorization of Fido as a dog and Garfield as a cat may make us believe that it is not a matter between Fido and Garfield alone but between dogs in general and cats in general. But what happens to this conclusion once we find out that we erroneously thought Fido was a dog and that in reality it is a cat? What may happen is that we continue to believe that dogs chase cats even though our concept of dog has shifted (and a supporting instance of our beliefs has been removed). This sort of situation arises frequently. It may sometimes be hard to see what kind of proposition we are talking about because invariable we have to point at it using some sentence of a language, but that doesn't mean this is not a real phenomenon. For we may think that part

[98] Like "elm", as Hilary Putnam: The Meaning of "Meaning", in *Mind, Language and Reality*, Cambridge, Cambridge University Press, 1975, 215–271, once pointed out.

of the competence of a speaker of English is to be able to translate an English sentence into a proposition; and that proposition will have to be expressed using some form of language of thought. This view is somewhat naive in my opinion.[99] My main argument is that it may not be possible even in principle to translate every sentence of natural language into the language of thought. But the conclusion is not that none of them is translatable or translated. The conclusion is that the universe of concepts that we are dealing with is a mixture of externally controlled concepts (for example, the meaning of the word "planet") and those that are our own (our concept of "planet", which still includes Pluto, or our concept of dogness, for example). The advantage of public concepts is that they can be shared, in the form of marks on paper, pixels on a screen or sound waves, but the disadvantage is that we are not in control of their exact meaning. Sentences can change truth value by virtue of the fact that they change their meaning. If my own concept of planet includes Pluto, I need to translate the same sentence of English, say "There are nine planets.", differently into my own language. I might have to translate it as "There are eight planets, not counting Pluto."

Every language is thus subject to constant change. And since no one really is in control of that change and it can happen anywhere at any time, it invariably follows that the meanings that we associate with given words can change between speakers. My neighbor and I can disagree what kind of action is considered illegal, what it means to be poor, and what counts as a good teacher. And we could disagree not only because we know different things about the world but even when we know the same relevant facts. The latter kind of disagreement arises when we disagree whether or not to be poor is to earn, say, less than 1,000 Euros per month. For if we disagree on that then we really are subsuming different things under the word "poor".

Signs

If we conclude from something that we see that there is something that we do not see then we are using the visible as a sign for the invisible. We are looking in the sky and we see smoke; we conclude that there must be a fire. Why is this at all a reasonable thing to do? Obviously, it is reasonable because there is a relationship between fire and smoke. This relationship can be looked at in two ways. One is objective: smoke is caused by fire. Therefore, when we conclude that there is fire, we are making what is known as an abductive judgement. We proceed from what we see to providing a reason for it to be there. The second interpretation is a subjective one: for us, fire is a sign of

[99] Marcus Kracht: Gnosis, *Journal of Philosophical Logic* 39 (2010), 397–420.

smoke. That is in effect why we are capable of performing the argument. The objective relationship does not itself cause any kind of reasoning. It is only in an indirect way that this comes to pass. The objective fact that smoke causes fire becomes a signifying relationship the moment we become aware of its existence. This signifying relationship is embodied in a sign. We say that smoke is a *sign* of fire. We shall see that signifying relationships, or signs, do not need an objective basis in order to exist or even be useful.

For the present investigation the main question is in what form this sign exists. This will allow us to explain how it can at all be effective in reasoning. Our answer will be that it exists in some material configuration of the mind that makes the mind *think* there is fire when it sees smoke. Of course, we could say that there is no need for a sign inside the subject if the causal link is already out there in nature itself. But that would deny the subject any autonomous role in the process. For it is patently obvious that there are two independent processes here: the physical causation and the representation of this causation in the observer's mind. Obviously, if we do not know that fire causes smoke, we will not be able to conclude that there is fire when we see smoke. And that is independent of the question of whether there truly is fire. Knowing that there is, however, may cause us to do things that we would otherwise not do (and might therefore not happen) like calling the fire brigade or running for the exit. In other circumstances we might believe that there is a relationship when in reality there is none. The sign in its materialization in the subject is thus independent of the reality that it purports to mirror.

Now what could be that sign? In theory it could be many things, and I will argue that among many things it *could* also be the sentence "Whenever there is fire there is smoke." In other words, believing that a sentence speaks the truth about the world is tantamount to having that sign within us. So, the sentence is the sign whenever it is believed to be true. In this way language becomes a code by means of which we can generate signs. This is no small feat as we shall see.

Signs cover vastly more territory than physical laws. For one, physical laws are unidirectional. They are about causation. They are made to explain the way in which the world moves forward. They explain how it is that when there is fire there also will be smoke. The signs however go in all possible directions. They will go from fire to smoke as well as from smoke to fire. Of course, the nature of the connection is different; one is causation the other is explanation. However, in everyday reasoning the boundaries between the types of reasoning are very difficult to discern.

Seen in this way, logic and semiotics are actually trying to achieve the same goal. Both are describing processes of reasoning but they differ in the way they see their fundamental nature. Logic just has propositions to state truths about the world, and delegates the reasoning to an abstract system,

the calculus.[100] The proposition does not determine the way in which it can be used in reasoning. Semiotics, however, in principle allows everything to be a sign for something else. However, the sign is not symmetric. A sign has two parts, the exponent (e) and the meaning (m). While the sign is the pair (e, m) we often say that *e is a sign for m*. So, we might say, somewhat sketchily, that we have a sign (smoke, fire) to express that smoke means fire. What this does *not* express is that fire means (or causes) smoke. This is because it is *not* identical to the sign (fire, smoke). The roles of the two entities are clearly different; there is the subject (e) that points to or signifies the object (m), and not conversely.[101] This flexibility on the part of semiotics is an advantage. It allows to understand human reasoning as a whole rather than just a small part that is concerned with physical reality. For signs can also be established purely conventionally. In that case they exist without there being a physical law underlying it. But they can be effective on human behavior nonetheless.

The problem of semiosis is now as follows. Consider some member of an Indian tribe looking at footprints in the sand. He may conclude that some group of buffalos had just passed this spot two hours ago. This is because he concludes that the form of these imprints, their number, the sharpness of the contour etc. is such that exactly this must have happened. But how does he know he is right? There is no unique answer to that since the ways we can come to know something are quite different. Yet they all agree that there is something of a rule that connects events of certain groups of buffalos passing at some point with imprints of a certain kind and multitude in the sand.[102] This connection is hardwired into the brain of the tracker. He has trained himself or has been trained to read these things, and as a result his brain has been configured to execute this mental calculation. This configuration is the material trace of the sign in question. It is not the sign itself. The same is true for computers. The bits and bytes resident in the memory of a computer which code a programming language so that the computer can answer whether

[100] More accurately, there is not only propositional logic, there is also predicate logic and type theory. So, there are more kinds of expressions than formulae to express propositions. However, propositions remain the privileged kind of expressions in these calculi, just the same way as sentences are privileged in syntactic theory.

[101] I am talking about *subject* and *object* here as grammatical roles of the verb "to signify", but we may also take this as a metaphor.

[102] There is a glaring problem in the way semiotics expresses itself. On the one hand it says that some thing expresses a meaning. On the other hand it is clear that what we have here is not one particular type of footprint p_{17} denoting an event e_{33}, of 12 buffalos passing two hours ago. Rather, it is a function f from footprints to corresponding buffalo group passings, which among other things maps p_{17} to e_{33}. In computer science one works with schemata and their instances, but there seem to be many layers of abstraction each calling for a functional relation. For in the present case we can consider an event of 12 buffalos passing through that happened two hours ago as an event type, of which the one hypothesized to have happened is just an instance. But it, too, is merely an instance of buffalo groups of some size s passing two hours ago, or of buffalo groups passing some time ago.

or not "2+2=4" are just the material trace or implementations of the program. The program itself is abstract. Otherwise we could not say that two computers have implementations of the same program. The implementations could differ in many ways. It would be hard to describe what material configuration would have to be present in the computer for there to be an implementation of a particular program.

The material trace allows the sign to act causally in the world. Therefore, the stability of the causal impact of the sign is ensured foremost by the stability of its material trace. So, long as the trace remains the same, its causal action is the same. When it changes so does its effect on the world. The computer will continue to compute in the same way as long as the individual bits of the implementation remain unchanged. But the more of them get flipped the more chances are that the computer will compute something completely different. It is to some degree irrelevant whether we want to claim that the changed data no longer is an implementation of the program or whether we want to say that it still is an implementation though a corrupted one. The trace is embedded in a causal structure and as long as the computer calls that program even when it starts to perform faulty computations, it is there to act the material trace of that sign. In the same way as a user must assume that the value of the variable stored in his computer has not changed so must the Indian tribesman trust his own memory that it does not corrupt the implementation of the skill of reading buffalo tracks. But he can never know for sure whether that is the case.

The stability of the trace in time and how it secures the causal effectiveness of the sign have been discussed earlier in more detail (Chapter 2: Theory — Traces). It also makes possible an intersubjective cohesion. How it will do so will be discussed later. What it means is that two people are said to have the same sign just in case the trace of the sign acts in the same way. This does not mean that they have identical traces (they cannot be identical, and often they will not even be similar) and so it cannot mean that we can identify the traces, it only means that we talk about their visible effects. It is by and large impossible to predict from the traces what the actual behavior of a system will be. Consider a computer and the fact that it contains some billions of registers that are set to 0 or 1 in a particular way. How does that enable us to predict the behavior of the computer? Thus, while descriptively it is easy to say what the state of a computer is, it is not easy to say how it will behave. This is not only a problem of the magnitude; a memory of 1 GB introduces so many possibilities that their interaction cannot even be listed. It also introduces a problem of abstraction. What part of the state shall be considered the trace of the particular program and what part is irrelevant? Seen from below (the individual registers) the program is abstract. Seen from above (identifiable behavior) however it is the states that are abstract. It is

this latter problem that justifies high level talk. There is no sense in trying to ground the behavior of computers in the state of particular registers any more than it is meaningful to ground the behavior of humans in the particular connections between individual nerve cells. This is an impossible task.

But now consider again the question of having the same sign, say (smoke, fire). If two humans possess that sign, that is, if they behave in a way that lets us conclude that they *know* that smoke is caused by fire, then we may reasonably assume that they will continue to think so. And this is because the traces as part of their physical makeup are stable to a sufficient degree. Communication relies for its success on these two factors: the interpersonal coherence and the temporal stability of signs. It means that maintaining successful communication requires making sure that we meet both of these requirements.

Concretely, the English language contains the sign (dog, doghood), where /dog/ is the English word and "doghood" refers to the meaning of that word in whatever terms we want to specify that.[103] We usually say that a person A knows the meaning of "dog" if he has this sign. That in turn means that the sign has a trace t_A in A. This trace becomes a physical attribute of A. His brain is shaped in such a way as to have implemented that sign. To say that another person B has that sign similarly means that he has a trace t_B of it. While we cannot ascertain whether or not t_A and t_B are the same in relevant respects, we can identify whether or not A and B have the same sign. We simply ask them about various objects whether they are dogs according to them. In this way we guarantee that they subsume the same objects under the word. The traces are therefore established to be causally identical. However, notice also the following. Consider two children that do not yet talk English but seem to have managed the concept of dogness. We would then like to say that they have the same sign, because they do have a label (whatever that may be) that unites the traits of dogs under a concept. But in what sense do we really say that they have the same sign? The problem although the meanings are the same, the terms for them are not. So, strictly speaking they do not have the same sign. However, from an external, causal perspective no difference is observed. Thus, there is a subtlety that needs to be kept in mind, which is that to have the same sign requires *subsuming* the two traces under the same abstract sign, which can be done only if they are causally identical. For language purposes, another condition is identity in exponent.

How then can the interpersonal coherence be produced? How can we guarantee that two people will subsume two causally identical traces under the same sign? There have been attempts to produce non-materialist

[103] Some take this to be the set of all dogs (extension), or maybe a function that takes a possible world to the set of all dogs in it (intension), others still think of the meaning as the essence (or "idea") of a dog. I am not going into details here since they will only obscure the point.

criteria for identity. Most famously, Plato identified meanings with "ideas". To associate the correct meaning with the word "triangle" means to grasp the idea of a triangle and make that the meaning of the word. Other kinds of explanations (say, in terms of morphic fields) suffer from a simple problem. They suggest that the way signs are implemented is the way computer programs get installed. There is a common source from which the program (and the installer software) gets downloaded and then installed. The coherence is guaranteed through designing a proper installer software and the fact that there is a unique source from which the sign is being drawn. It would be similar to the way Aristotle describes the process of understanding: the mind become one with the logos of the object.[104] The copies are thus obtained from the original matrix, over and over again. Discrepancies may arise (copy errors and failure of correct installation) but they are of first degree only. They can be corrected by reinstallation. This is however not how language works. Words are being passed on from generation to generation without intervention from above, without there being a supreme corpus from which to draw them. This in turn lets us explain how it is that Dutch "slim" (clever) and German "schlimm" (bad) are related. They have a common origin. While the changes are small in every step, the cumulative effect can be large.

Language change is fascinating in this respect. It shows us that signs are slowly changing and that it is both components, the exponent as well as the meaning, that change independently from each other. None of the changes is lawless, and mostly they proceed quite slowly, so that there is every reason to consider the sign (*e'*, *m'*) the same as (*e*, *m*) even though not everything is the same. It may be that the pronunciation of words changes (giving rise to an *e'* different from *e*), and it may be that the associated meaning changes (giving rise to an *m'* different from *m*). None of these changes is completely random, though many of them appear to be arbitrary (and are in many cases arbitrary). Institutions as well as powerful people like monarchs have tried to stop that change from happening or at least control it; but they have generally failed in their attempts. Language, it seems, has its own will. It proceeds slowly, but with a certain determination that it the end leads to languages breaking apart into many languages whose speakers do not understand each other and more, like French and Italian speakers that speak a descendant of Latin, or Swedes, Dutch and Germans, who speak a descendant of Germanic.

The attempts to stop the change of language is doomed for many reasons. One is very simply that nobody actually knows the language in its entirety. Linguists have fought for decades to obtain a full description of certain languages (English, French and others), and typically they managed with

[104] See Raoul Mortley: *From Word to Silence I. The Rise and Fall of Logos,* Bonn, Peter Hanstein, 1986. The doctrine of Aristotle was in fact more subtle and avoids some of the problems of the Platonic ideas.

respect to the syntax and morphology. But what certain words mean is hard to come to terms with. Try to determine the precise meaning of "liberal", "faithful", "formal" and so on. Typically, this requires language, and will inevitably be circular if we try to define *all* meanings of the words. If we cannot write down the meanings exactly there is no way to prohibit change as it would not be detectable. Given that languages are not uniform across speakers—there are regional variants, class based variants and so on—the language an individual speaks is a compromise between many requirements they have to meet simultaneously (being clear, being polite, speaking correctly but not taking too much time) and this puts pressure on the language as a system. And, not the least, language is also aesthetic, speaking gives us pleasure. In sum, language wants to be free.[105]

Looking at language as it exists in our own heads we also recognize its nature as an *archive*. The brain is no different from a clerk who can decide matters only by means of the papers that are being placed in his desk. The brain gets data by the sense organs. They provide a representation of the external world. The brain then acts on the basis of this input as well as the myriads of past experiences. These past experiences have been encapsulated in the form of signs. This is why we speak of *internal semiosis*. However, it is not necessary to have had experiences in order to have a sign. The Indian tribesman does not have to have experienced herds of buffalos passing by and subsequently looked at the traces they leave in order to know how to read tracks. It suffices to learn that form someone else.

This is a huge advantage, which is primarily due to the fact that we possess language to produce the effect that direct experience would otherwise have to play. However, it has dangers of its own. In the same way as language changes over time so will the factual capabilities of the tribesmen change. Learning is never perfect. Some of what we have to explain in order to transmit the knowledge exactly will either be left out or learned imperfectly for various reasons. It is therefore of importance that skills are subjected to some reality check. If the engineer does not learn how to build bridges it may happen that they collapse. If the farmer does not know where and when certain plants grow he will not get enough harvest. And so on. It is in this connection somewhat disturbing to note that on the one hand we need exact knowledge and on the other we do not want the outcome of the reality test to be negative. We do not want bridges to collapse, the farmer does not want his harvest to decline. And we do not want people to get killed just so that forensic medicine can advance its knowledge.

[105] See §12.4 and §13.8 in Jürgen Trabant: *Weltansichten. Wilhelm von Humboldts Sprachprojekt*, München, C. H. Beck, 2012.

This in part explains the conundrum that we face. Humans generally learn from negative experience, it is said. What this above discussion suggests is that this is by and large the only way they can learn in these types of problems. If something fails we know that we need to act, while when it does not, there is no reason to act. The best procedure would be that of the autopilot: constantly move a little left until too far left and correct to move rightward, until too far right etc. However, that would incur the risk of calculated imprecision.

We can now understand why language needs to be transmitted from generation to generation. Given that the ability to speak a language is not given at birth, only the ability to *learn* to speak it, every infant has to go through the phase of learning to speak. All this requires effort. Learning to speak a language takes them several years depending on the level of mastery. When speaking to others, though, we need to assume that they too speak the same language as we do. Otherwise communication makes no sense to begin with. But once we do, how can we assume language to work?

There are various models of that. I will mention one process because it makes us understand some details of the general process of semiosis. When a word is heard, say "sister", the inner processor starts to call out the first letter as soon as it is heard.[106] It is /s/. At this point all words starting with /s/ wake up and register. The processor then calls out the next letter, /i/, and from the registered words those drop out that do not start with /si/. This process continues until one candidate is left, and that is when the processor determines that the word has been recognized. At this point it can go on and determine what this means, for example.

It is important to realize that the number of signs is so vast that it makes no sense to keep them awake equally. Rather, there are procedures that make certain signs more readily available than others. In the case of word recognition, it is the phonetic shape. In other circumstances matters are less clear. Take as an example the theorem of Pythagoras. All children learn in school that in a triangle with one right angle, the sum of the squares of the cathedes is equal to the sum of the hypothenuse. Generally, they would know how to prove that, but they would be able to remember it if only in the somewhat cryptic fomula $a^2 + b^2 = c^2$.[107] So what makes this theorem at all useful for us? Just knowing it definitely is insufficient. We must be able in the relevant circumstances to recall what it says. What exactly is to be understood under "relevant circumstances" needs to be determined. It certainly is the case that a mathematician somehow learns to call this theorem at the right moment, while most people have no idea when to do

[106] Actually, it is not letters, but sounds that we talk about here, but I shall ignore that detail.
[107] Cryptic because they generally do not remember what the letters stand for.

that. In other words, it is not the sign itself that is different, but the semiotic processor that makes the decision when to use a sign. This is a subject that unfortunately rarely gets covered in semiotics.

What the discussion suggests is that signs are not implemented in the human brain by means of building a trace from an original sign. Rather, it is a process of creating a trace from the causal effects of some other trace. In computer science jargon this is called backward engineering. Free software clones of commercial software are made by recreating the behavior of the software. They are not the same program and the algorithms may be wildly different. And so, they are not obtained by buying and implementing the original software, because that is generally made impossible by encryption strategies. They are instead obtained by judiciously checking how the program reacts on user input and then creating another software with that exact functionality.

Humans are exactly in that situation. They do not possess the master plan. They do not know what ideas are associated with a given word and then grasp these ideas. They do not intuit something by reading out a morphic field. And they do not in general remember the sign. What they do is monitor as closely as possible other human beings and then build copies of their signs in their own heads. This creates a problem for us. Suppose that person A possesses a trace t_A of a sign S and that person B produces a copy trace t_B. Is this an instance of the same sign? And how do we determine that?

Evidently, in order to know that we need to control the process of copying itself and institute certain reality checks. This is basically the mechanism by which children get taught. They will be taught certain words or skills and then the teacher checks whether they have learned them correctly. This in itself is problematic, too. For the teacher may with some justification be trusted on that matter. But many matters do not come with an officially sanctioned arbiter. Or consider the problem of having the children do homework and then determining whether they did it right. While it is easy to determine even for a layperson whether someone has learned to play piano, it is much harder to determine whether someone has learned mathematics. If you listen to someone playing piano, you can roughly see how good they are no matter how well you have learned it yourself and how well you understand music. But to understand whether someone is good at calculating numbers or proving theorems you most likely need an expert. And then we have the formidable problem of determining who will be chosen to be an expert.[108]

[108] This is Plato's problem. To rule his republic he wants only the best men, but it is far from clear how they can effectively be chosen.

This is precisely the situation in language. We shall return to that discussion in Chapter 4: Language — Context and Interpretation. Suffice it to say here that the idea of subjecting signs to empirical tests is not always workable. Consider the name "Aristotle". What does it take to know who Aristotle was? Is it enough to know how he looked like? And what is it that makes our talk of Aristotle connect to the (supposedly) real Aristotle who lived more than two thousand years ago in Athens? Could there not have been two Aristotles? Or none at all? Recall that after the second world war there appeared a series of mathematical textbooks authored by a certain Nicholas Bourbaki. It so turned out that the authors were a French collective (plus some people from Chicago) who chose this name as a pseudonym. However, and that is the twist, there really *existed* a person with that name, and they expressly used his name for their purposes.[109] So, when you check that name you may find that he was a general in the German-French war from 1870/71. And you would erroneously think that this man was a talented mathematician.[110]

This phenomenon is not restricted to names. All kinds of words often just create illusions of understanding. Consider the slogan "compassionate capitalism" used by George W. Bush in his electoral campaign. What exactly is meant by that? Can we tell? Or take the words "sustainable" or "organic". When you hear them, you are not likely to consult a dictionary to see what exactly they mean. You register a meaning together with some emotion (positive if you sympathize with the green movement, for example). You understand that it has to do with a technology that we can in principle use indefinitely, or some way of growing food that does not use chemicals. But can you get closer? Can you determine what "indefinitely" means, what kind of chemicals are good and which ones not? Most likely the attempts at zooming in will stop very soon. Seen this way it is a miracle that we understand each other at all.

I do not want to go more into detail of the nature and relevance of signs in our culture. It is not necessary in what is to follow that we know exactly what a sign is. What is important to understand is the enormous effort a culture must make in order to secure its own continuity and that this project must always fail to some degree.[111] The problem of self reproduction of signs and their significance is a deep one. In order to understand what counts as a copy of a trace of the same sign we need to get at the essence of the sign. Two traces fall under the same sign regardless of the way they are coded if they preserve the essence of the sign (in terms of causal relationship). But in order for us

[109] See Amir D. Aczel: *The Artist and the Mathematician*, New York, Basic Books, 2007, for an entertaining portrait of this group.

[110] Saul Kripke has concocted many puzzles of these sort of twisted identities, such as the "Pierre Puzzle". It is not necessary to go into them in detail, though.

[111] See Bernard Stiegler: *Die Logik der Sorge. Verlust der Aufklärung durch Technik und Medien.* Frankfurt am Main: Suhrkamp, 2008, a topic to which I shall return in Chapter 6: Passing on Knowledge — Generations.

to determine that these two traces preserve some essence we need to have a concept of that essence.[112] This essence is itself implemented as a trace. This is evidently circular. We therefore must rely on a master code by means of which we can code all essences. This code is called language.

Codes

Unlike sign systems in general, which can arise in any context, codes have a special status. The signs of a code are arbitrarily created. The connection between some exponent and its meaning established by the code is not that of causation, similarity or independently existing cooccurrence. It is established through the creation of the code itself. The biggest such code is language. Before we turn to language, however, we shall take a look at other codes.

Knowledge that is not the effect of learning from experience has been coded. Typically, some natural or formal language has been used to transmit the knowledge. But codes and languages in particular are likewise subject of knowledge. There is no fact of the matter that determines that a certain kind of animal must be called "cat"—indeed, other languages use different words. Hence, the knowledge that some animals must be so called is itself something that we need to know. The code is therefore resident in our brains. Moreover, it is constituted that way. Searle in his discussion of social facts takes the example of money. A banknote is a particular piece of paper. But to be a particular piece of paper does not constitute being a banknote. Something else is necessary. In fact, banknotes are issued by (central) banks through an official process. They are individually registered (they all carry a serial number). And so there is a database of banknotes, and for a piece of paper to be money it has to be registered in that database. Earlier, there was no such system, so there was no reliable way to tell whether or not something had been issued by the authorities. But it can function like that if we believe that it is. And that in turn requires that this fact, that such and such pieces

[112] This consequence is nowhere found in the works by Peirce. Peirce like no one before (or even after him) studied semiotics as a discipline of human reasoning. At his time though people where not really interested in how language is represented in the brain. Even if there was no way to empirically verify hypotheses about the nature of the brain, there was no attempt to picture a way in which humans could at all institute the semiosis. Computer science is now closing the gap. With unimaginable speed it recreates all kinds of reasoning abilities (including translation) without necessarily knowing how they work. Brain sciences are still trying to determine the basic makeup of the human brain, while the mechanics are already being implemented in computers.

of paper are money, is represented in our mind.[113] The arbitrary nature of the monetary system is captured by the fact that any object can be turned into money as long as one reaches agreement that it should be so used.

Money like language is a way to mediate the effect over distances in space and time. To convey the fact that there is milk in the refrigerator I can of course open the refrigerator and show some container of milk. But this requires that me and my interlocutor are present at the location. Similarly, to convey that Caesar won over Vercingetorix over 2000 years ago we need neither be able to show Caesar nor Vercingetorix. We can pass on knowledge from generation to generation precisely because language enables us to dispense with the need of having the objects that we talk about present at the conversation.[114] This is partly because the relationship between signifier and signified is arbitrarily established. We can choose an appropriate object to be the signifier so that it can be shared with ease. Words and symbols are cases in point. This book is testimony to the extraordinary power of language to share ideas with minimal material effort.

Codes are relations between some object and a (mostly propositional) content. The objects can be letters, pictograms, road signs, bar codes, proofreader's marks, and so on. Once we start looking, we can find hundreds of codes that play a role in our lives. Their complexity varies enormously. They can be as simple as comprising a few signs or as complex as shorthand. Many others are not called codes but technically function in the same way. These are administrative tools like account numbers, license plate numbers, ISBN, IP-addresses and more. To be able to use these codes you need to know at least in part what things it relates to what symbols. How the codes are maintained is of lesser importance.

Here is one example. The ISBN (International Standard Book Number) is codified through some standard (ISO 2108). It determines how these numbers are composed. The initial standard consisted of 9 digits plus a so-called check digit that is a digit from 0 to 9 or the letter X. The code is maintained at some official site. It consist of an agency (the International ISBN Agency in London) and some subagencies, called group agencies, from which one can buy (!) these numbers and then assign them to a particular book. The assignment is recorded in a data base to which all parties have access. Recently, in 2007 it was agreed to increase the length of the number to 13 digits. Again, the first 12 are genuine numbers, the last is a check digit. This effectively creates a new code. However, since it would be a rather massive undertaking to start a new

[113] John R. Searle: *Mind, Language and Society: Philosophy in the Real World*, New York, Basic Books, 1997; Searle: *Social Reality*.

[114] Recall the people of the learned society of Laputa in Gulliver's Travels by Jonathan Swift who decided to carry around the objects in order not to have to waste energy through speaking.

code from scratch, the following system was introduced. Given a 10-digit ISBN, take the first 9 digits, prefix that with 978, compute the check digit and put that at the end.

There typically is not a single code to name things. We already explained that there are two different book numbers. Luckily, the length of the number reveals the code that has been used. In other circumstances this is not so easy, though. Clothes are measured in a variety of ways. Each country has its own "standard". This is why the numbers for these various systems are often preceded by country codes. Otherwise it would be hard to understand since clothes are sold internationally. This is somewhat similar to units. If we want to measure energy, for example, we can give the number in Joules, Watt-hours, barrels of oil equivalent, and more. The number itself is therefore completely meaningless, we need to have the unit as well.

Codes can be cascaded. To encode a certain content, you can first decide to use your verbal language skill, and then translate that using alphabetic writing into a sequence of letters. Or you may use shorthand to record a speech. You can assign a 10-digit book number to a book, then recode it using 13 digits, and then write it down in hexadecimal. There is no limit. Each of these steps may be reversible, but need not be. Typically, codes eliminate some information on the way.

One vital aspect of codes is that they must be known to the communicating parties. Suppose that A wants to send B a message, m, using a code C. A is called the *sender*, B the *receiver*. A must know how to encode messages via C. Minimally, to convey message m he must know what its code is. Let it be e. So, A needs to know that C contains the pair (e, m). B on the other hand needs to know that same fact in order to understand the message that A has sent. The code C is considered to be a set of signs. The sender *encodes* m by choosing an appropriate e so that (e, m) C. The receiver *decodes* e by finding an m' such that (e, m'). Communication is successful if $m = m'$. This is not guaranteed. (Consider ambiguous sentences.) The algorithms for encoding are distinct from the algorithms for decoding. There are codes for which knowledge of encoding does not guarantee an ability to decode.

In normal communicative settings it is both guaranteed that A and B know the common code C. Knowledge of the code should be shared and common. There are many instances where this breaks down, however. Children are assumed by parents to understand the language yet it often turns out that they are missing crucial parts of the meaning that is conveyed. When people speak at conferences they likewise assume that everybody speaks the official language(s) of the conference, minimally English. Again, this expectation can be disappointed.

Communication in a group works only if all members can rely on all others to know how to use the code. Language in this respect is a very complex and

difficult code. Learning it means learning the rules that govern it. This does not mean that we can tell what it is that we are doing when we speak, nor does it mean that we can explicate the rules by which we speak. All it says is that in practice we are able encode and decode.

Mastering a language is a prerequisite for mastering writing. Since most writing systems are alphabetic or at least based on a language one cannot learn a writing system before knowing the language on which it is based. But it is possible to be fluent in a language without knowing how to read or write. In fact, until very recently, this was the case for almost all people on earth. The knowledge or writing was restricted to the intellectual elite. There anyway was little that it could be applied to. Around the year 1800 most households had only one book at home (if at all), the bible.[115] Knowledge or reading and writing was necessary in addition only for dealing with the authorities.

Writing systems would merit a special discussion. However, for our purposes it is not so much their logical structure that is important. What matters here is the fact that the message becomes durable and shareable beyond space and time. We can use writing to preserve a message so that we can later recall it. This is useful in innumerable ways. Just one noteworthy fact is that with so many meetings we engage in professionally and privately we need to keep track of the things the we said or agreed on. We therefore have to make protocols. These protocols are then stored. Until recently paper was the preferred medium, but that is changing rapidly.

Language and Worldview

Language is a code; not any kind of code, but—as I said above— the master code. The one that is learnt with no help from other codes, and that subsequently is used to learn all kinds of other codes (including other languages). It has not been so obviously that language really is a code to begin with. The Stoics used to believe that the names they used for things were derived from Nature itself.[116] This claims that the objects carry their names in them so that they can be discovered from the object themselves without mediation from other people. This is not to be confused with the definition of truth by Tarski according to which "Snow is white." is true just in case snow is white. At first this sounds trivial, almost like saying that the reason we call a chair "chair" is because it *is* a chair. The triviality disappears as soon as we apply that to

[115] Protestants were notorious for insisting that everyone has a bible so they understand what is written. To the dismay of the authorities, for example in Catholic Austria under Maria Theresia, they went into the valleys to sell bibles to their fellow protestants, those that still held out against the ban against protestantism.

[116] See the discussion in Chapter V of Mortley: *From Word to Silence.*

another language. Tarski's definition now says, for example, "A hó fehér." is true (in Hungarian) exactly if snow is white. Similarly, to say that Hungarians call a chair "szék" because it is a chair is downright strange. There seems to be no connection whatsoever.

Ferdinand de Saussure is often credited with discovering that language is arbitrary. Although this is not quite true in that he was neither the first nor the only one in his time to claim this,[117] he nevertheless wrote quite a lot on the nature of linguistic signs. As one of the founders of linguistics he cared a lot about the nature of the linguistic code. What he stressed was not only that this implies that we need to have dedicated linguistic knowledge in that we need to know what the words of our language mean and how we get to the meaning of expressions, he also stressed that the language through accidents of various nature makes certain things related. In German for example you call a bench "bank", which is the same word you would use to refer to a bank. The words are called *homonyms*. What is the connection? Well, pure accident. While some homonyms appear motivated—like using the word "space-ship" for something that does not actually travel the seas even despite the fact that it does look like a ship—others do not. Yet others did make sense some time ago and then ceased to do so because the world around us has changed. An instance is the word "bookkeeping" for accounting, which was coined when merchants did actually use books to do the accounting. Nowadays they all face a screen and what they are keeping is data files or some such.

This view on language is still consensus. Chomsky does not disagree even though for decades he has championed the view that humans have a special faculty, called language faculty, whose initial state is genetically determined, which is geared towards learning and understanding language, especially the first language.[118] This is no contradiction. The genetic program allows to learn language but there is enormous leeway as to how languages can look like, in particular what words they use for what purpose. Even if one believed that there is a small set of basic characteristics of human languages[119] the fact that there are close to 6,000 languages that are sometimes totally different from each other speaks to the idea that there is no predetermined language to begin with. It does not matter for that purpose if we believe that language is God given or genetically programmed. The evidence is overwhelming. And this in turn means that humans need to spend a lot of effort learning a language. Not only their mother tongue, but any other language they wish

[117] Aristotle believed that language was essentially arbitrary. For similar views, see also John Locke: *An Essay Concerning Human Understanding*, 1690.

[118] Noam Chomsky: *Rules and Representations*, Columbia, Columbia University Press, 1980. See also Steven Pinker: *The Language Instinct: The New Science of Language and Mind*, London, Penguin, 1994, for a popular account of these views.

[119] As used to be the tenet of the so-called principles and parameters framework, which is no longer the central dogma in generative grammar.

or have to learn later in life. But then if languages are so different it seems like a miracle that we can learn to speak *any* language whatsoever. Imagine yourself as an adult trying to learn a foreign language without any instruction: that is approximately the task of newborn babies. This is a task that takes them several years if ever. This is why Chomsky proposed that the human mind is born with a certain predisposition, a program, that facilitates the process. And that as a result limits the possibilities for languages, since some imaginary language may not be learnable.[120] Chomsky's assumption, which has not gone unchallenged since, is that the input children get is too poor to nail down the adult language to a sufficient degree.[121]

Now, while Chomsky was concerned about syntax, the most evident counterexample to his thesis is actually semantics. When we look at the meaning of words, even simple ones like "car" or "house" we find that there is a lot of detail one needs to pay attention to. How do we distinguish different types of cars or houses? Are tractors or trucks also cars? How big is a house as opposed to a mansion or a skyscraper? What does it mean to borrow and when does borrowing and not returning actually become appropriating or even stealing? When we translate between languages we become acutely aware of the fact that single words often have a dozen or so translations, each with their own added nuance, words, that by themselves mean something different in the target language. The distribution of German "bringen" and "holen" is quite different from English "bring", "take" and "get". Nevertheless, all of them can at times be the translations of each other. This is awkward because it means that there is a lot of detail involved. And it does seem that even foreigners can pick up the differences.

It has been suggested that language also is a lens through which we perceive reality. A particular example is the so called *Sapir Whorf Thesis*.[122] It claims that language through its morphological system or vocabulary introduces a web of beliefs, biases or dispositions. As an example it was claimed that the Hopi do not distinguish future and past because their language only distinguishes present from nonpresent or that the Inuit would be able to distinguish many types of snow because they had so many words for it. This thesis has suffered many defeats over the years. For example, it was noted that Inuit languages did

[120] This may well be the case with artificial languages like Loglang. I have heard from two people trying to teach their children that language, but I presume that when they made that decision the children were already fluent in another language so that instruction would actually be much easier. Looking at the website https://mw.lojban.org/papri/loglang it appears that it already went the path of human languages by having many descendants.

[121] This is dubbed the "poverty of stimulus" argument. Additionally, for most of the early life children get no negative data on what sentence is *not* used. They only get it when they understand language and then can be told explicit rules.

[122] See Penny Lee: *The Whorf Theory Complex. A Critical Reconstruction,* Amsterdam, John Benjamin's, 1996, for a critical assessment.

not have more words for snow than English.[123] And I do not think that the Hopi did not know about temporal order. For there are many situations where it does matter (getting dressed, building a tent, and so on) and it is inconceivable that they did not cope. On the other hand, we do know that some effects of the sort claimed by the Sapir Whorf Thesis do exist. It is established that children acquire what is known as categorical perception of sounds. They learn to pay attention to some differences in sounds as opposed to others. And it is very difficult to undo this perception. For example, speakers of European languages generally have a hard time perceiving tones in Chinese. Another example is languages that require what is known as dead reckoning. In Guugu-Yimithirr, an Australian language, there are no words for left and right, you only have north, south, east and west. This makes it necessary to always record the absolute directions, otherwise one cannot tell a story properly. And it is observed that when they act out what happened to them, speakers will observe the orientation of things. If the crocodile attacked from the north, then when they are telling the story they will make sure they show the crocodile as coming from the north.[124] This is not the same in English, for example. So, one language may force us to record in some detail that another allows to omit. In German you need to record the gender of a person in many instances. There is no single word for "doctor" ("Arzt" for males, "Ärztin" for females) or "pupil" ("Schüler" for males, "Schülerin" for females).[125] The net effect is that from some texts we can infer the sex of a person. However, it neither means that English people do not remember the gender of a person (they generally do) nor does it mean that they cannot tell the difference.

There are a myriad of other categories of this kind. In Pima, a Uto-Aztecan language, you need to be specific about where an event happened with respect to you (in front, on the sides or in the back), in other languages you need to specify whether you saw something or whether it is just hearsay. All these things are simply enforced by the rules of the language and so also influence what gets communicated. But it is not clear how much this influences what we perceive. So, it is not the perception which is at issue but the nature of code itself. It may in some cases interfere with our judgements. In English, when you hear that a murderer escaped from prison you might conclude that this

[123] Geoffrey K. Pullum: *The Great Eskimo Vocabulary Hoax and Other Irreverent Essays on the Study of Language*, Chicago, The University of Chicago Press, 1985.

[124] See Stephen C. Levinson: *Space in Language and Cognition. Explorations in Cognitive Diversity*, Language, Culture and Cognition 5, Cambridge, Cambridge University Press, 2003. As Ed Keenan (p.c.) told me, while Malagassy speakers have the same problem, they do not in general learn dead reckoning. When visiting another village they would for example ask someone to tell them where north is so that they could express themselves correctly.

[125] There are efforts underway to change this. New forms get introduced to obliterate this difference like "Studierende" for students, which is gender neutral but only in the plural and "LehrerIn" which at least in writing is a merger of both forms. However, these are reformist attempts, they try to change the language and do not invalidate the point that is being made here.

a man, because that is what incarcerated murderers look like. And if police arrived at a crime scene where someone has been murdered, they would most likely look for a man as a suspect. German police officers would start off with the word "Mörder", but do they really firmly assume that the crime was committed by a man? And do they really believe that it was one man rather than two? Or a man and a woman? There certainly is a justified statistical bias. But there is also a simplification involved that may at times be hard to correct. We need to decide between man and woman, between singular and plural, and the choice we make sets us on track.

Problems get far worse when we turn to abstract words. Can we explain the meaning of "culture", "honor", "capitalism" or "god"? It seems that we cannot, or not to a sufficient degree. Books have been written about the nature of democracy, for example. It is not clear what electoral system counts as democratic. We consider the EU democratic but many people insist that the election of people of the inner circle is all but. Who is to say? The difference is not one of facts. The disagreement will persist even when the facts are laid out. But then what does it mean to have a debate between people whose languages are actually different if not in syntax then in meaning?

To answer this we have to recall the situation of infants. At first they are unable to understand a single word. Yet by listening to others talk and by being talked to they actually learn to speak until they are as fluent as anyone else. So, it seems that speaking performs two roles at the same time: That of communicating about the world and that of communicating about language. When we point at a car and pronounce the word "car", we are in effect not telling them that there is a car (they can see it, too) but that what they are seeing is called "car" in English. So, there is no communication of knowledge about the world but there is communication about the way the language works. As their skills progress, this will change. More and more language can be used to communicate content. But we will in actual fact never stop trying to adjust our way of talking. In debating whether or not the EU is democratic we may eventually not only agree about the facts but also refine or revise the meaning of the word "democratic" itself. It is interesting to note that while concrete objects can at least be pointed at, abstract objects cannot. It is not easy to imagine how we can teach the word "democratic" by pointing at something. Instead the method of instruction is verbal. We give examples or a definition. We explain. And we argue. We use language to show how language works. We use language to explain more language, to expand and modify it.

And thus we return to the problem of coherence. Consider the word "poverty". In addition to the intuitive meaning there are many technical meanings (like having an income of half the mean income, of sixty percent of the mean income, or maybe using the median instead of the average). If there is limited time to learn language we expect a natural variation in

the way people talk to begin with. As this variation widens over time, it may also impair learning. As a result, meanings are learned only in approximation. An additional problem is that there are many technical variants, which get used by statisticians because they have to have a precise definition so that they can use numbers. But that now creates a problem for communication. If the words we read do not match in meaning with ours, this can distort the facticity of our beliefs. For example, we may read that one third of the children live in poverty, but the word poverty does not have to mean the same as what we take it to mean, and that, say, if our meaning were employed it would be only one quarter. But not knowing that we now believe incorrectly that the fraction is one third. On occasion, though, the error gets noted. It is then that a conversation begins about the different meanings of the word "poverty".[126]

This notion of difference in meaning has achieved fame in the philosophy of science under the guise of *incommensurability*. Thomas Kuhn, who actually invented that term, discovered that idea when he was trying to make sense of Aristotle's *Physics*.[127] What Aristotle wrote there sounded utterly wrong, until Kuhn noted that Aristotle used the Greek word for "motion" in place of "change". Suddenly, everything fell into place. Kuhn then thought that the misunderstandings between physicists could in part be traced to precisely that problem. The terms they used had different meanings. This is what he dubbed *incommensurability*. It was that a theory could not be understood in terms of another theory. It had to be understood in its own terms, so to speak.

However, Kuhn did not say that there could be no understanding between practitioners of different theories. For he did suggest that after all progress was possible, so we could get closer to the truth. This would be the case, for example, if the terms could somehow directly be translated as in the case of Aristotle's "motion"; or if the use of these terms could be avoided. For example, even though the word "straight line" takes on a different meaning in general relativity theory, there is a way to talk about them using the standard notions so that certain unexpected behavior becomes perfectly understandable.[128]

Not everybody agrees with Kuhn on this last point. What Kuhn was getting at, though, was that to the degree necessary for achieving a consensus at all what was necessary was that one could agree on the phenomena that the theories would talk about independently of the particular languages that they used. This is pretty much the position of an experimental physicist. If general relativity said that light travels on a straight line but that space was curved through gravity this predicted that light would be able to move

[126] See Marcus Kracht – Udo Klein: Notes on Disagreement, in Daniel Gutzman – Jan Köpping – Cécile Meier (eds.): *Evaluations – Denotations – Entities. Studies on Context, Content and the Foundation of Semantics,* Amsterdam, John Benjamin's, 2014, 276–305.

[127] See the account in Horgan: *The End of Science*.

[128] The idea is that of embedding these curved spaces in Euclidean spaces.

"around" stars if they were big enough.[129] The calculations can be done, and we can make an experiment to test the validity of this claim. Indeed, this experiment has been performed. It showed that general relativity could predict what Newtonian mechanics could not.

The observation that incommensurability is a deeply rooted problem is due to Quine. He suggested that there is no way we can find out what the words of a foreign language actually mean. This has a radical consequence. Since children do need to learn a language, they too do not know what the words actually mean. In fact, nobody knows. Everything is up for grabs, and the analytic-synthetic distinction collapses. Of course, that would mean we can no longer talk of "the meaning" of a word.

This radical pessimism is not inescapable, though. To the degree that we are all humans we share certain ingrained ways of viewing the world (Chomsky uses Descartes' notion of "innate ideas"). And these innate ideas are uniform. Hence, basic linguistic intuitions can appeal to these ideas and thus guarantee a certain, let's call it, standardization. It would follow, however, that there are languages that humans cannot understand, a consequence that Chomsky fully endorsed. What constitutes such a language, and how much variation is theoretically possible, is anybody's guess at this point.

It has been suggested to me by Yannis Moschovakis at some point that semantics can make no progress as long as it does not have a notion of closeness of meaning. We can see exactly why that is so. Quine's position has just one problem, namely that humans do seem to communicate just fine. How is that possible if meanings cannot be determined by observation? The innate ideas are not all that it takes, for there is enough variation in meaning. What does the trick is that use of language allows us if anything to get *closer* to the actual meanings.

To sum up, language is actually not homogeneous. While Chomsky initially assumed a homogeneous community of speakers to derive learnability, it is obvious that language is all but homogeneous. This applies as we have seen to meanings. And it applies to syntax and, as everybody notices, pronunciation. In contrast to Chomsky, cognitive linguists assume that our language is the product of our individual learning histories and therefore almost by guarantee different, and the way the cohesion is guaranteed is actually by constantly using language to talk to each other. But if that is so, talking to

[129] It requires sophistication to understand how the notion of behindness is defined if straight lines are not straight any more. I will not venture into that here. Suffice it to say that general relativity predicts that any star is visible (that is reachable following light trajectories) even though in a standard three dimensional universe, anything that is directly behind the sun is not.

each other is more difficult because we also need to keep track of the ways in which other people can (and do) differ in their use of language. This leads to the topic of the next section.

Context and Interpretation

Reading a text or listening to a talk can be demanding at times. This has many reasons. One is obviously the content. A text that explains how special relativity theory works will be difficult to understand because the matter discussed is difficult. It has nothing to do with the language. At other times a text is difficult because it works inside a context that we don't know about. We understand the meaning but do not really know how to place the discussion. And finally some texts are hard to understand because of the language. This is immediately apparent if they are written in a language that we do not master. Or if we do understand the language but have difficulty reading, as would be the case if one learnt to speak but not write Chinese.

But all this aside there are inherent problems of language use that can make understanding a text difficult if not impossible. Take a word like "I". In a face to face conversation we know who this word refers to—the speaker. But in other contexts this is less clear. If we get a phone call from an unknown number, or get an anonymous letter then we do not know who is speaking to us—which may in some cases even be intentional. There are many more words of this kind, called *indexicals*. Other instances are "here" and "now". The magic these words perform is that we need to know the surroundings of the physical utterance in order to know what they are referring to. Say Socrates tells Plato at some of their discussions: "I disagree with what you just said." Suppose the next evening he uses the exact same words. Even though he technically used the same sentence, what he is saying is something different. He is voicing disagreement with a different thought. And this is just because the phrase "what you just said" identifies a different utterance by Plato. If at the first evening Plato had claimed that there are no irrational numbers while the next time he had said that knowing is not remembering then the first utterance commits Socrates to claiming that there are irrational numbers, while the second one commits him to claiming that knowing is remembering. And if Theaitetos had said to Plato on the first evening "I disagree with what you just said." then it would not have been Socrates who is committed to believing that there are irrational numbers, but rather Theaitetos. These are indeed very different things.

We can see therefore how much we need to understand the context in which an utterance takes place in order to know exactly what has been said. This is not a matter of nuances but a rather essential feature, because we

need to know to whom to attribute what belief. To see how important that is consider a lawsuit filed against some company resulting in the confiscation of thousands of emails, as it now occasionally happens. If the emails were dumped on the prosecutor with no indication of who said what and when and which email was in response to what other email, then quickly the task of finding out what actually happened would be hopelessly impossible. Indeed, it is not the content of the emails alone that gets copied but also the so-called header. These headers contain a lot of information: the email address of the sender, the date, the path the email took, the identifier of the email to which this is a response, and so on. So, each email unlike everyday utterances actually has an identifier by which it can be uniquely identified.[130] And this is precisely the information that is so vital. If someone simply writes to express agreement we need to know what email he responds to in order to understand what opinion he is expressing.

When the context gets lost, though, we are not without help, but the game is a radically different one. Suppose we have a manuscript written in Greek with no indication of authorship. Of course, we would very much like to find out who that person is. There are real cases of this kind. In ancient times there were many letters said to have been written by Plato, but we nowadays believe that only the seventh is actually by Plato himself, all others are from unknown authors. The matter is not trivial because if we are interested in the philosophy of Plato as many people in fact are then it is of utmost importance to know which books or letters he actually authored. For example, in his latest work, *Nomoi*, Socrates no longer participates in the dialogue (it is "an Athenian") and also ideas, an important element in Platonic philosophy, are absent. This is important if we contrast this with his most known work, "The State". Observe though that the matter is circular to some extent. We can claim that a certain manuscript has not been written by Plato only on the basis that we are sure that some other manuscripts were actually written by him. It is by scrutinizing them carefully (doing statistics or analyzing the content) that we conclude that the other manuscript has a different author. If the initial set of original manuscripts had been different we would probably arrive at a different conclusion.

When some work gets discussed by other people, especially when they are critical of it, then what they write will almost certainly not be a misrepresentation. In this regard it is a remarkable coincidence that we can reconstruct the philosophy of Chrysippus even though we do not have a manuscript by himself, but only discussions of his logical writings by Sextus

[130] This information can be faked, though. But that is beside the point.

Empiricus.[131] The reason why we are certain that Sextus Empiricus did a good job is because what he writes is clear and coherent. And the logical rules are certainly correct.

There are ingenious ways to check who is the author of a text, for example a statistical analysis of the most frequent words, but they completely fail on inscriptions. They are too short. When you go to Budapest and visit the Roman remains, you can read, for example, about many people that were buried there in the place once called Aquincum, but you don't have the foggiest idea of who these people were. Archeologists need to do many comparisons with other sites, and with some luck they can recreate the life of some person.

The context of an utterance is also important in identifying the language or code that is being used. Recall the case of the astronomers deciding in 2006 that henceforth Pluto no longer is a planet. Then it is important to know when some astronomer made a particular statement involving the word "planet" because he might be saying different things. Unlike the case of indexicals, there is no way to point to the language that we are speaking. Or rather, we rarely mention it. An astronomer may make reference to a different usage (by using phrases such as "in the terminology of 1890") but most conversation shies away from doing so. Language variation is actually quite an intricate issue. Often words allow to identify the time of utterance (because certain words are known to have been coined at a particular time), the region of the speaker (if he is speaking with an accent or uses certain region specific words), and so on. In turn, the knowledge of these things allows to identify what has actually been said.

Yet another problem is that of philosophical writings. If the uninitiated tries to understand a particular work by a great philosopher, say Heidegger, Husserl or Whitehead, then sometimes there are insurmountable problems. Part of it may be that we do not understand what the real problem is they are trying to explain or solve. And we might also be somewhat bewildered about the ideas that they are discussing. It may well be that we find the problems uninteresting or even nonexistent. But at times we simply can't understand what it is they are getting at. When Whitehead goes on about the "concrescence" he seems to say the unsayable. It is hard to understand the word itself in the way it is used in the text. Similarly, Plotin is obsessed with the "one" but it is one of the most difficult things to understand what it actually means. Especially in philosophy the idea that one can understand texts bottom-up, just by building up the meanings from words to sentences to paragraphs, is simply unworkable. One will have to read an entire book several times before the meanings of certain key terms become clear and we can begin to read the text as it was meant to be read.

[131] Benson Mates: *Stoic Logic,* California, University of California Press, 1973.

And this is not, to repeat, because the words are inherently difficult to understand. It is mainly because they assume a special meaning with the particular authors. These meanings must be intuited by the readers, and it sometimes needs a lot of prose until they get clear. Of course, these meanings change as the people change their philosophy. This is why it is of major importance to understand who is the author of what work.

To see where the difficulty lies I shall use an earlier example. As Ludwik Fleck pointed out, the word "syphilis" changed in meaning over the centuries. In the beginning, this disease like many other veneral diseases, was thought to be the punishment by God for immoral behavior, whence the German word "Lustseuche". Doctors grouped many diseases with similar symptoms under that name until it was possible to identify the antibodies. After that it was evident to doctors that the name was to refer to the disease that was thus identifiable independently of the symptoms. The medical research has thus not only changed our arsenal of knowledge, it has also changed the meaning of the terms themselves.

We are reminded of the twin earth thought experiment by Hilary Putnam. Putnam pictures a world just like ours with one exception: that the substance that looks like water, tastes like it, can be used like it, is made not of hydrogen and oxygen but rather of X, Y and Z. So, it chemical terms it is XYZ and not H_2O. Now, the question is: would we still call this substance "water"? This is in fact impossible to answer. In a society where all we have is our senses to check the properties of a substance we could not identify the difference, and we would therefore not think that it mattered much what chemical composition it had. Such was Western society until quite recently. In a modern society, where we get taught that chemicals are made from atoms and that the chemical formula determines the properties of a substance we would be less inclined to call the XYZ substance "water". We have to understand about this experiment that the word "water" actually turns out to have a gap in its meaning. The new hitherto unknown substance introduces a problem: what is it precisely that makes us call something "water"?

This is not a small detail. Kogge claims that Wittgenstein's philosophy has eventually gravitated around the problem of understanding concepts in the sense that all that is left for philosophers to do is to judiciously explore the terminology that is being used either in everyday life or in the sciences.[132] For after we know what a word denotes, clearly and unmistakably, we are ready to do science with it. Knowledge only comes after the code has been identified completely. But that, I might add, is an impossible job. The code evolves in tandem with our knowledge. Meaning and use cannot be separated.

[132] Werner Kogge: *Experimentelle Begriffsforschung. Philosophische Interventionen am Beispiel von Code, Information und Skript in der Molekularbiologie*, Berlin, Velbrück Wissenschaft, 2017.

PART II

MATERIAL CULTURE OF KNOWLEDGE

CHAPTER 5

CONSERVATION OF KNOWLEDGE

Archives

The poet Simondes of Keos is credited with the invention of mnemotechnics, an art used by poets and orators of the Antiquity, since speeches had to to be given without the help of manuscripts. It is told[133] that Simonides wanted to teach his art to the general Themistokles. He said, the advantage of his art was that one could make oneself remember anything that one wanted to remember. Themistokles inquired whether it was similarly possible to forget what one wanted to forget. Simonides had to admit that that was not possible. So, Themistokles kindly rejected the offer. Themistokles pointed at a fundamental aspect of knowledge: it could seemingly only grow. Our mind seeps up data like a sponge and does not let go of it. Nevertheless, the idea of Simonides was not pointless. While we do not forget things abruptly, our memory fades over time. We are not good at remembering exactly what we want or need to remember. So, what happens when we cannot remember anymore that which we think is essential to remember?

Knowledge is like the tea of Zen master. We have a limited capacity for storing it but it keeps flowing in.[134] So there eventually comes a point where we need to store some data somewhere even if temporarily. We can do that by telling another person to remember it for us. This is indeed the way early societies organized their knowledge. They had a few specialized people: a medicine man or someone who knew all the ancestors by name. Caesar reports that the Celts did not use any written records, but learned everything by heart.[135] The amount of knowledge that can be stored in this way is of

[133] Cicero De Finibus II.104. See Francis Yates: *The Art of Memory,* Francis and Kegan Paul, 1966, on the history of mnemotechnics.

[134] An anecdote tells us about a professor who goes to visit a Zen master to learn about Zen. The master is preparing tea while the professor is talking without end. When the tea is ready the master pours tea into the professor's tea cup but does not stop even when the tea starts flowing over. The professor complains to the master: "Don't you see that the cup is full? Why are you pouring more tea?" The master replies: "Don't you see that your mind is full? How can I teach you Zen if you do not empty your mind?"

[135] De bello Gallico VI.14.3. The education of a druid lasted twenty years according to Caesar.

course limited.[136] At a certain point the human brain was not enough. That was the point at which external media was being used to the purpose of storing data. This is what we today call an *archive*. Archives consist of material objects that we manipulate in a certain way so that they can become carriers of knowledge and data in general. The history of these objects is not only one of human culture at large but also one of human technology and invention. An enormous amount of ingenuity was necessary to arrive at the modern archives. It is important to take a good look at this technology in order to see how ephemeral our system of knowledge has in fact become.

The first archives were made of stone; at least these are the only ones that remained from the early era. A famous example is the Codex Hammurapi from the 18th century BC[137] or the inscriptions of Oinoanda from the 2nd century of our era, where the teachings by Epicure where "publicized" by carving them in stone, six meters high. At the time of Hammurapi people also used clay tablets. Their advantage was that one could write quite quickly and also erase something, as long as the clay hadn't dried. Clay tablets could also be burned.[138] For short term use Greeks and Romans used wax tablets. Their use is documented even until the 19th century. At that time chalk boards were introduced into the classrooms.

Apart from stone and clay people in Antiquity used parchment and papyrus for writing. The scroll was the medium of choice for longer texts. It ensured that the individual fragments were kept together. Parchment was the medium of choice until the Middle Ages, when paper got introduced. Paper was developed by the Chinese and brought to Europe via the Arab merchants. It was something of a revolution since it meant that the supply of the material, which was used both for packaging and writing, was practically unlimited at the time.[139] The next time this happened was with the introduction of the electronic data storage. Soon this led to an enormous increase in documentation, especially for commerce, and communication in the form of letters. In the Middle Ages the book had already replaced the scroll (made from papyrus rather than parchment), now paper replaced the papyrus. Paper, papyrus and parchment are not very much different from each other. What set them apart from stone and clay was their ease of use; you can create manuscripts easily with them, and you can carry them around. Already in Antiquity the intellectuals and their students circulated their texts all around

[136] How much data our brain can store is debated. John von Neuman estimated the capacity to be 2 times 10^{20} bits, roughly the size of the internet, see also Chapter 2: Theory — What is Knowledge? — Information and Entropy. Today the estimates are lower, around 3 to 1,000 terabytes (approximately 8 times 10^{15}). This is still a lot, but it is to be taken into consideration that most of it is not available for our own manipulation.

[137] This is not the oldest inscription but certainly the most widely known.

[138] It seems that many of them were just accidentally being burned in fires.

[139] Lothar Müller: *Weiße Magie. Die Epoche des Papiers,* München, Carl Hanser Verlag, 2012.

the known world, the Mediterranean. By way of example, the letters that Saint Paul sent out to the fellow Christians went to all places in the Orient. So, long distance correspondence did not start with the invention of paper. What changed was the amount that one could produce and the ease with which one could do that.

The city of Alexandria was known throughout the Mediterranean for its library filled with scrolls and manuscripts. Some estimates say that it contained half a million of them, which amounts to 80,000–100,000 modern books.[140] This was the largest library of its times; not even Athens or Rome had a larger one. When it burned down, the loss was unimaginable. In the Middle Ages it was the monasteries that had libraries and copy rooms, where they made sometimes multiple copies of manuscripts. Of course, manuscripts were copied all along: the kings of Alexandria searched every ship and confiscated the manuscripts, which were then diligently copied and handed back—the copies, not the originals![141] But it was only in late Antiquity that copying of manuscripts really took off with the establishment of the scriptorium. It wasn't only the bible and sacred texts that were being copied there. Basically the entire literature of Antiquity was being preserved by the efforts of scribes. An early example of this effort was the Vivarium of Cassiodorus in the 6th century AD.

In addition, there were no publicly available materials and texts for students of a university. So, they had to either copy it themselves or pay someone to copy it. Again, monasteries copied manuscripts in exchange for money.

The printing press changed all that.[142] On the one hand it ended the practice of manual copying. From now on manuscripts could in principle be issued in any number of virtually identical copies. On the other hand it offered rich people the opportunity to establish their own private library. It would still take centuries until ordinary people could afford the possession of books. Initially it was one book that they could at least get with a modest amount of money: the bible. This already was a big step given that the most important of all books was pretty much verboten for laypeople before the reformation. Initially, printing books was quite expensive. But at least the dissemination of new ideas was greatly enhanced since it was possible to mass produce books and send them around the world.

Printing was done by assembling movable letters into blocks that were then pressed onto paper. In the 19th century the assembly of letter types was mechanized. Around 1950 phototypesetting was introduced, which

[140] Scrolls could contain any amount of text, and their size was only limited by practical considerations. They could not be very long or else one was not able to use them.

[141] Martin Lyons: *Books – A Living History*, London, Thames and Hudson, 2011.

[142] Elizabeth I. Eisenstein: *The Printing Press as an Agent of Change*, Cambridge, Cambridge University Press, 1983.

dispensed with the metal types altogether. This made printing far easier and cheaper. Not long after that the first copy machines were produced, which allowed to copy an existing manuscript or book in no time.

The development of printing introduced its own demand. The 17th century saw the appearance of the first newspapers and journals. From then on their number and distribution skyrocketed. The first scientific journals got introduced in the 18th century. Today each university discipline counts hundreds of journals. While teaching material was very hard to get in the early days of university, today it is ubiquitous. The photocopier allowed to reproduce material with almost no effort. And even paper is slowly losing ground. Increasingly, teaching material is offered electronically only, but can be printed at any moment.

The costs of producing books have dropped significantly. Every year 100,000 new books appear in Germany alone. The national library had 27 million books in 2011 (compared to 100,000 in Alexandria 2,000 years earlier). Since that many books consume a lot of space the microfiche was introduced, which is nothing but a miniature photocopy produced on film. It could be read with a big instrument that was nothing but a specialized magnifying glass.

The electronic age has multiplied the possibilities once more. The new storage media, hard disks, CDs, DVDs, magnetic tapes and storage chips, have a capacity that varies between a few gigabytes and a few hundred terabytes (in the case of high density magnetic tapes). Just for comparison: this book consumes in its electronic version (PDF) about 1 MB, so these types of storage devices can carry from a few thousand to a few hundred million of times as much information. The library of Alexandria can be fitted onto the hard drive of a laptop.

The expansion is continuing. The next phase is the cloud and the data-centers, which I will discuss later. The cloud is a network service, which connects all users on this earth with the data-centers. Google alone possesses hundreds of thousands or maybe even millions of disks in a dozen of data-centers that store all information that Google is able to collect, including information that allows to efficiently search for and in the documents.[143] The cloud itself is just the management engine that interfaces this hardware with its users. It could not operate in this way without all the data-centers. These are the modern archives, they story almost all information available on this planet.

Archives were developed because of the need to store information. At the same time they have themselves created an expansion in stored knowledge, that made further developments necessary. The modern age of print-on-demand that can produce a book at the cost of a few thousand dollars

[143] https://www.quora.com/How-many-data-centers-does-Google-have?share=1, which claims that in addition to the dozen official centers there are hundreds more. See https://royal.pingdom.com/2008/04/11/map-of-all-google-data-center-locations/ for a different statistic.

has made the publication of ideas and texts feasible for practically everyone. Where a Babylonian priest would have to communicate face to face with his students, there is no more necessity for anyone to move for the purposes of getting information no matter how arcane the subject matter and how remote the person that could tell us anything about it.

The ubiquitous availability of knowledge has its downsides, though. The massive creation of new content reproduces the old problems, only at a more elevated scale. Digitalization projects (Project Gutenberg, Google Books) are not only tools to make information available. They also serve as a protector of knowledge. They preserve it.[144] With growing volume of content the problems of access grow as well. It is not enough to have the data somewhere; it must be properly administrated as well.

Administration

An archive consists in a collection of documents and a system of storage so that documents can be found. So, it is more than a simple storage. Only the archive makes the embedded information actually accessible. The definition is wide and encompasses a great variety of physical implementations. It can consist in a single book, several books, or several documents. The books can be the books of a merchant, a company, a single person or a library; the documents can be legal documents, they can be administered by the church, the town, a university, an NGO, or a private company. They can contain contracts, letters, search warrants or shipping lists. The physical form is likewise not restricted. All the above can be made of parchment, paper, or simply be in electronic form. In most instances these documents are instruments of administration.

Administration is not only something for institutions or companies. Every individual has to do a fair amount of administration. It starts with managing a check book, or a household shopping list. And it continues with letters to various offices, most notably the tax office.

The management of an account is one of the most important affairs. It used to be done by merchants and banks, and they were mostly the only people capable of doing so. The word "bookkeeping" reminds us that the art of keeping score of the financial balance between the parties involved was by means of a book that contained a record of what happened. It was a special art to maintain it so that it could function in the desired manner. Today, practically everybody learns to read, write and calculate in school and is therefore assumed to be capable to perform some of that bookkeeping for themselves.

[144] We shall return to this subject matter in Chapter 5: Conservation of Knowledge — Durability.

The amount of paperwork that even an individual has to perform has increased over the decades. The same goes for institutions. The number of documents and the detail at which information must be supplied is sometimes staggering. The ease with which this can be done depends in large parts on the order that is maintained in the archive. Any piece of information must be retrieved, and the larger the body of data the more time one needs to find it. In the worst case we need to scan all the documents linearly until we find what we are looking for. This would be the case if we would simply throw all documents on a single pile.

An interesting parallel exists here. In Jewish tradition, religious texts could not be destroyed even when no longer needed. Instead they were thrown into a room, called the geniza.[145] This was a room next to a synagogue, whose sole purpose was to house these documents. As there was no intention to read them again. They were simply thrown on a pile. Evidently, if ever the need arose to reread one of them, the search would have been extensive.

A geniza is not an archive. What it lacks is order. It is not designed to contain documents *for use*. And that is an important criterion. An administration, at least one of modern proportions, cannot work without an archive. The archive is its memory. And as such it contains information together with a system that allows to retrieve this information and change it when necessary. So, the archive is simply a device by means of which an administration can function. Administration in turn exists in two senses. In one sense it is an organization with a specific purpose, in another it is an act that is being performed in that function. A civil registry, for example, contains documents on marriages, births, and deaths and other vital information on people and their relationships to each other. The archive contains documents, and these documents can be sometimes modified, and new ones can be added. The administration is tasked with deciding what kinds of documents must be added or changed, and with handing out information concerning people on request. The administration has to make sure that the archive is correct. For example, if a baby is born, a document must be entered stating the obvious information: name, date of birth, parents, and so on. A library likewise has an archive, containing information about the books that it owns and the people that may take out books, and which people currently have taken out which books. Every time a book gets taken out or returned, the records must be updated. This is a delicate process especially when the number of books is in the tens of thousands. No wonder libraries were one of the first institutions to push for automated processing of information.

The tasks that a computer has to perform seem easy given the speed with which computers operate. However, data not only needs to be stored but also processed so that information can be retrieved easily at any moment.

[145] Müller: *Weiße Magie*.

The algorithms that perform this task are hidden from the user, so it appears as if the computer is actually only storing the data. There is more than one discipline in computer science which concerns itself with how to store documents and link the information.

The task of maintaining an archive often incurs a great power on the administration. For it is often the only place where a certain piece of information is being kept, or the only trusted party in a dispute. Max Weber was so impressed by a functioning administration that he declared the lifespan of an empire to be more or less determined by its ability to maintain a functioning bureaucracy.[146] The documents that are being kept in a civil registry and similar institutions however are not simple records. In actual fact they often constitute the reality that they are supposed to document. For example, if you have to prove you are married it is not enough to show photographs or other memorabilia. What is often necessary and sufficient is a document issued by the civil registry that certifies this. In principle the registry could fail to issue the document or worse never enter the fact that you married into the archive or simply not find the documents once entered. But that hardly means you could substitute the document; it only means that one has to make sure that the registry's archive is in agreement with the facts as far as your marriage is concerned and then proceed.

A spectacular case of failure has been the so-called "Robo-Signing-Scandal".[147] The banks together with their attorneys had failed to create the necessary documentation once a mortgage contract had been signed by the relevant parties or simply signed off on papers without properly checking the information. As a result thousands of people ended up with a mortgage they never asked for or with their house being taken for alleged failure to pay the mortgage. It took years to correct the records.

These records are as essential as they have been in Roman times. The Roman cadastre was invented as a public record to establish which piece of land belongs to whom.[148] Not all countries or empires did something similar. To this day, Greece has nothing of that sort. It is simply known which land belongs to whom. The cadastre is obviously very helpful in disputes over land ownership. It is said that the Middle Ages began at the point the Roman cadasters got destroyed. For at that point, land ownership was up for grabs (again).

[146] See Max Weber: *Wirtschaft und Gesellschaft* (5th ed), Tübingen, J.C.B. Mohr (Siebeck), 1985. Essential in his view was the ability of a bureaucracy to perform its tasks without being influenced by other powers. His role model was the Prussian state which created a bureaucracy consisting of paid independent bureaucrats whose salary did not depend on the whims of other people with greater power.

[147] See https://en.wikipedia.org/wiki/2010_United_States_foreclosure_crisis.

[148] The oldest cadastre that is still preserved is that of Orange in Southern France. It documented what lands have been given to which soldiers as a gift in reward for their service.

The administration is supposed to maintain the archives. In computer science jargon it is tasked with truth maintenance. That means that the archive should be a consistent and truthful mirror of reality. For example, if a person is getting married, the administration has to see to it that its records get updated so that the person is now recorded as being married. Although in principle the archival function follows the world, this is not always guaranteed. Another problem that often occurs is that information is needed in several places. In the pre-computer era this was especially difficult since that meant that several copies of documents had to be made that had to be the same and distributed to the various places.

We can still see reflexes of this process. In many countries before you enter you have to state through which seaport or airport you wish to enter. You cannot change that. This is because the document used to get sent to that exact location. If you tried to enter elsewhere there would be no records that showed you were allowed to enter.[149] Also, birth certificates were issued only at the place where you were born and the original records are being kept there. Presumably this is to ensure that you get a copy from an original only. Again, in the old days there used to be no central archive of any sort.

The amount of data that is now being stored for administrative purposes is staggering. For example, companies are monitoring the movement of items, money and information and store them for several years. Typically, they pay specialized companies to do that for them. The amount of data thus generated and stored for just one company quickly exhausts the capabilities of simple storage devices.

The biggest concern however in addition to the archival function proper, namely truth maintenance, is the maintenance of consistency. Every time data gets copied there is a risk of error. Be it manuscripts or IBAN numbers, be it text files or birth certificates, every time a new copy is being issued new errors may slip in. Of course, one could compare the copy with the original. But this is not always feasible. A copy of a manuscript is too long to allow for effective checking. Any proofreader can attest to the fact that errors are not so rare (a professor once suggested one error every five pages is the best you can do, everything beyond that costs way too much effort). With birth certificates the problem is being reduced by producing copies from the original alone, so errors do not compound. With account numbers (books, bytes, banknotes, and so on) one has introduced a system of error correction that I will explain later.

[149] Note that the visa in your passport is one half of the documentation. The other half is what the custom officer is shown on screen and has been sent to him by the authorities.

Durability

Knowledge can age in the sense that it is superseded by newer, better knowledge. This is often the way we see scientific progress. Progress results in falsifying or modifying previous ideas. Though this rarely happens in mathematics since proofs if correct constitute unshakeable evidence for the truth of a claim, empirical science relies heavily on experiments, whose results are not always replicable. And often new technology allows for better measurements or new kinds of experiments. When the new facts appear we simply say that what we used to call knowledge no longer can be called knowledge because we found out that it is not true. It is then it is not what you read in the books that has changed; and it is not the case that people did not believe earlier that it was true. We say that *we now know that it it not true (and never has been).*

However, knowledge can also age in another sense. It becomes less clear or it simply disappears. The clearest case is our memory. We do not remember things that happened a while ago in all detail. In fact, as time goes by, more and more details get lost or appear somewhat vague. We remember talking to someone and cannot recall where we said that or what exactly we said then. We remember a car provoking an accident but cannot remember how exactly it looked like. We may even fail to recall its color.

Quite similarly manuscripts decay over time. What was once a completely legible manuscript slowly loses this quality over the course of centuries. The parchments breaks, little parts fall off and disappear, the colors of the ink are becoming subdued and merge with the colors of the parchment.

This is not the process of knowledge suddenly failing to be knowledge in our eyes. It is not the content that has changed. It is the physical nature of its carrier that has changed and that makes it more and more difficult to know what the actual message has been. The laws that govern this process are the laws of thermodynamics, which we discussed briefly in Chapter 2: Theory — What Is Knowledge? — Information and Entropy. The global message is simple: in the long run, chaos prevails over order. More precisely, it is not chaos that increases but uniformity. Order in thermodynamic terms is absence of uniformity. However, thermodynamics is a rather precise theory. It puts actual numbers to this. It defines *entropy* as a measure for order.

$$S = k_B ln W \tag{5.1}$$

Here, k_B is the Boltzmann-constant, W the number of microstates and ln the natural logarithm. Since the unit of k_B is J/K (Joule per degree Kelvin) the same holds true of the entropy. The entropy is therefore increasing with the number of microstates. Now here is one of the fundamental laws of thermodynamics.

2nd Law of Termodynamics: In a closed system entropy never decreases.

Crucial to this law is the definition of a closed system. It is a part of the physical space such that no energy or matter passes through its boundaries. Earth, for example, is not a closed system because energy is sent continuously from the sun at an average of 1,361 W/m^2 or 1.361 kW per square meter. That is actually a lot of energy. On the other hand, earth also radiates out energy.[150]

When a system is closed, the total energy contained in it cannot change (that is the first law). All that can change is the manifestation of that energy in terms of heat distribution, for example. No system is closed in the technical sense, but some of them get close to this ideal. Thermodynamics can also deal with open systems, as long as the energy and mass flows can be calculated.

There is a statistical interpretation of this law that interests us here. It rests on a distinction between macrostates and microstates. We will see that macrostates are the equivalents of information, but for now we need not understand that parallel. The more microstates are contained in a macrostate the more stable it is. This is because the system is in permanent fluctuation. The system changes its microstate since little changes happen all the time. These are the forces of chaos if you will. All these little changes will go unnoticed, however, when the macrostate is insensitive to the changes. In other words, while the microstates change the macrostate can remain stable as long as the new microstate is in the same macrostate. The trick in thermodynamics is to treat order as equivalent to energy. That means we can create order in a system by injecting energy into it. This can happen in various ways. We can use a printer to distribute little colored dots on a paper to produce a copy of this book, for example. Or we can use a monitor to continuously send photons or electrons of different wave length so that the eye will see a similar image. What is relevant for the discussion is that the individual letters are decomposed into little cells, so-called pixels, and that these pixels get assigned a single, uniform color. These are the dots. Now, if the pixels are sufficiently small our eye cannot see them as dots. Rather, all it sees is a combination of lines and shapes that constitute a letter.

We can now fix a rectangle of pixels and run through all possible distributions of black and white, say, and check whether they constitute a particular letter, say 'a'. This forms the set M. M is our macrostate. It represents the letter 'a' in that it contains all manifestations in terms of microstates. Let us now track a particular instance of that letter just printed on paper. It is probably quite clearly recognizable. Over time however its shape deteriorates. Some part of

[150] The concern over climate change stems from the fact that the greenhouse gases inhibit this radiation. So, the atmosphere cannot cool in the same way as it used to. The energy balance of this process is actually reasonably well understood in contrast to the many subprocesses that create the overall phenomenon.

the ink spreads to adjacent pixels so that they too become black. Or some of the pixels become lighter and lighter so that we start seeing them as white. In this way, the microstates change in a sequence *m*, *m'*, *m''*, ... As long as they are in *M* we shall read them as 'a'. Once they are outside *M* we no longer read them as 'a' and we start guessing. Maybe the letter we are seeing is actually an 'o' or some other letter? Or it is not a letter at all?

There is a tradeoff between stability and amount of information. We can put more information on a page by using smaller font. But a font metric[151] determines the number of pixels that make a letter. The smaller the font the smaller the number of pixels that constitute a letter. The smaller that number, the smaller the number of microstates that the macrostate contains. And so, the number of moves we need to make to a legible version of 'a' into an illegible version (or a version of some other letter) is reduced. In other words, if the print is smaller, not only is the text harder to read, it also decays faster. It takes less time for the manuscript to reach a state where it gets hard to decipher.

The case is no different in a handwritten manuscript. Although we cannot speak any more of fonts, pixels, and often not even about individual letters, the case remains fundamentally the same. At the core there a process of distributing ink on paper in a certain way that could in principle also be recreated by a printer (for example, just scan that document and print the resulting file). You can now start subjecting that handwriting to the same process of "smearing" the ink slowly over the paper. Depending on how legible the writing was to begin with it takes decades or centuries until the legibility starts being severely compromised. There are of course other factors that exacerbate the problem. The quality of the ink or paper deteriorates as well. And the paper eventually breaks, and some of its parts may get lost. This makes reconstruction of the original state of the manuscript and its actual content sometimes an impossible task.

If entropy cannot decrease it will eventually increase. And that means in increase in disorder, or, as I said, an increase in uniformity. So, order is subject to decay. It follows also that order is equivalent to energy, and to create it we must pay with energy, that is, disorder at some other place.[152] Knowledge is threatened by this process. This process is not the same as the willful destruction of the order, for example when in the case of a palimpsest the ink

[151] The font metric states how big the individual letters are physically.

[152] Evolution on earth was possible precisely because of solar radiation. The constant influx of energy made it possible for plants to harvest this energy to create an order not possible elsewhere. However, the point discussed here remains. The order I speak of here is order in the sense of what humans want the world to be like. Nature does not understand or respect that order, and so humans must maintain it themselves. The price is consumption of energy.

was scraped off to make space for another text.[153] This willful destruction could be kept at bay, though again a certain price must be paid. However, the thermodynamic process is not controlled by will and is unavoidable. Physics cannot be bribed into turning a blind eye on something. Instead, halting the process of decay means taking appropriate measures, and to be effective they must be paid for in terms of energy. Every archive is subject to this process, be it manuscripts, books, CDs, DVDs or hard drives. They contain the traces of the knowledge. Once the trace disappears, the content disappears with it. And there also is no difference to how our brain works. If its inner structure gets damaged, we may not remember anything any more.

We can witness the passage of time on any object. Even a message carved in stone is decaying over time. Maybe the time span is not any more centuries. Now it is millennia. The forces of chaos, small as they may be, will erase any message over time. And the direction is clear. It is not that order can emerge out of chaos. It is always the converse. This borders on the trivial. The message is there not because physical forces brought it into being. It is there because someone had the intention of putting it there. But surely no one will think that the message can be issued by just producing some random array of matter and then just patiently waiting for the order to appear.

It is even worse: even if the order would appear no one would notice because it would be entirely without meaning. I mention in passing that it is somehow similar to spontaneous revelations. If God would reveal the secrets of life to a child or an uninterested person it would be of no use whatsoever. Or if God were to speak an unknown language. Hence the idea that order can come out of chaos is misleading. For it is not order that we are interested in but meaning in human terms. And we have to define the terms first before the meaning can be revealed to us in an intelligible way.

Thermodynamics does not respect our human order because it is statistical in nature.[154] Essentially, if you have a microstate all changes are to some reasonably adjacent microstate. However, that process cannot be controlled. When you throw your mail onto a heap you cannot expect it to pile up neatly over time because statistically this is unlikely. There simply are too much other states it could go into. And most of them are as chaotic as the initial state.

[153] As I explained earlier (see Chapter 2: Theory — What Is Knowledge? — Information and Entropy) the actual writing of a bit of information into a device is the energetically costly act. The reason for this is precisely *because writing onto it means destroying the information that resided previously in that information cell.* Thus, at least theoretically, it is not the data itself that costs us energy but its preservation. A quantum computer can calculate at no cost, as some physicists claim, but the results if they are to be of any use must be written onto some storage device because the quantum state loses its coherence. And that is where energy is needed.

[154] Arieh Ben-Naim: *Entropy and the Second Law. Interpretations and Mis-Interpretations*, Singapore, World Scientific, 2012; Arieh Ben-Naim: *Entropy Demystified. The Second Law Reduced to Plain Common Sense* (2nd ed.), Singapore, World Scientific, 2008.

Now, as much as the process is open in the future it is also open in the past. For a given microstate *m* there is any number of microstates from which it could have been produced. And so, even though the manuscript evolved from some definite microstate to the one we are seeing now, this microstate cannot be reconstructed. The longer the process lasted the more states it could have potentially been in at the beginning.

This bring us to a paradox of recent and not so recent technology. The exponential increase of storage capacity witnessed in the recent decades has a downside. The information must be packed very densely onto the material. It thus gets harder to read. You need highly specialized equipment to write or read out the information. Moreover, the energy consumption of these devices is rather modest given that people routinely store gigabytes worth of data. And that means that the lifespan of that data is severely limited. Indeed, while the Codex Hammurapi is still legible (after more than four thousand years), the life span of hard drives is a mere five years. Of course, the worry is not new. The above mentioned Cassiodorus introduced the copy rooms in order to preserve the knowledge by copying the manuscripts. Charles the Great had a special letter type introduced (the Carolingian minuscules) partly in order to ensure legibility also over time. And abbot Johannes Trithemius of Sponheim complained in 1494 about the spreading of paper replacing the parchment, because, as he said, these documents were supposed to be made for eternity, while paper did not last as long as parchment. Compared to that our modern equipments are pretty short lived. The notions of sustainable storage or even of preserving our physical culture in the internet or cloud seem somewhat strange to say the least. Given that hard drives last approximately five years so that Google has to replace tens of thousands every day it is appropriate to ask in what sense data is being stored sustainably.

What this suggests is firstly that most of the data is not as essential as it first appears. Much of what we store even in terms of paper loses its value after five years. So, it could easily be dispensed with when the hard drive gets replaced. However, if there is such a thing as perennial knowledge and tradition, how is this possible to maintain? How could mankind have preserved the odes of Horace who said that he had erected a monument more durable than iron? Did he not know what kind of work he demanded of future scribes to copy his works letter by letter so that they could indeed last for the millennia to come? In all likelihood, he did not care so much about millennia, but at his time, when parchment was still being used, it lasted longer than anyone could imagine. Now however we need to care. When CDs were first invented it was said that they last a lifetime. But soon libraries discovered that 15 or even 10 years is more accurate. So, they now routinely replace them. What is taking a toll on them is something as simple as dust, known at least to engineers as an aggressive substance. Dust may be a nuisance when you have

a lot of books, but they will last more or less a lifetime regardless. Paper, especially archival paper (which is in particular acid free) easily last a century. Parchment lasts for several centuries.

According to Moore's law the density of storage is doubling every 12–14 months.[155] This is good news, as material gets manufactured that can store more and more data. And the energy costs for managing that data is declining as well. The capabilities of a modern day laptop surpass those of the computer of an entire research institute 50 years ago. The speed at which they compute has increased according to estimates over a million times in the last 30 years. One factor of a thousand is due to better algorithms, and another due to increase in speed. The downside is that computation and storage are more error prone at all levels.

The storage chips manage to squeeze more and more information onto the surface. This means that the conducting lines have a thickness of several atoms. We can now start to think of atoms as pixels in the layout of the circuit. It is clear then that as long as we design the chips that way we must make sure that the atoms do not get disturbed. Consider however the size of an atom against the size of a dust particle and you can see why preserving the integrity of a circuit is not trivial. Another known factor is the radiation. Radiation consists in particles (photons, electrons, protons, and even bigger particles) that fly around at varying speeds. Radiation is capable not only of disturbing computers, it can compromise the layout of a circuit.[156] This puts an end to the idea that computers cannot make mistakes. This is physically impossible. Not only can there be faults in the manufacturing of microchips, not only can there be problems with their design. Even if everything is perfect the presence of radiation means that every now and then some bit gets randomly changed.

Of course, once the problem is acknowledged one tries to counteract it. One method is to use error detection and correction. A pretty simple method is the 7-bit ASCII code. It works as follows. Of the 8 bits contained in a byte only the first seven are significant. Hence the code contains $2^7 = 128$ symbols. The last bit is calculated so that the number of 1s is even. (In binary this means that the bits of a code word sum to 0.) So, if the original word was 0011010, then a 1 is added to that the number of 1s is 4, an even number. Thus, 00110101 is the correct code word. Crucially, 00110100 is not a code word. This means the following. If a code word is transmitted and a single bit

[155] Or at least it used to grow at that rate. Rumors have it that in order to keep up that speed, data would have to be stored vertically onto data chips. To this day, they are basically horizontal arrays of storage. However, to manufacture such layered designs totally new types of factories would have to be built, at very high costs.

[156] See https://en.wikipedia.org/wiki/Data_corruption. It quotes a study by CERN saying that data corruption is in excess of 1 in 1016.

is randomly changed the result is recognised to be a non-code word. It does not matter where the change occurs. Thus, if the third digit of 00110101 is changed we get 00010101, which contains only three 1s and thus cannot be a code word. This code is useful for error detection. However, it cannot correct the error since it cannot detect where it occurred; and it will not detect a double error, as the number of 1s is once again even.

However, just as there are methods to detect errors there also are methods for error correction. Here is a simple one. Instead of writing the message in plain, say 101, we write each bit three times, like this: 111000111. The number of bits in the transmitted message is divisible by three, and it can easily be segmented into triplets and the triplets be replaced by their first member, say. Now let's suppose one of those triples gets manipulated; so we receive 101000111, with en error in the second digit of the first triplet. Then it is easy to see what the original message has been, as the two other bits are in agreement. So, the 0 is corrected to 1. And this methods corrects any mistake as long as there is at most one per triplet. If there are two mistakes in a triplet, then the correction will fail. So, if instead of 111000111 two errors occur in the second triplet, for example 111011111 the sequence will be corrected to 111111111 and thus it is assumed that the original message was 111 rather than 101.

It is easily seen that the price of error correction is a multiplication in the amount of transferred data. With n bit of message the actual transfer requires $3n$ bit, three times as much. In practice, leaner versions exist.

Now in addition to the ephemeral nature of data because of its increased density there is quite another problem. To manipulate the data we need to rely on quite sophisticated technology. While a book can be read by anyone with good eyesight, already the microfiche is too small to be read with the naked eye. However, a device for reading them can easily be manufactured. Not so with the recent storage devices. There is no way to make devices for reading CDs, magnetic tapes and others from scratch. The technology is so sophisticated partly because of the minute size of the data blocks, that we have to rely on a few specialized companies to make them for us.

Finally, consider also the problem of encryption. It is not even enough to be able to read out a device in raw form, say as an incredibly long sequence of bits. You may for example try to look at the raw code for documents.[157] This raw code is pretty useless. And that is due to a number of factors. One is that data may be intentionally encoded. Another is that data is mixed with markup that

[157] This is possible. I use VI, an editor that lets me look at the raw data (in hexadecimal form, which is just a convenient form to show the individual bytes). The present manuscript was originally written more or less directly in this form, with layout instructions interspersed. In other document formats such as docx or even PDF this is not possible. They are incomprehensible.

contains meta-information, layout instructions or hyperlinks. And then all this is beefed up with information that is being used by the computer for error correction and other administrative purposes. To make matters worse these encoding processes are cascaded so that at the end it is virtually impossible to retrieve the actual information.

This is no small problem. If you think about the fact that administrations worldwide use data that is encrypted using firmware you will realize that they have become totally dependent on the maker of that software. If the service is discontinued they will not be able to access their data. So, far these companies have not tried to use their monopoly, probably for fear of losing it eventually. But there is a lot of of so-called malware that works in this way. It encrypts data and decrypts it only when the owner pays the ransom.[158]

The problems of storage have been acknowledged. For example, a group of researchers have tried to raise awareness for the problem of long term storage.[159] The problem is being discussed there at length. However, it starts with a premise that this book does not share: that the electronic data storage is here to stay forever. It discusses the preservation of culture at large through the lens of a techno-optimism. *If* technology progresses in the same direction then we need to talk about how it can be made sustainable in the same manner as the parchment based culture has made itself sustainable through the copying of manuscripts. *If however* electronic storage will not expand or be inherently limited then we also need to talk a scenario where the electronic storage needs to be sustained by nonelectronic means.

The signs are already there. At the time of writing there has been a conference of librarians in Germany where they talked about the costs of digitalization. With funds for research and libraries not expanding there is now little room to digitalize manuscripts or convert other media into electronic form. They will then remain non-digitalized and we need to ask whether it is prudent to abandon the technology altogether if it means not being able to use them anymore.

Knowledge as a Mass Product

Modern industrial society has among other things introduced mass production. It consisted not only in increased numbers but also in standardization. This has not only revolutionized production. It has also changed our mental

[158] Exactly this happened to British Hospitals and other companies such as Deutsche Bahn with the so-called WannaCry ransomware. It is not clear whether or not money has been paid to have the data decrypted.

[159] See www.berliner-appell.org and Paul Klimpel – Jürgen Keiper (eds.): *Was bleibt? Nachhaltigkeit der Kultur in der digitalen Welt*, iRights.Media, 2013.

frame. For one of the effects of large scale production was what is known as economies of scale. This is the effect of rapidly decreasing costs per unit as the numbers go up. For once a factory has been built it can produce goods in hitherto unimaginable quantities: cars in the hundreds of thousands, computers in the millions, and so on. This means that the costs of the factory matter less and less since they can be distributed over the units produced. And, increasingly, humans are being replaced by robots, which are cheaper as they work around the clock without interruption.

The rapidly falling costs of computer equipment are a case in point. In my childhood a pocket calculator was an engineer's dream. Programmable pocket calculators existed but where very expensive. As a student I witnessed the introduction of PCs for the mass market. They were clumsy and had very basic screens. Then came the laptops, the tablets and after that the smartphones. Each of these products was hailed as a means to further the democratization. Especially the internet was touted as a means to give virtually free access to all information worldwide. And indeed, Wikipedia is an example of the powers of these new machines. Despite its drawbacks and the suspicion with which it is looked at, the information it gives in areas that are not en vogue or political is pretty robust. And the amount of topics which are being covered are enormous, not to mention its offsprings that use the same technology (wikis) for their own purpose.

But every technology has its downsides. There is a common process of innovation, which I call rationalization. At first a new technology is introduced because it solves certain specific problems.[160] In the first round, when all goes well the invention indeed simplifies the lives of people. However, as time goes by and more and more people start making use of the technology it is more and more taken for granted. Others simply assume that we have it and use it and therefore they rely or make us rely on its services. At this point there is no turning back: you simply have to use it because otherwise you are too slow or inefficient. Consider for example the introduction of the mobile phone. In the era of landlines you had to be near the telephone for other people to reach you in person. This is a problem for people that have to go around such as merchants or workers. The mobile phone solved their predicament. Now they could carry the phone with them and everybody could reach them by phone if they so wished. Initially, having such a phone was a clear advantage over others. You could be reached by phone when they couldn't. So, everybody tried hard to get themselves a mobile phone. At a certain point, though, it was taken for granted that you had one. And, as more and more

[160] See the fascinating book Beniger: *Control Revolution*, which asks precisely that question about the computer: why did it arise at the moment it was invented? In other words, Beniger wants to know what problems people were trying to solve by means of the first computers and how they did that.

people used it the number of phone calls one was getting through the phone was skyrocketing. The phone became a nuisance to some. And bosses quickly realized that it provided a means for checking on their subordinates or for giving them orders at odd hours. The workplace started to invade the homes of people. There was no more question of whether or not to have a mobile phone. It was compulsory.

The same goes for knowledge and information. Fifty years ago it was not such a trivial matter to get access to information. Even finding out where a particular office was and when it was open to the public was a nontrivial matter. The internet solved that issue overnight. All businesses and institutions started putting their information on the web. It became easy to find out about them from any location at any time. However, once information became available online it became compulsory to use it. This places a lot of people at a disadvantage. The lady at the office in Budapest being told that she had to get herself a waiting number online was visibly irritated. How, she asked, could she go online to request such a number? That however did not matter to the clerk. For him the matter was resolved once and for all. And that is just one example of many that I have seen.

But more is to come. Consider the idea that you have to know all the rules. Before the law you cannot plead ignorant. According to Roman law, the law applies whether you or not you know about it ("ignorantia legis non excusat"). This principle is important and is applied in common law as well as other legal systems. It would be otherwise too easy to bail oneself out by claiming one did not know that what one was doing was actually illegal. Yet the more laws get introduced and modified the more problematic this becomes. The ideal according to which someone acts in full knowledge of the law is increasingly becoming unattainable. In fact it may increasingly occur that people believe that what they did was not even immoral. A case in point is provided by laws against harassment. A harassment policy declaring that offering a date to a person who has previously declined one such offer from that person is an offense is simply arbitrary. Most people would not think that such behavior is in the least immoral. So, they have no reason to think that it would be under a policy to prevent harassment. There are many more such things. The point here is not whether or not that particular policy is good. It is that once the policy is issued it is assumed that people at least know about it. For what else would such a policy be useful for?

Unless we accuse lawmakers for arbitrarily introducing laws in order to provide enough pretexts for punishing people, passing laws can only make sense if the relevant parts of society come to know about it and have a chance to adjust their behavior. However, the speed with which laws are now created is mind-blowing. At one point it was even argued that the German parliament should allow to vote on a bill without discussing its content.

However, if even discussing a law before ratifying it is such a time consuming affair that lawmakers are considering its elimination we are in dangerous territory. In the presence of an increasing trend to regulate all domains of human behavior people will resort to a cost-benefit analysis. How important is reading this particular piece of legislation in terms of the risks I incur by not reading or understanding it? Can I skip reading the conditions and terms of use for this software or is there something in it that is useful or essential for me? Coleman speaks in this connection about an asymmetrical society.[161] More and more it is not one individual against another individual but rather one individual against a corporation. Whether you are buying a product or want to claim money from an insurance company or file a tax return, whether you are applying for a job or a place at the university, you are always faced with an institution that knows all its rules, and has made them to fit its own purposes while you have to find out what the rules are and how they affect you. This is why in the end the individuals band together to form lobby groups that help them so that the battle is fought between two equal competitors.

It is often claimed that the internet finally democratizes the entire society since all knowledge is openly available. But that is not even true. Entire countries can and do prevent their citizens from accessing the internet; companies keep most of what they do secret. Even though information is mostly stored electronically this does not mean that everybody gets access to it. And, it is still true that there are effective tricks to hide the truth even if it is out there in the open. One way is to tell as many different stories as you possibly can. Or simpler still: provide as much detail as you can so that people who want to know what is going on simply have too much to read. Documents with plenty of information cannot be read quickly, and the trick is to hide the dirty facts in the main text and keep quiet about them in the summary. Most people only read the "executive summary". They will never notice what the document really says.

Of course this is a known situation. Educators in schools and universities have concluded decades ago that it is no longer adequate to try and give people a comprehensive education. Biology, physics, chemistry, mathematics, they all have progressed and sometimes changed completely so that maybe it makes no sense to teach anyone the complete basics. And new subject areas appeared on the horizon that seemed at one point as important as the classical subjects. Instead, the emphasis was being placed on learning how to get at the facts and synthesize them into a coherent picture. In place of weekly assignments came the projects. Projects allowed to concentrate on a particular narrow area and as a result delve deeper into it. The disadvantage of this method is

[161] James S. Coleman: *The Asymmetric Society*, Syracuse, Syracuse University Press, 1982.

that no matter what project one chooses there always are basic facts that one needs to know in order to understand the literature and be able to estimate how valuable it in fact is. In the absence of fundamental knowledge, a project on climate change, for example, will become a survey of opinions with no possibility of knowing which ones are more plausible. But anyway, since the project is there only to teach *how* to get at the facts, the facts themselves and the lesson we should draw from them become secondary.

This is not to be taken lightly. As I wrote at length in the previous chapter, the capacity of using language effectively depends on it being used in talking and arguing about the subject matter. The more one engages in learning and discussions the better one gets at reading and understanding what is being said. If however projects are simply short lived experiences then they will not lead to a build up of linguistic competence; and since that competence is vital in understanding what is being negotiated, students in the end not only lack the knowledge; they also lack the basic skills to learn it. Unfortunately, this is the state of the art in universities. Time for basic introductory courses is being cut and regular attendance is deemphasized leading to a breakdown in communication.

Unfortunately, the development did not stop here. The internet in its early phase did require at least to know *where* to find the content. To request a document or read about some particular topic one had to know the internet address so that the browser could request it. That proved to be cumbersome and inefficient. With the content growing at an exponential rate it was impossible to keep track of all the sites on which information was available. Not long and people starting thinking about search engines and machines that would collect all these sites on request. Thus, the search engines were being constructed and perfected. Now you can type in a few key words and then the machine will search for sites that provide the best match for your query. These engines became so sophisticated that they could also find sites in different languages pertaining to your topic; or correct spelling mistakes. And they would rank the results so that they could find matches even when not all of your criteria could be satisfied, for example, when they could not find a document containing all your search words. All this taken together with the explosion in the numbers now lead to the problem that one gets hundreds of websites where most of them are pretty irrelevant. The user will have to think of ways to cut down their number.

It is all too easy to blame the state of affairs on an excessive curiosity or an all too diligent administration. Research is producing data and knowledge on a massive scale as well. Here are some examples. When physicists run large experiments, the amount of data produced is staggering. The LHC (Large Hadron Collider) at CERN has a measurement device that produces 1 TB (10^{12} bytes) of data per second, more than can actually be transmitted let

alone by stored.[162] Thus it filters most of that so that the transmitted data is in the magnitude of gigabytes per second. But even that amount of information is too much to be handled by the people at CERN. So, it is actually once again computers that analyze this data until they have extracted something that is fit for human consumption. A second example is the weather. In order to create a prognosis for the next days, institutes need various measuring devices, weather stations, satellites, and so on, and great computer power. And all that is nothing compared to the programs being used to prognosticate the development of our climate, as is done by the IPCC (Intergovernmental Panel on Climate Change). And, finally, consider research on language. Linguists started with corpora, collections of sentences with annotations, often created by hand. This grew over time in size and number. Corpora are now measured in gigabytes (of text). Now that it is possible to archive entire videos, the amount of storage space is growing fast. Research groups routinely create terabytes worth of data.[163]

As these examples show the volume of data is so big that it is not feasible to have humans analyze it. Thus, we once again use machines to deal with the data, which was mostly machine generated to begin with. Similarly, the output of measuring devices is no longer read by humans, but is directly fed into other machines for further processing. Of course, this is a reasonable strategy since most of the detail is only needed to synthesize a picture. For example, the many weather stations and satellites are needed to get an overview of the exact temperature and wind data at as many locations as possible. They will be synthesized eventually into a picture that can be scanned quite quickly and efficiently with one's eyes. The raw data would however be impossible to read or understand.

There is a longterm trend, however, in this. The world is now no longer experienced or understood as it presents itself to our senses. We are not in an immediate relationship with the world; rather, we are increasingly seeing it through devices. While initially this provided a double access to reality in that we used the devices and then checked on them using our own senses, increasingly our perception of the world is provided by machines only. The machines are being left alone to do their job and we trust them unconditionally. We go to walk with an umbrella not because we think that it is going to rain but because a machine tells us so. We turn left with our car

[162] Jürgen Schnack, p.c. Voit: *Not Even Wrong*, writes that each of the two detectors was expected to generate 1,000 TB of data per year.

[163] Our own research cluster generated 80 terabytes over the course of 10 years. The costs of keeping this data online (a requirement by the funding agency) are substantial. I have read estimates that half of the data that is generated in this way by publicly funded research groups is offline because data services are underfunded.

not because this is where we know we need to go but because the machine is telling us to turn left. We are paying the amount at the checkout without knowing whether it is the one we ought to pay or not.

This can sometimes end very tragically. Air France lost an airplane in a severe thunderstorm.[164] After years of research into the causes of this crash the following emerged. The low temperatures had put a layer of ice on the sensors so that they started to malfunction. They declared that the aircraft was sinking. No matter what the pilots did it continued to complain. When they finally realized that they were *not* going down, it was too late. The problem was that the pilots had no way of knowing the instruments were wrong. The were only allowed to see the world through their instruments. And this time the instruments were wrong.

In the 1980ies it was once suggested to show simulations of physical experiments in schools rather than experiments themselves. This in a sense is the logical conclusion of this development. If anyway we trust the machines why ask what is in store for us in reality? Instead we should ask what the machine thinks is in store for us. The only difference being that physics would no longer be called an experimental science. The computer is not running experiments on real objects, it only simulates them. But if it is the machine that deals with reality, there is no need for us to understand it.

This may be far fetched with respect to physics. But the administration is playing that game with us already. It is virtually impossible to double check on an administrative process. We are helpless in telling information and disinformation apart. If you are trying to check your phone bill or your tax return you will quickly understand that transparency is an illusion. If you want to understand how you can save energy or eat healthy food you will quickly realize how much information you need to collect before being able to make a reasonable plan. The manufacturing of light bulbs, computers and even food is so complex that it is next to impossible to understand the effects of one's choices. Hundreds of studies are being published on any given topic, by people of any persuasion that it is impossible to read them all. And that is problematic because a common excuse of politicians against effective regulation has been that it only requires an "enlightened consumer" to stop this. If, for example, people valued the environment they would choose the environmentally friendly products and there would be no need to regulate anything. The only problem, however, is that the enlightened consumer has difficulties telling truth apart from propaganda. It is not possible to understand using your high school education whether genetically modified food is healthy or has any undesirable side effects. Any study that says it is unhealthy is countered by others that say it is safe and vice versa. This is done

[164] Flight 447 on June 1, 2009.

on purpose.[165] The uninitiated will eventually resort to the following strategy: pick individuals or institutions who you trust and listen to them. You do as they say, unless there is massively clear evidence that they are wrong. And all that is not only because the claims are contradictory. It is also because there is no way one could even begin to gather direct evidence by doing one's own experiments. While everyone could have replicated Galilei's experiment by dropping two two canon balls from a nearby tower to find out if it is really the case that they drop equally fast independently of their weight, no such experiment is remotely possible in genetic engineering.[166] As a consequence we are dealing with the matter in the manner of bureaucrats. We collect a few studies, plough through their executive summaries and then vote of the matter.

The explosion of data and data handling machines are corrupting our abilities to deal with the world. And they tend to corrupt the regulatory process itself. I give an example from my own experience. In Germany universities are increasingly asked to compete against each other. The government mandates success indicators and collects the data to compare the success. In the same vein the universities are now making their own departments compete. But it is not easy to compare the success of the department chemistry with the department of education. So, a formula got invented that calculates the success of each department in comparison to similar departments and then compares that relative success measure of the two departments in question. It goes without saying that the primary data is pretty hard to get let alone evaluate. And that no one can do the calculations by hand. In effect, what has been touted as a transparent method for evaluation is all but that. It is transparent only to the degree that one knows what formula is supposedly being applied. But no one can actually check whether that is the case.

This sort of problem appears everywhere now. As Bardi emphasizes, models will not be accepted by humans unless they allow for mind-sized versions.[167] His example is the world model provided in the limits to growth analysis. At the time of the publication of the book the method was quite new. It had been invented by Jay Forrester, the supervisor of Dennis Meadows. However, despite its simplicity the models exhibit sometimes quite counterintuitive behavior and so a lot of effort was spent in understanding why they are right, and how it is that humans can essentially come to terms with the behavior of

[165] Naomi Oreskes – Erik M. Conway: *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*, London, Bloomsbury Press, 2010; and Steven M. Druker: *Altered Genes, Twisted Truth: How the Venture to Genetically Engineer Our Food Has Subverted Science, Corrupted Government, and Systematically Deceived the Public*, Salt Lake City, Clear River Press, 2015.

[166] Given its dangers that is actually good. But still it means that we have to rely on other people.

[167] Ugo Bardi: Mind Sized World Models, *Sustainability* 5 (2013), 896–911.

such complexmodels. Reducing them to "mind size" is a solution, but it suffers from a drawback: we have to trust that the behavior we see in the reduced version is scaling to the big one.

The enormous volume of computations needed to verify claims acts as a smoke screen. It lets us see the world only very sketchily. Of course, when this happens with the LHC at CERN this may not be such a dramatic thing. Whose lives are touched by the question whether or not the Higgs-Boson exists? However, in other cases the mistrust generates deep problems because what is at stake may in some cases even be the lives of people.[168]

[168] See for example *Quand l'austerité tue* by Sanjay Basu and David Stuckler, Le Monde Diplomatique, Octobre 2014.

CHAPTER 6

PASSING ON KNOWLEDGE

Surpassing the Limitations of Space

Knowledge is limited, both in time and space. Not surprisingly therefore a lot of technological inventions were designed to overcome these limitations. The previous chapter dealt with the temporal limits. This chapter will deal with the spatial limits. Of course, these questions are intertwined. For knowledge to be preserved over time we need to maintain the integrity of the artifact that carries it. For a while this can be done by just fixing details here and there, especially when as we explained the information decays gracefully. But as time goes by the entire carrier of the data deteriorates. The paper of the book wears out, the hard drives refuses to work because it can't administratively handle new data, and so on. In the end what we do is copy the data onto another carrier. So, one way of securing the permanence of our knowledge is by transferring it. On the other hand, transfer of data often occurs for entirely independent reasons.

The transfer of knowledge or information requires communicating it. To be communicated it must take a certain physical shape, a material sign that is easily transmitted. Our ideas must be put into words, the words into sound waves, which find their way to another ear where they are converted back into words and, finally, back into the ideas. The distance between two people that talk to each other is limited, though. Just a few dozen meters and the voice of another human is no longer audible. If you want your message to travel further, you need to send a messenger and have him remember everything. It was only with the invention of the telephone that oral communication was allowed to transcend the limitations of the voice. So, much so that now one can speak to any person on the planet almost instantaneously. The only hindrance is the time zone. The internet and its ability to transmit pictures as well was once again a big step but not as big as the telephone itself.

Writing changed the situation dramatically. Even though one still had to use a messenger in order to send the message, now the sender did not have to know the message. The message now received a status of its own, it could be stored, copied, and moved around. Also, writing forced one to think about the nature of the linguistic code. One had to understand that the sentences

were composed from words, that words in turn were sequences of sounds, and that it was more economical to write words on the basis of their sounds than based on their meanings.[169] Subsequent coding schemes such as Morse code are just one further step in the cascade.

The mobility of knowledge was greatly enhanced by an increase in information density. This may be expressed either in terms of bits per square centimeter, bits per cubic centimeter or bits per gram. While until quite recently the rates were rather modest, only a few bits per square centimeter, they have literally exploded in the last decades. This is one side of Moore's Law, which predicts an exponential growth for the information density.[170] Nowadays it is possible to store several gigabytes on a chip the size of a coin. However, the material itself is not mobile, something else must be added and that is the transmission of information. In the beginning the information carrier was moved directly. Later, it got converted into electric signals and then sent through a cable, or transmitted through the air. Actually, fast transmission was invented in the 19th century, long before it was possible to store information electronically. It enabled the invention of telegraphs, telephones, radio and television, before the computer was even invented. The roots for long distance communication had been laid back then.

When we talk about communication, three questions arise: *how fast* can the information travel, *how much* will we be able to transmit, and *how reliable* is the method of transmission? To start with the first question, up until the mid 19th century, the fastest way to send a message was using a horse. So, the speed of a horse was the limit. That allowed messages to travel up to 150 km a day. However, that method was restricted to very important messages and it limited the amount that could be transmitted. In terms of mass communication at that period the medium of choice was the horse carriage. It was substantially slower but could carry much more weight. The telegraph, introduced 1839, changed that overnight. Messages now travelled infinitely fast for human purposes. All that was needed was to connect two points with a cable. Post offices started to be equipped with a big tower. These towers were interconnected and allowed to send messages between them. Furthermore, as early as 1858, intercontinental cables became operational, though only for a month. Initially, it was used only for telegraphs. Phone cables appeared as

[169] See Florian Coulmas: *Writing Systems. An Introduction to Their Linguistic Analysis*, Cambridge, Cambridge University Press, 2002. Suffice it to say that early writing systems were only partly based on sound. The sound based writing systems took a long time to develop. So, culturally speaking this was an immense step to take.

[170] Moore declared in 1965 that the density would double every year, and in 1975 that the density would double every two years. For the forty year period from 1975 to 2015 this meant that density became 1 Million of what it was in 1975, while the period from 1965 to 1975 accounted for another factor of a thousand. The compounded factor is 1 Billion, roughly according to reality.

late as 1940. The technology changed from copper based to optoelectronic, but that affected only the bandwidth, that is, the amount of data that could be transmitted, not its speed.

The next revolution was the wireless transmission, using radio waves. This allowed communication not only between arbitrary points on earth but also into space. It should be noted though that the majority of data traffic on earth still uses cables. Their number has been growing exponentially after 1990. The reason is that the capacity to handle large volume of data needed to be expanded in lockstep with the dramatic increase in computational speed. For it is the change to the computer that now operates the transmission rather than humans that has led to the explosion in data traffic. It has basically eliminated all barriers to communication. The present manuscript can be transmitted to any place on earth in less than a second.

The change from telegraphs to telephones and to video conferences also has changed the nature of communication. It is no longer necessary for a person to travel. It is enough to install a video camera. This is now done on a routine basis, and it can connect any number of people and make them communicate using either video conferences or webinars. Importantly, the information flow is not one way, as in radio and television. It is actually almost a real time conversation, if one discards the inevitable delay of the radio signal. The data transmission rate is obviously much larger than for ordinary phone conversations.[171]

Telecommunication has allowed information to really flow. No need for postal relay stations where you had to change horses. All you need now is computers or smart phones and sever farms. On the way to establishing this infrastructure not only technological problems had to be solved. It was also necessary to understand how the messages could be coded into a format suitable for transmission. As a case in point notice that many different types of codes had to be invented for various purposes, starting from the Morse code to the Unicode of today. It has led to many other benefits on the way.

At this point I pause to consider the development of numbers and number systems. They provide an interesting insight into how practical needs drive the cultural development. When the cattlemen were being sent, to make sure they did not pocket some of the profit they were given a sealed clay box containing a little toy animal for each animal they had to deliver. Additionally, the number of animals was inscribed on top of that box. When they delivered the animals it was an easy matter to compare the live animals with the puppets. In set theoretic terms, a bijection or one-to-one correspondence was being established between the puppets and the live animals, which showed that

[171] The rate is between 128 and 768 kbps (kilobits per second). 128 kilobit run at approximately 16,000 letters since a letter of the English alphabets standardly consumes 8 bits. High quality voice transmission needs 64 kbps.

they were equal in number. Sooner or later it was realized that if the number was written on top of the box there was no need to create the little puppets.[172] This is the origin of the waybill. A waybill is a document that accompanies the freight indicating details concerning the origin, destination, route, and mode of delivery. In short, everything that the sending and receiving party as well as intermediaries such as customs need to know about the material or goods that are being shipped.

The need for this is clear: when the society expands beyond a certain point there is less and less possibility for face to face communication. Hunter gatherers did not need to count, they could simply show their prey to the others. But as soon as the early societies began to expand beyond a tribe, or likewise, as soon as they began storing food for later use they had to establish means of communication across space and time. And they also had to make sure that they were not taken advantage of. For although it would have been enough to just send the cattlemen off with the animals to be delivered, one could not always rely on them. It was better to install a system of control.

These mechanisms that were put in place took a long time to evolve. For any communication across space and time requires taking into account a changed context.[173] For example, reference to time could no longer be made using the words "now", "tomorrow", "later", and so on. For it was not always clear when the message would actually be delivered. So, a system of time synchronization had to be put in place. Some mutually known reference point needed to be established. Otherwise it was impossible to coordinate action. Given that even several thousand years ago the empires spanned hundreds of kilometers, it is inconceivable that no common method of talking time and space existed. Ruling an empire in the Ancient World required coordinated efforts of thousands of people and a complex administration. This administration relied on written records.[174] This among other things was a driver in the establishment of a calendar. Calendars were needed above all for the regulation of farming. At a certain point it was necessary to coordinate the efforts of farmers, especially when they had to be taxed.

Similarly, it was important at any time to ensure that the right quantity and quality of stuff was delivered. This required to have common measures and standards. Whereas time could reliably be measured in days, which requires no imagination of any sort, it was much more complicated to establish a measure, and even more a standardization process. Consider the question of determining how much grain had originally been put into a storage. It was

[172] Peter Damerow – Hans J. Nissen – Robert K. Englund (eds.): *Archaic Bookkeeping: Early Writing and Techniques of Economic Administration in the Ancient Near East*, Chicago, University of Chicago Press, 1993.

[173] Cf. Chapter 4: Language — Context and Interpretation.

[174] Most of the clay tablets of the Middle East contain administrative records.

not enough to know that so many boxes or containers had been filled, as long as it was not clear how big each container was. Uniformity in measurement is a big issue.

Building complex structures, palaces or pyramids required timely coordination of delivery of materials. If the requested materials did not arrive no work could be done. Beniger puts emphasis on the fact that the introduction of the railroad made raw materials travel faster than the messenger.[175] However, the facilities to produce steels were dependent on a constant supply of iron ore and coal at a reasonable price. They desperately needed to be able to send around information faster that the train. This would allow them to place orders ahead of time, and also to do that at a moment when the price was best for them. Otherwise they would be dependent on the producers.

The measurement of time became even more important when traveling by boat. Determining one's place on earth required knowing the time and date, and that could only be achieved using exact clocks. Where Romans and Greeks in addition used mainly sundials or water clocks, and later in the Middle Ages mechanical clocks were being used, they were not of sufficient precision. Enormous awards were being offered to people that would manufacture clocks that would measure the time exactly when in adverse conditions. The British government promised a reward of 20,000 pounds to anyone who could determine the longitude accurately (which required above all, as I said, measuring time as exactly as possible).

Later, with the introduction of the railroads, another problem was being tackled. Technically each longitude has its own daytime due to the rotation of the earth. Until the 19th century each village determined for itself at what moment it was midday. This was done locally, of course. However, running trains across the US required a schedule, and such a schedule could not be organized without synchronization. While in principle this only required a listing of the differences, the calculations were pretty involved. Sooner or later it was agreed to establish time zones, so that time differences would not be calculated in minutes but only in hours. This still is the situation today. It is noticeable that some timezones do not run exactly along the longitudinal line, for administrative reasons.

Interestingly, when Albert Einstein was still working in the Swiss Patent Office in Zurich, one of the concerns they had was synchronization. This was very much the topic that he was working on in relativity theory. Basically, while the work at the patent office required evaluating devices and techniques for synchronizations, he started asking a fundamentally different question: what if the process of synchronization was at heart a physical one? What would happen if there was no cosmic overlord synchronizing the clocks but instead

[175] Beniger: *Control Revolution.*

all physical processes had to communicate to do that? Given that experiments were showing that the speed of light was *constant* no matter how fast the observer moved, Einstein took this to indicate that light was the messenger of the physical universe. Synchronization requited communication. But if communication could not be instantaneous synchronicity would become a rather complex thing. The implications this had for physics where enormous. But we should also draw a lesson for communication here. Synchronization is inextricably linked with communication. In daily life, although we hardly notice that anymore thanks to our electronic infrastructure, communication in turn requires synchronization.

This is evident in GPS systems, for example. They work on the basis of communicating the geostationary satellites. The problem is that the systems have to receive signals from these satellites but in addition they have to synchronize their internal clocks with that of the satellites. This is an interlocking process.[176] In theory radar systems are independent devices for establishing position, but in practice many manufacturers have made them dependent on the GPS system for synchronization.

To close I shall discuss the last of the three questions: how reliable is communication? In passing, I have noted that sending messages makes them vulnerable to manipulation. When the message is not delivered by the same person it is anyone's guess what actual message is being delivered. But apart from willful manipulation there is also the accidental manipulation. The message can get altered or corrupted. Again the most simple case is that of sending someone to deliver a message. You can make the messenger learn the message by heart so that you fix the exact wording. But you cannot guarantee that he remembers everything well enough. Thirdly, and very interestingly, there are occasions when you want to make sure that only certain parties are able to read and understand the message. While one method is making sure the message is not available physically, when this is not feasible one needs to think about encryption. Especially in wartime this is the method of choice to ensure that only trusted parties can read the message. Caesar was said to have used a method of encryption, known today as the Caesar. It involves rotating the letters by a fixed amount. However, also in business and financial transactions, encryption is used a lot. This ensures that the communication is not tampered with, and that it is only between trusted parties. Interestingly, the latter is a very important point that is easy to overlook. With modern communication involving dozens of real stations that route or resend the signals it is vital to guarantee that the end-to-end communication happens with the right person. One way of

[176] Synchronization of machines is a highly nontrivial affair, see Donald MacKenzie: *Mechanising Proof. Computing, Risk and Trust*, Cambridge (Mass.), MIT Press, 2001, for an enlightening historical account.

ensuring this is to use protocols and encryption. For if an untrusted party receives the message it will not be able to read it as long as it does not posses the means of decryption.

The establishing of trust has been a very important topic in accounting. In the modern era this is achieved by a method that is called blockchain technique. Basically, what it ensures is a distributed form of storing data (no central authority). Different storage points are agreeing to store data. Each time a move has been agreed on, it is simultaneously implement on all of them. It uses an encryption mechanism based on the previous entries. The encryption is used to make sure that no entry is being changed retroactively.

Networks

Knowledge and information are generally distributed. Nobody knows everything, no library has all books, and no server has all data stored locally. It is neither possible nor does it make sense to have all knowledge simultaneously everywhere. It is impossible because the amount of data is so large that it would exceed the capacity of any storage device. Even if we tried to store everything (as Google once did) it inevitably leads to the same problem: the data must be distributed over several devices. Storing the data at one location is also not always desired. Consider all the private data stored away in your files. Surely you do not want all of it to be stored elsewhere even if you are being assured that no one else can get access. No one really knows what happens to data stored on other locations. Physical control is one way of keeping access limited.

Sometimes however the opposite also is desired. Data is subject to decay and manipulation. Thus, it risks to lose its integrity. In order to minimize the risk of this often the same data set is stored in different physical locations. Nature does that: the DNA is stored basically in every one of the cells. All major institutions do that. They mirror the data in one location in another location. In order to do that, they have to transmit it.

This leads us to the notion of a network. In the most simple form a network is collection of locations, called nodes, together with communication links, called connections. The connections will for simplicity be thought of as involving two nodes, and being bidirectional. Information in a connection can thus be sent from one end to the other, and back. An example is a telephone network. Each phone is a location, and the locations are linked using wires. Now even if nowadays there are no wires, the notion of a network is still useful. For the phones do still function as if there were a wired connection. The message is being addressed like a letter and then sent to the location with

the address, and nowhere else. Now think of each phone as actually being a person or some machine holding data. In either form the network can be used to exchange information between the nodes.

There are other kinds of networks. Radio signals are from one source to any number of receivers. And they are typically one way only. But all this detail is not necessary for understanding the effect of networks in storing knowledge. Consider two nodes, A and B and a single connection c between them. Assume that the knowledge that A has is $W(A)$, and that the knowledge that B has is $W(B)$. We may think of $W(A)$ and $W(B)$ as sets of propositions. For example, we may have $W(A) = \{P\}$, which means that $W(A)$ contains (only) the proposition P. Now introduce another symbol, $W(A) \vdash P$, to say that *A knows that P*. Knowledge comes in two flavors. We may know a proposition because we know of it immediately, for example, in this case A knows P because P is a member of $W(A)$. This is called *explicit knowledge*. However, we may also know something because it follows from something that we know explicitly. This is called *implicit knowledge*. Implicit knowledge requires deliberation, or reasoning. If, as in this example, $W(A) = \{P\}$, then A knows $P \vee Q$ ("P or Q"), though only implicitly. For on the one hand it is not in $W(A)$, so A does not know it explicitly, on the other hand, if P is the case so is $P \vee Q$, so A can deduce $P \vee Q$ from what he knows explicitly.

Now consider the case that $W(B) = \{P \to Q\}$, that is B knows (explicitly) that if P then also Q. If A and B have the option of using the channel c to send around queries and get answers to their questions, A can send a request to B asking whether $P \to Q$ is the case. B will actually confirm that, upon which A now knows both P and $P \to Q$. And so A can conclude Q. We now have the following situation. Neither A nor B know explicitly nor implicitly that Q. So, symbolically it is $W(A) \nvdash Q$ and $W(B) \nvdash Q$. However, $W(A) \cup W(B) = \{P, P \to Q\}$, and so A and B *together* explicitly know P and $P \to Q$; and they together implicitly know Q. This is a common scenario, and it is useful to elaborate on it. First, for a group of individuals we may define the explicit knowledge as the set union of the individual explicit knowledge of its members. If the group consists of A and B we put $W(\{A, B\}) := W(A) \cup W(B)$. So, the group knows φ explicitly if any one of its members knows φ explicitly. The implicit knowledge of the group consists of all the logical consequences of its explicit consequences. In this way we have defined explicit and implicit knowledge of a group. The implicit knowledge of a group is different from what in the literature is called *common knowledge*. P is common knowledge in a group if every member knows it (implicitly) and knows that every member knows it and knows that every member knows it and so on. In a sense, common knowledge is the common shared knowledge (a kind of lower denominator), while group knowledge is the joint knowledge of all members, pooled together.

The crucial difference between a group and an individual is that the pooling of knowledge comes at a price. Suppose we ask A in this example whether or not Q is the case. Then in order to establish the answer A would have to ask B some question. That in turn requires communication through the network. Thus, the way in which the answer is established is by using several rounds of discussion, where members are being asked questions and either answer them straight away or initiate a new round of questions and answers. Here is an example. Take three individuals, A, B, and C. Assume $W(A) = \{P\}$, $W(B) = \{P \to R\}$, $W(C) = \{R \to Q\}$. Suppose A wants to know whether Q. All members are linked by a connection. So, A asks B whether or not $P \to Q$. B does not know the answer. Instead, B decides to ask C whether or not $R \to Q$. C will confirm, leading B to conclude that $P \to Q$. So, B will now confirm to A that $P \to Q$, and A finally has an answer. This time the protocol involves several questions and answers. The request by A initiated a request by B, which was answered, and only after that the answer to A was established.

These rounds of questions and answers can get arbitrarily long. As is immediate from these examples, there are two obvious problems. First, when we do not know the answer to a question it may not help if we simply passed on the question to some other person. If A in both of these examples had just asked whether or not Q, there would have been no answer. So, A had to come up with another question to ask. In this case the question was whether given what he knew (P) one could determine Q. So, there is some ingenuity involved in asking the question. Second, it is important to know whom to ask. If we ask the wrong person we may not get an answer. In the second example, B did not know the answer. The procedure could have stopped right here, and often enough in real life it does. Only if the other party is helpful and sufficiently ingenious it may initiate a further round of the game to establish the answer. This happens quite frequently. You enter an administration building wanting to solve a problem. You do not know who could help you with that. You decide to ask at some desk. Then depending on circumstances that person at the desk will tell you who to direct your question to, pick up the phone or shout to other end to see if the person there can help you with your problem. Or they might simply declare that they are unable to help.

Let's look at the second problem in more detail. It concerns two separate questions. First, what is the network like? What are the nodes and what are the connections? Second, what do the individual nodes know and how will it solve my problem? These questions may sound trivial in the modern age, but even today they are quite relevant. First, concerning the layout of the network. Not only are there many people that can answer a question that we may have, these people live in different places, and it may not be such an easy affair to get hold of them. If you want to get help with your tax return, you need to know who the people are that can help you and then find out how best to reach

them. Before the phone was invented one had to go there in person; now it is either the phone or the internet which we can use to establish a connection. However, there is an additional problem and that is the problem of who has what knowledge. Depending on the nature of your tax problems you want a specific person to handle it, so you also need to know whether they actually *can* help you.

An anecdote will highlight this. A few centuries ago, a certain Antonio Magliabecchi was famous throughout Europe for the fact that he knew about all libraries what books they contained and where. This person never travelled and hardly ever left his room. People visited him from all over Europe for two reasons: one was to update his knowledge, and the other was to ask him about books on certain topics. In a time where information travelled on horseback, it took a lot of effort to get information. Obviously, before one even considered to visit a library to get hold of some rare book it was a good idea to run the question by Magliabecchi to see if it did have that book, or if in fact there was a closer library.

As libraries got larger, it was a question of time until even a Magliabecchi would be unable to get all the necessary information. Instead, libraries started to print catalogues of books and sent them around to other libraries. So, by visiting any one library you could get to know what books the other libraries possessed, and so you could decide which of them was worth a visit.

Now let us look at a somewhat different type of network, a Boolean circuit. A Boolean circuit has nodes and connections, but there is a difference. First, information goes only one way, not both. And second, per time unit there are only two possible signals: 0 and 1. The nodes, or gates, compute very basic functions. Consider an XOR-gate. It has two ingoing connections, i_1 and i_2, and one outgoing connection, o. If at time t the input signals are different, then o will send out a 1 at the next time; otherwise it will send out a 0. This function, called exclusive or, is nothing but the addition of two binary digits: 0+0=0, 0+1=1+0=1, and 1+1=0. (In the latter case there will also be a 1 added to the next higher bit, but that does not concern us here.) There are more types of gates, but it is possible to construct every Boolean function from just these XOR-gates. For example, to compute the negation, the function that sends 1 to 0 and 0 to 1, all you have to do is send the the same signal into both i_1 and i_2, which outputs 0, no matter what the input is. Add one XOR-gate, and feed the signal of i_1 and 0, which will give the desired result.

Now let's focus on the gate. They are agents that operate on a fixed schedule. They are not supposed to choose or do anything other than what is written in the description of the gate. If your gate produces 1 from inputs 1 and 1, it is no longer an XOR gate. So, the individual gates are pretty dumb; all they know is how to add two binary digits, according to the definition we provided in

Chapter 2: Theory — What Is Knowledge? However, they can be connected to form a network of any complexity whatsoever. The complexity resides almost entirely in the design of the network.

Neural networks have been pictured in a similar fashion. The have nodes, the individual cells, and various connections, ingoing and outgoing, which compute a rather simple function. Signals are 1 or 0, or "fire" and "do not fire". One of these functions is a threshold function: if the number of ingoing connections that fire is larger than a given number, then fire out to all connections, otherwise do not fire.[177] This somewhat generalizes the Boolean circuit, so it allows to calculate any boolean function, and hence any arithmetical function.

However, it was quickly realized that either the brain had a layout fixed from birth or the brain is capable of changing the structure in response to input. The first of these options would lead to a fixed brain that would not learn anything. The second necessitated a change in the definition of neurally computable functions. To allow for smooth learning it was assumed that each connection was given a weight, a real number. The neural cell would then compute the weighted sum of the incoming connections. Furthermore, the threshold value was variable as well. For example, consider ingoing connections i_1, i_2 and i_3 to node n, with weights 0.7, 0.4 and 1. Let the threshold be 1.5. Then if both i_1 and i_3 fire, the value is $0.7 + 1 = 1.7$, which is above the threshold, and so the cell fires. If only i_1 and i_2 fire, the sum is 1.1, and so the cell does not fire.[178]

Given a fixed structure of nodes and connections the network is nevertheless adaptive and its power to learn connections are immense. We can take this to a higher level and allow the network to restructure itself, as the brain actually does. In the first nine months it actually builds the entire network, and it does that on the basis of the input it is receiving. Major learning algorithms in computer science are based on variations of this theme.

The point of the exercise is twofold. On the one hand it shows that complex behavior can be broken down to very basic patterns. Second, the architecture of the network can in some cases be all there is that determines the behavior.

There is a deep lesson in this. We may apply the analysis of intelligent action at any level. We may decide to look at a single brain and treat it like a giant network with little neural cells. Or we may see it as a network of a few dozen autonomous units, specializing on different tasks (language, vision, and so on). Or we may think of each human as a single node in a network the is the society or a big company.

[177] Warren S. McCulloch — Walter Pitts: A Logical Calculus of the Ideas Immanent in Nervous Activity, *Bulletin of Mathematical Biophysics,* Vol. 5 (1943), 115–133.

[178] This in an artificial example. The different versions of neural networks are legion.

It is the same reality that we are describing, but the granularity of the resolution is different. The parts may also have more or less complex functionality depending on their complexity. Also, depending on what the network is like it acquires a different agency. While we do want to say that humans have intentions, it is not clear that we wish to say the same of the individual neurons. It seems that the neurons are pretty deterministic in their behavior. The input determines the output, there is no choice. However, we can imagine that each neuron is something of a small bureaucrat having learned how to act for the purpose of the network, with action schedules being changed every now and then. Thus, bureaucrat can be charged with adapting the weights, so there is an element of learning that he can do for himself. However, this bureaucrat will never come to understand the higher purpose of all this. It is never revealed what the masterplan of this structure is.[179]

On the other hand, the network does not seem to know much either. It shows a composite response pattern without knowing how it got there. Learning simply happened. In fact, it is still a mystery what consciousness is.[180] But this may be true only insofar as we require the network of its components to make sense of its own behavior. We can leave that to us humans to do that. And we do that actually all the time, attributing higher order or even intelligent behavior to complex systems.

Now apply this to complex social organisms and the situation is again different. With a well-defined control structure we can to some degree understand what the system is doing. When the manager of a company tells us about his plans and decisions, we can get some indication of where the company is headed. But many societies are self-organizing. A university has no fixed person in command, even most companies are not in private ownership. This actually means that the company assumes a life of its own that is independent of the people that work in it. This is because everyone is entangled in a web of expectations and payoffs, with rules that tell them what to do and what not to do. Their freedom is limited for the whole to gain predictability. If the job is to build a skyscraper, the orchestration needed to complete this job is mind-blowing. And it requires virtually all people to work according to schedule.

In these circumstances most people actually are aware of the higher goal that is put in front of them as a collective. And they know how that percolates

[179] This is reminiscent of the Chinese Room argument put forward by Searle. It serves to debunk the idea that the bureaucrat can actually know what the entire system is up to. Specifically, if it acquires the ability to correctly combine pairs of Chinese characters, this does not make him know any Chinese grammar.

[180] See Horgan: *End of Science*, Chapter 7, for competing views by the experts. See also Tetens: *Gott denken*.

down to the individuals, themselves, and they can partially understand why it is that they have to perform a particular job. But for the whole to function, that knowledge is not always necessary.

In many other situations however there is little or no insight into the motivation. And there is therefore less predictability. Psychoanalysts have firmly established that humans can deceive themselves about what they really want. How that is possible in terms of intentionality is a subtle matter. If part of me wants to be a musician, or go on a world tour, while other parts want something different, what is it that *I* want? And how exactly can some part of me want something? Maybe in this way that if we deny that which it wants then we will suffer in the long run. So, there is no consciousness involved, only a subject that watches itself over time to discover that which makes it thrive. This is a point Peterson stresses.[181] There is according to him typically no knowing what we want or how we behave other than by acting it out, or absent that possibility, by discussion or speculation. Fictional stories are good at helping to sort out our mental frame.

Societies are no different. If we want to award some intentionality to society we have the task of identifying that which it actually wants. But how can we find out? Somehow this is even harder because there is no obvious person to ask. Maybe the solution is just the same: we let it evolve in time and study its evolution. This is why we study history, for example. It may not be clear what lessons we can draw, but the interest in understanding is certainly there. Maybe there is no will or want in a society, but there certainly are patterns of behavior, because at some lower level of resolution there are groups or individuals that have an expressed will for something to happen or be achieved.

Finally, there is a point to note about knowledge. Consider the efforts of achieving scientific excellence. This is not a matter of entrusting the advancement on some bright individuals. There is an entire infrastructure of institutions that needs to be developed in order to achieve this. However, it is not always clear what exactly allowed success. Was it the determination of some gifted people? Was is the quality of the education system as a whole? Maybe there is no final answer. However, as societies evolve they always run a risk of undermining that which brought the success, partly because they do not know that this is what caused it.

Consider the cornerstone of a democratic society, the freedom to choose one's profession. In Western Europe this is translated as providing places for students according to what they want to be in later life. Whatever the rationale is for that it risks to generate not enough skilled people in one discipline and too many in another, seen from the society's overall perspective, or, as it is typically phrased, according to the market. This always threatens the market

[181] Peterson: *Maps of Meaning.*

balance. The reason is clear: there is no way for individuals to effectively guess what their real chances are, they are running the risk of making less than optimal decisions. In theory, the fluctuations will be ironed out by the market, but in practice I am not so sure that this is realistic. The decline of attractiveness of sciences is not a short term trend and will probably not be stopped by offering higher salaries. By the time we understand what drives that mechanism, it may be too late.

Generations

As the saying goes, every time someone dies an entire library is dying as well. The traces of knowledge will be wiped out forever, at least from our modern point of view. In Antiquity what mattered was the fame that one had acquired in life, and that provided some consolation. As Horace put it: I have built a monument stronger that iron. He was right. For even if he had to die, his manuscripts remained, and people still liked to read them. So, much so that they copied them when they started to show signs of decay. Who would not like to be that kind of person whose fame survives his body for generations to come? After all his works seemed to be so valuable that they put their efforts into preserving it.

And so what is interesting about this is that the treasure of knowledge needs attention. We cannot leave it to chance. We must see to it that it does not leak. For we are in constant danger of losing it, either because people forget, or because they die. Or because the things that store the knowledge for us decay. Whether it's Google throwing out hard drives or monks copying manuscripts. Given that copying a manuscript by hand is a time consuming process, and that each copy had to be made individually, we owe the monks enormous gratitude for the service they provided. For we are speaking of an enormous stubbornness with which they were sitting in the scriptoria day in day out making sure nothing of value gets lost. You just need to consider the thousands and thousands of manuscripts that we have today, most of them Medieval copies of originals that were long lost. Today, it is the libraries and printing presses that are continuing this effort. And the museums who see it as part of their duty to preserve the cultural heritage.

But what is the use of all this effort if there is no one that will actually read the books? Of course, when you preserve some manuscript you are not in the first instance thinking about that question. You could be excused for thinking that the book will find its readers when there is a need for it. But that assumes that they *can* actually both read and understand it when the need arises. We have seen that this is not guaranteed. Languages change over time, and the manuscripts are without the context that gave rise to their existence.

If you read texts from Roman or Greek writers they are at times not really clear without a sound knowledge of what was going on at the time. This is why critical editions are necessary. They try to reconstruct that context for us, apart from trying to recover the original text as well. All this however is only possible if enough other manuscripts are preserved that allow to reconstruct the situation of the writer. Often, when all we have is just a few texts in a language (Etruscan or Gaulic are cases in point) it becomes next to impossible to understand what they are saying. This same problem may also arise in the future. A dramatic case is provided by the so-called nuclear semiotics. In 1981 the US Department of Energy together with the Bechtel Group created the Human Interference Task Force to investigate how one could effectively prevent future generations to inadvertently meddling with dangerous nuclear waste. Since the waste would retain its deadly radiation for thousands of years this was a pretty daunting task and no promising solution given the enormous time span has been found let alone implemented. Basically, apart from relying on human nature it needs to rely on passing on a fair bit of knowledge over dozens of generations.[182]

Modern societies have created a number of institutions to make sure knowlegde is passed on to the next generation(s). These include schools, universities, but also libraries, museums and recreational centers with an educational mission, for example national parks where one can not only spend time in nature but also learn a lot from the exhibits. These institutions are important. While we can expect that virtually all children learn to speak, all other knowledge seems to be dependent on various factors. But surely language skills are just a fraction of what humans need to function well in a society; and it is likewise a fraction of what society needs individuals to know. At school they therefore get taught mathematics, reading, writing, a foreign language or two and the basic of natural sciences, physics, chemistry and biology. Besides that there is music education and some general introduction into the culture. A century or so ago most of what children learnt they were taught by their parents, relatives and close neighbours. And still today school leaves a lot out of the curriculum that is nevertheless vital: cooking, taking care of one's health, dealing with children and other people, learning to drive, managing one's finances, doing the tax return and more. There is a growing trend to make school take care of all these issues as well, and some have responded to that, but it is not possible for school to take of one's life unless we are talking about boarding schools. This is why there are many other institutions, from sports clubs to libraries, that exist side by side with schools.

[182] When I look at another project, the so-called Doomsday Vault, it does confirm my overall skepticism. Basically, the wheather got so warm that the permafrost around the Vault is melting, which should not happen. It seems then that we are not able to build something strong enough to last for a few decades.

After school, learning continues. However, increasingly people are specialising in what they learn. University education does emphasize a need for universal learning, but it nevertheless asks students to choose some subjects as a specialty. When they go on to do a PhD what they are dealing with is not even an entire discipline, only a fraction thereof. This is necessary as the body of knowledge is so large that one can only know a very small fraction. However, it is necessary to know everything around a certain issue where one wants to make progress. Fundamentally, progress in science cannot be had without mastery of much of the knowledge that exists. This is more than just an issue about whether progress is predicated on knowing the blind alleys or earlier attempts. It also about knowing whether at all we are making progress, and for that we need to know what we have known so far.

The education sector has seen a dramatic expansion in the last hundred years. Back then, around 1900, half of the population was working in the fields. Today that figure is lower than 2 %. The mechanization of the agricultural sector has produced a dramatic shift in the societal structure. In Germany, with a population of 64 Million there were some 55,000 students enrolled in the universities. Today, with a population of 82.5 Million (2016) the number of students is 2.8 Million (2017).[183] 55,000 is less than 1 permill of 64 Million, 2.8 Million is more than 3 percent of 82.5 Million. In other words, there is a thirtyfold increase in the relative number of students. The same goes for schools. The proportion of children going to schools that prepare for entry to universities has been steadily increasing. This is a self-enforcing trend. More education breeds more education. People leaving universities are looking for employment that matches their education. And, at least until recently, technological progress meant that such jobs were indeed on the rise. This trend does have its limits, though. The world is changing and the perception of the world is likewise changing. Not everything that science is digging out stands the test of time. This is the case with elementary particles, genetics, as well as quite mundane truths. As a child I was told that spinach contains iron and that margarine is healthier than butter and that Caligula was a deranged leader of Ancient Rome. A decade later I learned that none of this is true. Apparently no study was able to find higher concentrations of iron in spinach the first study had proclaimed; the benefits of margarine were now claimed to have never been established and were just claimed to make people buy it; and Caligula got his bad name because certain authors, among which Suetonius, simply disliked him.[184] Each generation thus has its own world view and thinks that earlier generations were simply to eager to believe certain stories. Of course, the next generation will think the same about this

[183] Numbers taken from http://www.destatis.de.
[184] Alois Winterling: *Caligula: Eine Biographie,* München, Pantheon Verlag, 2007.

one. In view of this it is not a good idea to just get rid of all earlier science books once a new theory is established. Nobody knows how long that new theory will be around, and which theories that have once been proposed will at a later point be revived. Thus, whether we like it or not, progress in science does not mean that we can dispense with earlier knowledge.

Indeed, in view of the fact that scientific progress is temporary it is not wise to implement a winner-takes-it-all approach. Instead, Finke advises an ecological approach to knowledge in contrast to monoculture.[185] Knowledge, Finke claims, is indeed de facto an ecosystem, and therefore should be treated like one. Rather than placing one's bet on just one theory one should let all of them grow until a consensus is being reached. As an example Voit cites the enthusiasm of physicists over string theory and the influence of big names like Witten who discouraged students from studying other theories.[186] It did not help much, by the way. When a new theory loses traction we would like to be able to return to the other ones we did not pursue for a while. Just consider this case. After World War II many farmers started adopting the new chemical approach to farming, with fertilizers, pesticides, and so on, to the dismay of their fathers. Now increasingly I am hearing more and more stories of the reverse happening: farmers are turning to ecological farming. And their parents are shaking their heads in disapproval.

Passing on knowledge to the next generation is an incredibly complex affair. It is a social phenomenon and does not function like a network of data connections. The human brain is not a hard drive. We do not learn by storing facts. The mechanism by which we learn is much different. First, it is not simply a matter of being exposed to something only once. Learning requires repetition.[187] To master the forms of the French verb you have to invest a substantial amount of time. Not only because there are many of them that are exceptional here and there. But also because it is not a matter of simply knowing what the forms are. You also need to be able to produce them on the spot. The necessary fluency requires substantial effort. Repetition is key. The idea of repetition has somehow been in disrepute for a long time. Somehow it was thought that motivation is enough, but it can only be the catalyst. Most people do not learn instantaneously. This is no surprise if you think about what the brain is actually doing. Consider learning a bunch of phone numbers of your friends by heart. Ordinarily, one would think it is just a matter of memorizing some numbers and names. But the end result is something completely different. Not only are you able to say which friend has which phone number, you may also be able say which phone number belongs

[185] Peter Finke: *Ökologie des Wissens. Exkursionen in eine gefährdete Landschaft*, Freiburg im Breisgau, K. Alber, 2005.

[186] Voit: *Not Even Wrong*.

[187] This is because of the continuous and slow self adaptation that the neural network performs.

to which friend. And you can do that even if you are just given a fraction of the number, or some part of the name, or even just an indication where that person lives.

The brain is not a set of data. It learns by associating the data, potentially linking anything with anything else. This is why you can have memories of a person or a place when hearing some music, why seeing a particular book can bring up memories of something seemingly unrelated, why you can establish relationships between words of different languages that seemingly aren't there, and it is why you are able to learn things that you have never been told. The brain is trying everything it can to get the most of our experiences. And this is what makes learning so time consuming. To establish these connections is not a one time affair. They are established slowly, because they need to be corrected when necessary because not every observed connection is in the things, it might just be a one time affair.

Thus learning is hard work. For modern societies this is measured in money. It hires teachers, day care workers, university professors and many more people in order to guarantee that everybody is learning as best as they can. In addition, rules are being set up for schooling and education. Most countries have made it mandatory to attend school for a number of years, and have made it free of charge in return. Teacher are told what to teach, and often enough also how to teach. They have to issue reports and consult with the parents to make sure rules are being kept. Universities are regulated as well. They are given the right to award certain degrees, they get money for instruction and research. Increasingly, however, governments are not automatically doling out money but are asking schools and universities to have students be tested, to compete with each other, write proposals, and provide detailed figures on request. This can be seen as an effort to guarantee a standardized education.

These standardization efforts are in my opinion a remedy for a shift in our society that might have severe consequences. Recall that the attitude has been changing some fifty years ago. Knowing facts was considered a rather mindless business. Today, educators are asked to separate knowledge from competence, and increasingly it is the competence that is being considered as the only thing of value that should be taught.[188] This is a questionable value in itself for much of what we do requires a lot of factual knowledge. Speaking a language is a case in point. The competence of speaking French consists in knowing a lot of words and morphological tables. Underlying, however, the idea has been that the lack of knowledge really is driven a lack of desire to

[188] I can attest that at the university where I teach this is in full force. The lectures or modules must be accompanied by a separate description of the factual knowledge and competences that are being taught. If a student wishes to transfer credits to another module administrators are instructed *only* to look at the competences.

learn, and so it is the teachers and professors who, rather than insisting that certain things have to be learned, are being asked to instill the motivation in their students so that they will want to learn. Some countries already discipline professors if some students are being failed.

But why is this necessary? In a short essay, Stiegler reflects on the ills of education.[189] His ire was sparked by a French advertisement for a television channel. The picture shows a young man, a middle aged man, supposedly his father, and an old man, his grand-father. It reads: "He deserves better than that." The message is: television is better for the youth than the (boring) adults. Nothing could be further from the truth. Stiegler explains this as follows. What we value as human beings most is relationship. All the wealth we are accumulating is worth nothing if we are alone. The value system we have does not appear magically. Look no further than psychology. Freud claimed that children learn from their parents what it means to be a human being. They learn how to behave and what to value. They learn what other people value in them that makes them good company. This is what Jordan Peterson is getting at when he puts up the rule: "Do not let your children do things that make you dislike them."[190] Because the children *want* to be told what is right or wrong, and the task of the parents is to show them what makes them being liked by others so that they have some company.

This now provides two lessons. One is that the idea that school can basically teach anything that needs to be learnt in life is simply wrong. Parents are necessary whether they want to do the job or not. Second, the motivation of children and people in general is coming from the older generations, parents and grand parents alike. And no one can take their place. That is why the idea that certain new media, be it television or social media, computers or books, can somehow produce motivation is an illusion. In actual fact they do the opposite. They take away time from the parents to teach the children the most basic things in life.

In that respect our modern culture is at a cross road. The increasing tendency to let other people educate the children, often in the name of providing equal opportunities, may be responsible for the lack of motivation in them. To make teachers and professors responsible for motivating them is to ask for the impossible. Now add to this an increasing disorientation of the society at large about what values it ought to have, and the motivation for learning evaporates.

[189] Stiegler: *Logik der Sorge.*
[190] Peterson: *12 Rules.*

CHAPTER 7

THE USE OF KNOWLEDGE

Cost and Benefit

Knowledge is necessary. We need it to survive, literally. In the old days, you had to know which plants you could eat, how to avoid predators, how not to get killed by members of the other tribes, and many more pretty essential skills. Today the body of knowledge needed for survival is different. You need to know the dangers of certain chemicals, of electrical current, how not to drown in water;[191] and depending on where you live there are certain obvious things one needs to learn in order to not get into trouble.

To learn, you have to expose yourself to risks.[192] It is not possible to learn to climb a tree without also risking injury. Similarly, for riding a bike, climbing a mountain, doing a chemical experiment, and so on. The safety culture around children has had a detrimental effect on their ability to learn basic skills. The risk is always taking a toll on someone. Also, it invites reckless behavior later in life. If you have not been exposed to danger you can't understand it. Some knowledge is bought with a big sacrifice. I am imagining how early humans found out what plants are edible and which ones are not, and how they inadvertently stumbled over all kinds of plants with deadly or hallucinogenic effects.

At a higher level the scenario repeats. Science has made Western culture triumph in a big way. If you want to know the benefits of knowledge, it is best to start with an encyclopedia of the 19th century. There you can see what experiments were being performed, what the reasons were for performing them and how they helped mankind. The list is truly impressive, and my intention is not to diminish the successes. I am mentioning two areas where the impact has been beyond imagination. The first is health. The discoveries of the 19th century stand out in this respect. The importance of hygiene was discovered, the source of many diseases, and last but not least also the narcotics. Take away these three and you can see how revolutionary they

[191] I read somewhere that in Germany every year some 400 people die while swimming. Although people enjoy it many seem unprepared for the dangers that it offers.

[192] This is the theme of Jordan B. Peterson: *Maps of Meaning*, as well as the somewhat more accessible Peterson: *12 Rules*.

were. There would be almost no surgery, and most diseases that used to kill people would still be there. The modern successes appear rather modest in comparison.

The second area is transportation. The invention of the diesel engine and all that came out of it was truly a revolution. It allowed to plough the fields without horses, and freed many people from the drudgery of field work. Anyone hailing the next revolution in the form of flying cars should take note of how labour intensive farming used to be before the introduction of a tractor.[193] Similarly the work in the iron and steel factories. To get a notion of how the conditions were like in the US around 1900 you just have to look at the conditions of migrant workers in China or workers Bangladesh today.

But all the successes should not make us complacent. Literally everything has two sides, and progress is no exception. Knowledge is used not only to do good. Inventions, even if initially made for a good purpose, can be used to harm people. We all know that. However, there is a deeper problem that I'd like to address here. It comes in two forms. The first is that discoveries have side effects, and these can be pretty severe. Typically, there is no way to avoid them, one can only mitigate them. Some of them can be known in advance, others only emerge on the way. The second problem is that the advancement of science in and of itself can be problematic in that it produces hazardous materials or other unintended consequences. The difference between the two is perhaps one of magnitude. Inventing new technologies does not interfere with nature or create much waste. But when the technology is finally thrown onto the market this is generally done on a massive scale, and that will cause problems. As the saying goes, the dose makes the medicine.

A problem of the latter kind is exemplified by the car. In and of itself the car is a wonderful invention, and for a long time it really seemed that it would be able to deliver easy transportation to virtually everyone. Probably very few people imagined just how bad that would be. The space required for cars is enormous, not to mention the bad air and the road accidents. The list of negative effects is growing longer and longer. With transportation eating a third of the energy budget it is also one big source of greenhouse gas emissions.

When talking about side effects, a serious problem of modern societies is drug addiction. What is not commonly known is that—as their name actually suggests—most drugs originated in the laboratories of drug companies. At a certain time they were even completely legal, such as heroin or (met) amphetamines. The progress that allows to treat mostly pain is bought at a price that is increasingly difficult to ignore. A spectacular case in point

[193] For a vivid account see Robert J. Gordon: *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War*, New Jersey, University of Princeton Press, 2016.

is the opioid crisis. It has its origin in a rather successful drug against very strong pain that was being pushed rather heavily by the manufacturers. It ended up being so popular that when doctors ended the prescription people turned to "regular" drugs from the local dealer. And that typically sparked off a downward spiral.

It is hard to shrug this off as a case of greed gone too far. The link between drug making, drug use and drug abuse is so obvious that it is hard to imagine that one could have one without the other. Just how much abuse may be the question, but no abuse is impossible. No law on earth has managed to eradicate adultery or murder, why should drug abuse be any different?

So the story really is a different one. On balance, progress has made the lives of people better all over the world. This has allowed to ignore the side effects. The dominant story was that of the hero who would fight the superstition of the locals and eventually triumph over their skepticism because he had science on his side.[194] Indeed, science triumphed in the West like nothing did. And when it came to asking for funding, it was the natural sciences that would claim the lion's share because they had something to show. Growth in GDP was not achieved through editing the complete works of Brentano or Hegel, but through working out new products and better ways of mass production. The account of Gordon in showing what mattered in terms of growth is telling.[195] Arts and humanities do not figure there apart from film and mass media.

On the other hand, freeing people from hard labour also released a lot of potential for doing other things. New scientific fields were being discovered or worked out that initially held no promise of growing the economy: psychology, sociology, and anthropology are cases in point. They did advance our understanding of the human nature, but the economic potential was modest except perhaps in exploiting psychological knowledge in advertising.

However, the tide is slowly turning. The reason for that is quite prosaic. If energy supply is not growing especially if energy per capita remains constant or is declining, then engineers are experiencing problems in justifying their work. Sure, for the moment it is thought that we need even more of them since they are the only ones being able to make machines more efficient. But the potential for economizing in this way is rather modest. It is estimated at a mere 1 % percent a year. The rest will have to come from cutting real services. And that means reducing the need for machines.

[194] Just read the entertaining story "A Connecticut Yankee in Sir Arthur's Court" by Mark Twain.

[195] Gordon: *Six Headwinds*.

The contrast could not be starker. While the next revolution is now at our door steps threatening to eliminate half of the jobs,[196] the evidence is mounting that this revolution will need energy that we will not have in the future. It is unclear how the game will play out. Although potentially we could eliminate all physical labour, for various reasons I think this is not going to happen. Even though one could argue that the amount of extra energy needed to pull this off is modest in comparison to what we are already consuming, it is not clear where this additional share of energy will come from. Something else has to yield it. Second, it is not clear to me what the psychological effects will be of a production system that needs no physical labour. Humans are made for doing manual labour so they will have to find a meaningful outlet for this energy.

What is slowly taking shape is what I call the second reading. The idea is that when you read a book for the second time it may be a totally different book. Or when you hear a piece of music again and again it will sometimes take on a totally different meaning. Similarly, when we look at our own civilization. When the first machines got introduced, the enthusiasm was overwhelming. The era of Enlightenment had begun and promised to liberate people from the self inflicted slavery that was caused by their ignorance. Not that progress has always been unanimously welcomed. Far from it. Machines made people redundant, and the pollution was already then obvious. But as with the church, the grip of the narrative was so great that no matter what the grievances were they would eventually be forgotten in the name of the greater good that was being produced.

That attitude took a hit in the 20th century. On the one hand, the positive attitude towards our own culture was badly damaged by what some people have called the "Thirty Years World War". The Western world was basically at war with itself for some thirty years until a new power structure had emerged. Europe had lost, America had won. And science had demonstrated its potential to destroy all life on earth.

This is a theme that was meant to stay. While it was first the nuclear weapons that had the potential to kill all life on earth, the list of items grew longer over time. As of now there even exists a center dealing with what is known as "existential risk".[197] These are risks threatening to put an end to basically all human life if not life in general. Carl-Friedrich von Weizsäcker, physicist and philosopher and an highly influential figure in postwar Germany, wrote extensively on the dangers and dynamics of the nuclear arms race, as well as on nuclear power. Yet he once predicted that if the world were to survive

[196] Carl Benedikt Frey — Michael A. Osborne: *The Future of Employment: How S*usceptible Are Jobs to Computerisation?, Oxford, Tech. rep., Oxford University, 2013.

[197] See http://futureofline.org.

the 20th century it would rather be the ecology that would lead the list of priorities. He was absolutely correct. Here we are, the list of self-inflicted ecological problems growing by the day.

Now all that could in principle be blamed on bad behavior of some groups of people or institutions or on bad implementations of otherwise good ideas. That is at least the standard excuse of scientists and others. Their idea is that science is inherently good and should be promoted and instead we should discourage its abuse. Although I do not buy that argument, let's assume that it is a viable option. That still leaves us the destruction that is being caused by the research itself. To start, just consider how many nuclear weapons have been exploded in order to find out how construct them in a proper way. Their number is above 2,000! The human death toll due to radiation is estimated at 300,000.[198] The facility at Hanford has containers with sometimes unknown cocktails of highly radioactive substances, which are leaking into the ground. Cleaning the facility takes forever. And that is just one of many.

This list can be continued. Currently, experiments with GMOs means growing plants and studying their properties. However, it is unclear how genetic contamination of the environment can at all be prevented. Thus, regardless of the actual risks that a product sold at the counter poses, the research itself is not only risky but virtually guaranteed to involve major changes in the genetic makeup of natural organisms, which generally constitutes an irreversible change.[199]

More broadly, research consists as we have noted in controlled disturbance of the natural course of events. This means that there is no hope of a science that does not involve sometimes irreversible changes. However, it should be noted that this applies to any human activity so the human imprint on nature will not disappear if we stopped doing science. Thus, it is not a question of whether science has harmful side effects but how much of it there is and whether on balance it is worth it because of the insights we are gaining.

As a human activity, scientific research has taken and continues to take its toll on nature. The institutes, libraries, computer equipments, chemical substances and other objects that are being used by scientists to increase the knowledge are just the conventional share. Even if scientists worked in a company they would still consume as much and produce waste. However, it should be noted that as scientists they have helped to create these objects that have now become a normal part of life.

[198] Nicola Armaroli – Vincenzo Balzani: *Energy for a Sustainable World – From the Oil Age to a Sun-Powered Future,* Weinheim, Wiley-VCH, 2011.

[199] As this site notes, the number of tests involving plants in nature is declining in Europe: https://agnuhaan.wordpress.com/2013/04/18/keine-gentech-freilandversuche-in-deutschland-mehr/.

Chapter 7

Big Data

The invention of the computer has helped to increase the speed of computation, and it has thereby helped enormously in our understanding of the world. Increasing the capacity of storage facilities had meant that the understanding of the world has not only increased in depth but also in breadth. Whatever the acceleration of computational speed meant for science, it was only one half of the equation. Computers could calculate faster, but if we cannot store much of the results, then all we gain is maybe just one more tool for doing more complex calculations. But the minute that we can store large amount of data we can approach many questions with a totally different eye.

Here is an example. In the past, translations were done by algorithms that were being hand-coded in terms of the linguistic knowledge that they used. The machine was being told how to translate particular words or phrases. There was no way for it to know that the Hungarian word "szék" means the same as English "chair"; and it could also not know that the phrase "committee chair" will despite that *not* involve the word "szék". Thus, translation software involved a general mechanism, language independent, for translating between languages based on a bilingual dictionary, and a collection of such bilingual dictionaries for each pair of languages to translate between. The fact that the European Union at one point required all of its documents to be translated into the official languages spoken in the EU member countries meant creating a lot of such bilingual dictionaries, and they actually provided a major source for machine translation efforts.[200] This approach is one of the past, however. With the introduction of parallel corpora, that is corpora, which like the Rosetta stone contain the same text in two languages in a completely parallel fashion, it was possible to have the machine discover the rules of translation for itself. All that is needed is a general learning algorithm and a parallel corpus, and the machine could use that corpus that learn for itself how to translate. This is what Google is actually doing.

The basic idea is quite similar to what we think is the basic mechanism by which animals and humans are learning. Children learn languages by being exposed to it day in day out. They are thus forming a large corpus of data from their experience that is being more and more organized as the brain discovers the inner workings of the language. Whether or not that is the entire story is not the question, actually. Chomsky has argued that humans need an innate faculty for language in order to be able learn as much as they are actually learning from the data. But that does not imply that there isn't a lot that children can *only* learn by exposure to language. And it also doesn't mean that the amount of data they are getting is small.

[200] Actually, one needed such dictionaries only for the official languages, so that translation from Italian to Polish could be done via English.

So the present situation we are in is as follows. Science has used large collections of data to find out about what is going on. Scientists were using their own brain as a device to learn what the rules of nature are. Increasingly, however, they have found ways to teach machines how to learn, and that now feeds back into science. The corpora are getting so vast that the amount of knowledge surpasses what any person could experience in a lifetime by several magnitudes. Computer aided diagnostics is a case in point. One of the methods of making computer programs that diagnose X-rays and other images of the body is by pairing X-ray pictures with the diagnosis, and leave the rest to a learning algorithm. All that is needed in addition is access to a large data base of X-rays.

Another area where big collections of data are helpful, even necessary, is evaluation. When a drug has successfully completed the trials and is eventually put on the market, it is important to monitor the actual effect that it has on the people to which it is being administered. Insurance companies can thus get a better overview of the costs and benefits of the existing drugs. In China this has led to a rigorous screening of patients. Anytime people are seeking the help of a doctor all their data will be entered into the system, including the medicine they buy.[201] This leads to electronic data on a massive scale that can be used to evaluate what treatment is actually efficient.

Other disciplines are benefitting as well. To sociologists social media is a gold mine. They do not have to go out and do experiments, they can instead focus on evaluating the data that already exists and has been provided voluntarily by the users. Where that is not so other tricks have been invented. For example, you pay people a modest sum for the task of classifying products, and then feed the data to a machine that can learn this task. Eventually, the human is out of work.

Big Data has indeed created a revolution. Although at its core pretty unsophisticated, scientists have risen to the occasion. It is clear that simply amassing data is not enough. I already said that the brain is not just a big storage of facts and experiences. Quite the contrary.[202] The challenge has been to extract as much as one can of the data. The first step is to create interconnections. This has been the idea of search engines. They do not simply store data they also massively interlink it. The next phase has been to establish mechanisms for extracting higher order knowledge. This has been the idea behind FuturICT, a project led by the physicist Dirk Helbig, which had been running in the EU flagship competition.[203] The idea was that on the one hand the machine would collect all available data, and on the other it would be providing algorithms to

[201] It is unclear how one can effectively monitor whether they are actually taking the drugs that they are being prescribed. That is a weakness of the system.

[202] See Chapter 6: Passing on Knowledge — Generations.

[203] This would have meant funding of half a billion EURO for ten years. The project did not make it in the final round, though.

answer questions on the basis of this data. Examples include using regional and sectoral economic data that would allow to detect market trends and crime statistics to help police spot emerging criminal networks.

This is quite useful. How many of the applications require such a big machine is debatable. Helbig himself used the example of the financial crisis of 2008 that may have been prevented using clever algorithms and big data. However, some economists claim that the crisis was easy to spot. What was much harder was mustering the will to counteract it. Indeed, a lot of modern research seems to have less impact on society than it should have. As a case in point look at the literature on gender differences. From a scientific perspective there are robust findings showing that they exist and what exactly they are. And yet, in political discussions the consensus seems to be that none exist.[204] This however has effects on the laws that are being implemented. The knowledge-based society sometimes fails to act according to its own knowledge, a phenomenon we know all too well from our own lives as well.

What the new developments bring for science is not so clear. More knowledge is not necessarily better knowledge. This is known also in modelling. Models with many parameters can actually be worse than models with few of them. And it is not clear that the data that we are so busy collecting is increasing our understanding of the world. The situation is a bit like that of the beginning of the 19th century. The era of exploration that began two centuries or so earlier created an immense body of new data.[205] At the same time this increasing volume of facts created an uneasiness. More and more things appeared, and it might have seemed that nature was simply a vast collection of curiosities. But scientists quickly set out to change that impression. All along they believed that the basic natural laws were simple and few. So, they set out to find them. By the mid 19th century the efforts began to pay off. In 1855 Maxwell published his famous equations, thermodynamics was being developed around the same time with its fundamental equations, followed in 1905 by special relativity theory and, two decades later, quantum theory.[206] In about a hundred years physics had added to mechanics a handful of new disciplines together with their founding principles.

The same could in principle happen again. The era of big expansion will probably soon come to a halt. But the effects of having created these big collections will be felt long after. In could be that the data will be used to gain insights into nature and create new disciplines and revolutionize existing ones.

[204] Jussim: *Social Perception*. There commonly is a distinction being made between sex and gender, and that while sex differences undeniably exist, all gender differences are proclaimed to be cultural. While the scientific evidence is overwhelming that gender is not entirely a societal construct, gender studies proclaim the exact opposite.

[205] Peter Burke: *A Social History of Knowledge II: From the Encyclopaedia to Wikipedia*, Cambridge, Polity, 2012.

[206] Voit: *Not Even Wrong*.

Technical Limits of Access

The huge scale at which data is being produced puts once again one question on the agenda. And that is the question of access. It should be noted that having knowledge written up in a book is not doing much good if that book can't be accessed or read. Thus, it is as important for a society to create knowledge as it is to actually give access to it. This is why in sociology of knowledge it is actually emphasized that knowledge is in part an ability to achieve one's goals. This distinguishes the sociological perspective on knowledge from the philosophical perspective. Granted, knowing means being in accordance to truth. Knowledge *is* defined in terms of truth and not in terms of actions. And yet, when we ask what knowledge is good *for*—other than just being a goal in itself—then we can no longer ignore questions of access. Knowledge is power but only insofar as the knowledge can be used when needed.

We have spoken about the efforts on the part of civilizations to give knowledge wings.[207] In preparation of the next chapter we shall take a look at the present situation. Data is being stored in numerous data centers all around the world. This has various reasons. On the one hand, there is so much of it that the data needs to be distributed. On the other hand, one needs a certain amount of redundancy to prevent loss of data. And thirdly, a lot of data is actually only of local interest and therefore intuitively control over the data is exercised by keeping it physically close. So, a company that needs to store its data will likely prefer a local provider over a distant one. It should be noted also that the US claims jurisdiction over all information that passes through a server on US soil, so some companies or public institutions are actually trying to avoid using cloud services or data storage whose location they cannot control.

Given that the data being spread out as it is it must be transferred to other locations when requested. To get access to a single document requires an orchestrated effort of machines just like sending a letter requires an institution that manages dozens of people and transportation devices so that the letter arrives at its destination. The electronic archives require round the clock access to energy. They also require maintenance of servers, monitoring of data storage devices for their functionality, and much more. What used to be a system of libraries and interloan facilities is now a web of machines that allows us to see what there is on earth, so long as it is in electronic form.

The scale at which this network operates is mind blowing. People are not actually aware of how complex it actually is. To give an example, the mobile phones are connected to a network of antennas that have quite short range.

[207] Chapter 6: Passing on Knowledge — Surpassing the Limitations of Space

This is partly due to the fact that the power of the signal that the phone emits is limited. The antennas define overlapping *cells* within which they administer the phone calls.[208] As people move around, they move in and out of these cells, and the antennas need to pass on their clients when they are leaving their cells into the next one. This mechanism works seamlessly. In fact it is completely automated. A similar mechanism is in place for wireless computer access (WLAN). Data transfer is not only done using wireless technology. Most data traffic is still being sent using cables since the volume of wireless transfer is limited.

The energy required to keep this service up is considerable. It consumes as much electric energy is India.[209] And this is not even considering the hardware that needs to be produced as well (embodied energy).[210] A recent study suggested that while the footprint of the ICT industry was a mere 1.5 % in 2015, this could grow to 14 % by 2040, with smart phones taking up a bigger share than computers, laptops and displays.

It would require an entirely new kind of science to describe the effects of data. The question is not really whether we have an actual answer to our questions. The question is what the energy cost of that answer is. In the age of mass data, affordability of knowledge is in part determined by energy costs of storing it. This can be broken down into two questions. The first is how much we need to actually store one bit of information. The second is how long this bit of information is available. These questions are interrelated, as I discussed in Chapter 5: Conservation of Knowledge — Durability. It costs a lot of energy to chisel some letters into stone, but they may last several thousands of years. By contrast the energy costs on modern storage are minute by comparison but the information will last only for five years on average. Quantum computations are virtually free but the coherence is gone in a fraction of a second. Multiplying bit with time, we get an equivalent of what might be called *information impact*. Since information can be measured in energy, the unit of this measure is actually that of energy times seconds. So, when we compare Ancient information storage with modern one it is not enough to compare the costs of storage. We need to compare the information impact.

[208] See also the discussion on P. 139.

[209] See *How Green is Your Cloud* by Greenpeace 2012. A single data center consumes as much electricity as 180,000 households. Greenpeace quotes the following values from the SMART 2020 Study for 2007. The IT sector is responsible for approximately 2 percent of CO_2-emissions, roughly 830 billion tons, with half (407) stemming from computers and the other from communication (116) and external storage (307).

[210] For a quantification of this figure see Barath Raghawan – Justin Ma: The Energy and Emergy of the Internet, in *ACM SIGCOMM Workshop on Hot Topics in Networks (HotNets-X)*, 2011. Roughly, in 2011 the power consumption in total was between 170 and 307 GW for the information exchange and the production and maintenance of the embodied hardware. This is between 6.8 and 12.3 percent of the electrical power production.

This is because the information impact determines the information flow. Every time the information threatens to decay we need to copy it. This creates data flow for the purpose of maintaining data integrity. This is independent of the question of accessing the data for a purpose. Data throughout can be divided into two parts. One is to obtain useful information, requesting that data be inspected or copied, and the second is to manage data for its purpose.

There is more to that, though. There is also what I call the *ambient structure* that is necessary for the physical process to have the significance that we want it to have. Consider writing the answer to a question (yes or no) in the form of a single bit onto some storage. How do you know which of the many locations contains the answer to your question? This will be answered by a correspondence table, which links your question to the location of the answer. But consider this: how does the machine make sure it can identify the bit through some address? Isn't it the case that there is a material structure into which the location is embedded and which guarantees that the address points to a definite unchanging location similarly to an address on a letter? This is especially pertinent in quantum computing, though the miniaturization of standard storage devices will surely create similar problems. What we need to guarantee is a certain stability of the ambient structure into which the information is dropped. This structure costs way more energy that the raw bits themselves. But it is essential in the architecture of information, because it is responsible for the reliability of the system.

Dead and Living Knowledge

In the previous section I talked about the technical side of knowledge. I concentrated on data rather than knowledge, since data is what matters from a technical point of view. But data is only a tool. It helps us to remember something. The data processing is rather dumb, though. It is a mere pushing around of bits. What is needed, in the words of Bennett, is intelligent use of data. After all, stored data needs to be meaningful to us. We do not write books for the sake of writing them. That is the exception. Books are being written so that they will be read.

But what is the purpose of reading books? Of course, reading books is a pleasure to some. But also reading books or papers is a way to learn. Knowledge is transmitted from one mind to another by means of writing. Writing has the same function as money has for trade. Money was a way to liberate people from the necessity to have something useful to trade at a particular moment. Money unlike many goods does not perish. Similarly, books are a way to preserve knowledge until the moment comes that it is being needed.

This has proven to be a very effective tool. So, much so that nowadays research output is being counted in terms of written output and barely in terms of discussions and teaching. This however is a dangerous trend. There is a distinction to be made between knowledge that is alive and knowledge that is dead. Any knowledge in books and anything embodied in machines is dead for the purpose of human action. This is because the quest of knowledge is pursued not primarily to fill pages of paper or to construct machines, but in order to solve a human predicament or in order to achieve enlightenment.

The first task could in principle be handed over to a machine. Who cares really whether the answer is being delivered by a human or by a robot? Yet, there is a lot of science around which is far too difficult to be picked up from the books. My favorite example is high energy physics and its latest offspring, string theory. Although there is no shortage of popular books on the matter, it is virtually impossible to learn anything from these books that would allow you to understand or apply the theory let alone make new contributions to it. The amount of mathematics that goes into these theories is mind boggling. Anything from Lie-groups and their representations, algebraic topology, cohomology and topology not even counting classical differential equations, is being used there. How does one understand all this in physical terms? What is the significance of all that?

And that is not at all a trivial issue. It is one thing to trust the physicists that they understand what they are doing. But we do not understand it and there is no hope of us ever getting close. What then is the difference between a physicist and a priest or a guru? Of course the physicists claim that their science is based on solid experimentation and what have you. But to an outsider there is not a shred of tangible evidence that he can name that makes the claims of physicists more credible. The particle colliders are awe inspiring monsters, but there is no way to understand what happens inside of them. The mathematician Grothendieck once quipped that science is the only religion that does not admit it is one. If you listen to what physicists say whether we should adopt this theory of the universe or another you cannot help but wonder whether they truly have a superior sense of the fundamental questions of being.[211] There is any amount of rival theories, and they are often not as beautiful as their proponents make them. Why is that inherently different from a religious argument?

[211] Actually, Voit: *Not Even Wrong*, tells the story of fierce debates of this sort inside the physics community. The enthusiasm for string theory is not shared by everyone for a variety of reasons. One is that it cannot be put to test; and another is that it needs so many parameters that it is hard to imagine it is a theory in the genuine sense of the word. The arguments that are being exchanged do seem at times to be purely religious in kind. Voit also is skeptical that string theory can be called beautiful.

On the other end of the intellectual spectrum lies philosophy. The problems that beset philosophy are quite opposite to those of physics. The questions are so simple that it is hard to imagine that philosophy is actually difficult. Only when you start doing it you find that the art is not to get lost. Actually, philosophy is the art of articulating one's fundamental structures of belief. And that is a very difficult exercise indeed. When a student writes a term paper on Kant or Quine in principle he could just search the library or internet for useful quotes on the topic and then just fill the remaining lines with connecting prose. It would maybe suffice to get some credit points and a degree in philosophy in the end. But the intention of assigning that paper was not to have some pages be filled with more prose, or even some new interesting facts be discovered. It was to provide an opportunity to deal with an important question. As it stands, philosophy is not so much about answers but about questions. There is no hope that the essential questions of humanity can be answered. Instead, every generation, in fact every single human being needs to answer them for themselves. What counts therefore is the way we look at the question and the way we let it guide us. The process of writing the term paper and engaging with the question is the real reason why it was being proposed; once the paper is written it is already dead in the water. For only the process is what matters.

We are now facing a true dilemma in the coming years driven by the fact that just about any human activity can in principle be handed over to a computer. When asked about the prospects of ordinary working people this is often shrugged off with the remark that it has always been like that. Technology replaces labour and people have to get used to that. But can they? The mathematician Eric Weinstein once said in an interview that people go to university to learn. The investment pays off when they have actually learned it; and the way it does so is by people repeatedly applying their skills. But when computers are always better at rote tasks, what is the purpose of learning anything?

The answer, I think, is that we need humans in control of the whole enterprise. Humans will be the overlords of the machines, so to speak. For that however they will have to fully understand what the machines are doing or else they cannot know whether they do a proper job, and, more importantly, humans cannot know how to improve the machines. This leads to a rather daunting suspicion. If the presently declining educational system actually did nothing to lower the performance of Western societies, then what is needed is not so much any more a broad base of well-educated people. That phase has come to an end; we have machines for that. What is needed is a broad base of people competing for the few top positions at which they can actually still

contribute to the development. It is not the number of people that try which is relevant; it is the number of people that succeed. The more sophisticated the system gets the more high class performers we need.

This however renders knowledge in the traditional sense meaningless. The ramifications of this are not well understood but will soon emerge. Consider linguistics. Until recently, linguistics was considered an important ingredient for making linguistic software. So, it was deemed expedient to learn syntax, morphology, and the essence of language structure. But if algorithms can learn to do the job themselves, what is the point of teaching and learning linguistic theory? One answer is that linguistics was and still is the theory of language structure. So, even if the computers can do that for themselves, the programs do not yield theories, they just yield some neural networks. But it would be ridiculous to teach such a network in place of a theory because nobody would understand the workings of such a network. So, we are experiencing a split. For the purpose of producing applications theory becomes irrelevant. All that matters is to produce the right kind of training sets and to train the machines until they are ready for use. In the long run this will render linguistic theory for the sake of applications useless. It will be confined to areas where we still need to instruct humans, such as school teacher education.

That would make theory pretty uninteresting. It would then more and more look like art for art's sake. Most theoretical sciences will fall by the wayside and be taught simply as general education like philosophy or art history. This will take away the pride of people if all they can achieve by learning, by and large, is being added to an increasing body of job seekers rather than professional performers. This will not help the motivation to learn. The embodied scientific knowledge in the system will then mostly be in the hands of machines, and hence dead. You can't open a computer and find out how it works like you could with a car in the old days. But if humans lose the impetus to learn for whatever reason then virtually all knowledge dies. It is a matter of time until the machines die as well.

Relevance

If the pursuit of knowledge comes at a cost what could justify it? I have indicated above that the rapid decline of costs of inquiry or data manipulation have not made costs of knowledge insignificant, the contrary is actually the case. We are actually spending more and more money. The target figures for the industrialized nations are specified in terms of percentage of GDP. Thus, spending on knowledge is growing in absolute terms as the economy grows.

When Simonides tried to persuade Themistocles to learn his art, he was telling him about the benefits. But he did not mention or was perhaps not aware of the fact that more knowledge is not necessarily better knowledge. But maybe that was not his point anyway. In his time, it was after all very important to memorize things; there was no other way to know them. So, there is an art of memory involved in any case, it just changed its manifestation. For Greeks and Romans, it was the head, now it is hard drives.

But if the body of knowledge is limited then it is always worthwhile to ask what takes priority. What is it that is really relevant or important? For other than just being beautiful, knowledge also serves some purpose. We know facts so that we understand something. For example, in the 19th century historical linguistics was considered one of the most important areas of research. Nowadays it is hard to explain what the motivation for this was. It probably had to do with the fact that in those days people learned Greek and Latin at school, and Indo-European was the ancestor of most languages spoken throughout Europe. The idea of finding out about the ancient past of our civilization was a great motivation. And the results that were established are truly remarkable. It led to the reconstruction of an entire language and of the path how it split up into all the languages spoken today. We know how Latin branched off, and we can trace the history all the way to modern French, Italian and Spanish. On the way we discovered about the culture, the religious beliefs, mythologies and much more.

It may be said that with time when a field matures many of the initial questions are being answered and this explains why the interest eventually disappears. This may be the case with Indo-European and a lot of other fields. After all, it is hard to motivate new students to engage in a subject that has answered most of its questions. It basically means that it has lost its impetus. Typically, however, it is not so much the lack of questions as the arrival of new questions and other fields that drives out the old fields. That was surely the case with linguistics, which changed from a study of the historical development into the study of syntactic structure and then into human language processing, among other things. Maybe it had to be done this way, one at a time.

Historical linguistics like many university disciplines was pursued out of cultural concerns. History, anthropology, literature, art, none of them can be said to have monetary value. They are done because this is what our culture is like. We are curious about all these things. It was also thought that one could not be a member of a society without taking a thorough interest in its culture.

Interest was certainly enough for a mathematician. I remember someone giving a talk about some strange geometric structure; to motivate it, he simply stated that he got interested in them and so started to play around with them.

That was perhaps an honest assessment and at the time I was a student it was considered good enough. Nowadays it is however considered unwise to openly admit that one is simply driven by curiosity. In fact, one of the main reasons why governments invest so much money into research is that they truly believe that the ability for a country to compete in the international markets depends on its scientific research output. And scientists, eager to get the funding, assert just the same. There are two ways in which research is important. One is that the research directly feeds into the market in the form of innovations. This is the Silicon Valley business model, where the university is situated in the midst of an area full of companies and entrepreneurs ready to jump at any idea that promises financial rewards. Another way is that the research done in universities makes enough students interested so that they will eventually enter the area. For that it is important to have motivated professors, who in turn get their motivation from the research they are doing.

It is not clear how good the correlation is between research output and overall success of a country. It all depends on what one is measuring. Neither is there is a uniformly agreed measure of educational success, for university rankings are looking only at part of the sector, nor is GDP per capita a widely shared measure of success of the country as a whole.[212] But nobody questions the idea that a country depends on a well-functioning education system for its success, financial or otherwise.

Yet, as growth is harder and harder to come by the question arises whether the schools and universities are doing the best they can to contribute. This can be seen as an attempt by the administration to monitor and standardize education and to be able to manage it. The underlying idea, apart from a certain outflow of neoliberalist thinking, is that the main parameter for success is research funding, and so if the money is directed to the most promising areas, it will deliver the best return on investment. Indeed, those who resent this idea will have to realize that this has always been the case. Those who needed money in Prussia two hundred years ago to do research they deemed valuable had to write to the king to obtain funding.[213] Later it was the minister, now

[212] As for educational success, the OECD has introduced many more measures and tests, especially for schools, so that by now a wide array of success measures exist. Needless to say they paint a totally different picture from the university rankings. In terms of economic success, some people consider the Gini-coefficient to be very important for it indicates stability. Other people think that GDP needs to be modified to exclude non-wealth-generating sectors such as banking, real estate and insurance, see Michael Hudson: *Killing the Host: How Financial Parasites and Debt Bondage Destroy the Global Economy*, London, ISLET, 2015. Given that some countries are making a lot of money with their banks this would surely rearrange the success ranking.

[213] Ursula Klein: *Nützliches Wissen. Von der Erfindung der Technikwissenschaft*, Göttingen, Wallstein, 2016.

it is the funding agencies or the companies. In each case, results were being promised in return for money. Unless one was sufficiently rich there was no other way.

Science is thus never neutral. Always decisions have to be made as to what is important and what is not. There is never enough time and never enough resources to do everything. Societies thus think scientists are accountable for the money they are spending. They also need scientists to evaluate what is worth spending money for. That is not a trivial matter. Is the money for the Large Hadron Collider in CERN well spent? What about nuclear energy or even fusion? The answers do not only depend on what we want but also on whom we ask. Some physicists think that the ITER fusion reactor will never work. Others think it will.

There used to be a common answer: if we do not do the research other people will pursue it and it will be our disadvantage. If we abandon fusion research, other countries are happy to take our place. But that is not always true. The research into magnetic trains has led to no big outcome. Basically, only two countries were doing research, Germany and Japan, and neither built any train worth mentioning. No matter how much money is thrown after technology, it needs to deliver in the end. Moreover, motivating research by pointing at others is simply just passing on the responsibility. There is no guarantee that this is working, and as in all other areas of life we remain solely responsible for our own actions. The next chapter will show that the argument is based on a false premiss anyway: that the others will be capable of anything that we are not. But this is not so. There are genuine borders of knowability. As we are approaching them the question of relevance becomes more and more important.

CHAPTER 8

THE FUTURE OF KNOWLEDGE

Science and Religion

I started this book with the question of what the likely future is of our knowledge based culture. I proceeded by showing that knowledge needs the material culture for its own survival, and that there is an inextricable link between knowledge and energy. If that is so the answers are starting to take shape; for it has become clear that the fate of our knowledge based society depends on its ability to provide enough raw materials and energy, which means in particular that the often heard arguments that our energy predicaments will be solved once we are transitioning into an information based society, are misleading if not downright false. But before I can venture into that territory I need to talk about some more problems that our scientific culture faces. This may hopefully show that the crisis of the new scientific world order, so to speak, is not a short term blip. There are problems of a fundamental nature for which no easy fix exists and which can only be solved through a thorough debate.

The concept of a knowledge based society is due to Robert E. Lane. Its characteristics are among other things that the people want to understand nature and its inner mechanisms, that they use science to do that and are willing to pay a substantial price in terms of the gross national product for the advancement and teaching of science. And they will use that understanding of nature to their own benefit, or that of the society at large. The method of choice is the scientific method. So, the knowledge based society will use science as its religion, so to speak. For it invests science with the task of interpreting the world by means of its methods of divination. And it will base its decisions largely on what the scientific consensus is. This also applies to questions that are at heart moral in nature. It is not religion that is allowed to make recommendations but rather scientists because only they are in a position to also provide rational reasons for their positions.[214]

[214] This is what is actually going on. In Germany there are various committees that are supposed to make policy recommendations on essentially moral questions such as abortion. However, a person will not be admitted to these committees if he or she identifies primarily

Modern Western societies profess to being knowledge based in this sense. Hence they adopt a thorough distinction between faith and knowledge, between religion and science. The separation between state and church also has had the consequence of reducing faith based arguments to being private affairs, while officially sanctioned choices must be backed by scientific evidence.

There are good arguments why this is a questionable way to run a society. First and foremost it is not true that science can provide answers to moral questions. And this is because the facts are amoral to begin with. All science can do is provide arguments why certain moral choices are contradictory or why choosing one way of dealing with things is morally better than another. All our knowledge of how the little child grows inside the mother does nothing to answer the question of whether abortion is morally wrong. Similarly, there is no scientific argument deciding on whether we should eat meat or whether we should keep the fast. All science can do is point out the consequences of our behavior and then we see for ourselves what the moral stance on these consequences is. This is known as the Is-Ought-Thesis by David Hume. Simply put it says that the laws of what is and the laws of what ought to be are independent of each other.[215]

However, I shall set that aside. For there is a much overlooked problem of this construction. It is that science is slowly loosing its capability to play the role that the society wants it to play. This has many reasons some of which I shall discuss now.

At the present moment science is losing ground in terms of credibility. As I have discussed above, there is a growing awareness that science is very often abused or misused for personal gain. I am not talking here about the fact that scientists have beliefs. I am talking about studies that are being fabricated (where the results are more important than the actual experiment), where counter-evidence is willfully ignored and doubt is being created where an actual scientific consensus exists.[216] Another problem is that the media are not any more taking their role as gatekeepers seriously. Financial pressure is taking a toll on investigative journalism. Niklas Luhmann said that mass media have an important role in shaping the perception of the world. That incurs a big responsibility. Newspapers and television are increasingly failing the public, however.[217] Lest you think this is an innocent affair, let me

as belonging to a church or a religious group. However, being a scientist seems to guarantee impartiality. However, given the sometimes grotesque naivety of scientists in face of the human predicament, this is not a good idea.

[215] See Gerhard Schurz: *The Is-Ought Problem: An Investigation in Philosophical Logic*, Vol. 1. Trends in Logic, Dordrecht, Kluwer Academic Publishers, 1997, for a thorough discussion.

[216] Oreskes – Conway: *Merchants of Doubt.*

[217] See for example *A Glitch in the Matrix*, by David Fuller, https://www.youtube.com/watch?v=zQCTeGKHsVc.

tell you an anecdote. In the 1970ies the West German authorities believed that one way of winning the cold war was mass media. So, they made sure that everyone in Eastern Germany could receive Western television programs. However, there was a region in Saxony that the waves could not reach so the people there were effectively being cut off. Inside East Germany this was known as the "Valley of the Ignorant".[218] Their image the West was in fact *more* positive and rose-colored, not less. Mistrust for their own media was enough to make them believe the contrary of everything they were being told about the West.

Science is an organized form of inquiry. Society pays people to search for the scientific truth, and it needs to trust them. Even though knowledge is the aim, it cannot be had without an even greater amount of trust. If we can't trust the scientists in what they are doing then they can't create knowledge for us. Seen in this way it is not irrelevant who is paying for the inquiry. Given that science has become very expensive and university funding has not increased governments have turned a blind eye on the fact that private money is flowing into the universities in growing proportion. At first this looks like a good deal. The additional funds allow to buy new and better equipment which in turn promises better results. But sooner or later loyalty conflicts will start to appear. What if the results show that the new technology is not as useful as believed or even dangerous? Will this get published despite the fact that this threatens the business model of the donors?

The current strategy is to keep the matter secret. By way of example, the university of Cologne had signed a contract with the company Bayer whose content was kept secret. Public groups have tried to get access to it but the judges had turned them down. This is just one case among many. What it does is to create suspicion. It will make sure people will mistrust the results coming out of the institutions because they can't believe that they were the result of an impartial inquiry.

This highlights a problem we have with knowledge. Only a small fraction of those possessing it possess independent or first hand evidence. All the majority does is believe the scientists. But in order to believe them you need to trust them. Thus, if the trust is gone all that knowledge has been created in vain—if it is knowledge at all. Hence, transparency is needed. In an ideal world the companies would do the research themselves and the universities would be paid by the public to evaluate their claims. And in that ideal world everything that the universities do would be available for everyone without limits. But as matters stand, companies pay universities to do the research and use the credentials of the scientists to treat what they allow them to publish as an established fact.

[218] "Tal der Ahnungslosen"

The genetic industry is a case in point.[219] At this point it is the industry that pays for most of the actual research, be it at their own labs or at universities. Most publications are thus paid for by interested parties. They also finance lawmakers and oversight committees. In a recent case, the Federal Oversight Commission in Germany[220] in a statement on the safety of genetically modified organisms copied from documents made by Monsanto. Similarly, when NGOs tried to understand what made the EFSA[221] consider glyphosate safe for use they literally had to sue them to get access to the documents. In this connection it does not really matter whether the authorities are right or wrong. Maybe what they wrote or said was correct. However, this is slowly becoming irrelevant. For what this behavior risks doing in the long run is erode the trust in the institutions.

This is no small deal. Consider an ordinary person trying to understand his situation. For example, consider a farmer trying to understand whether growing genetically modified plants is a good thing. After all, this may for him not only be a financial decision but a moral one, too. It used to be said that at the core this is not a moral issue. All that is needed is that the plants he is growing pose no risk to the humans or animals that are supposed to eat them. But the problem begins right here. The farmer may insist that it is a moral problem in that it is of wider significance. We might understand this as follows. *Scientists define what it is for these plants to constitute a risk. In particular they define it in such a way that they can at least get an answer. But who knows whether they know what to pay attention to?*[222]

In detail, is it enough to just ask whether the plants are safe to eat? How about the wild animals that eat from the plants? Maybe we should care about them, too. Similarly, when the plants get treated with chemicals it kills the insects that we don't like (called pests) while also killing all the others. And as a result the birds that feed on them remain hungry. And so on. In the end we can see how little we as well as the scientists understand of the complexity of the ecological system. And yet we decide to change its parameters in a big way. The scientist may think of the moral issue as an underlying scientific problem. He may think that once we have collected all the evidence we may be able to make a well-informed decision based on nothing but facts. But that is nothing but wishful thinking. For nothing indicates the we shall get at these facts fast enough. Scientists are quick to point out the promises of a new technology and promote it (and themselves) before it is in any way clear what the downsides are. Innovations are thus thrown into the world on a big scale before we know what risks they constitute.

[219] Druker: *Altered Genes*.
[220] The BfR (Bundesanstalt für Risikobewertung).
[221] European Food Safety Authority
[222] This is also a point made Fukuoka: *Farming*.

Hence in view of this it would be wise to respect the moral nature of the question. One reason could be humility. The story of the scientific truth that can be easily shown to be superior to the primitive superstition has lost its appeal. Today's questions are not of that sort. It is not a question whether we believe in evolution or in creationism. You can't make an argument for or against GMOs in that way. Even someone firmly believing in evolution can still be against them. What is more, you can believe in evolution and still be religious. Morality is orthogonal to the issue. This brings us to the second point. Most likely it is futile to try and win people over by discarding their moral fears with scientific evidence. As Jonathan Haidt has argued, people differ in whether they allow facts to override their moral judgement. Essentially, he claims that the Western type liberals are an exception in that they are thoroughly utilitarian. To them anything is admissible as long as it does not hurt anyone.[223] Others will allow no amount of facts to alter their moral beliefs. And, I may add, most people are somewhere close to the second type.

It is quite similar with so-called scientific debates. It is rare for people to start with an open mind. I do not even know how that is possible. An example I know quite well is nuclear energy, more precisely nuclear fission reactors. They were a big issue when I went to school. People were divided over whether or not we should have them. Opposition was fierce. Underlying most arguments for or against, however, was an emotional attitude. You either thought of it as a demonic force, unstoppable when released. Or you thought of it as a technology like any other that we will master eventually.

Scientists sometimes note in despair how hard it is to make people change their minds. But the question is: why should people change their minds about it? This is not an exercise in mental flexibility. Ideally it is about the truth. So, if you firmly believe the answer, that is, if you believe you know what is going on, why should you change your mind? So what kind of evidence will be needed to make a supporter of nuclear energy reconsider their case? Will one nuclear catastrophe like Tchernobyl do? Maybe it happened because the Soviet Union was not enough advanced technologically. Will another one like Fukushima do? Surely the Japanese were advanced enough. But maybe the Japanese pushed too hard on nuclear energy to begin with. So, it is really not the technology that is risky. It has to be run properly. And so on. If you take two countries, France and Germany, side by side, the difference in political outcome couldn't have been more different. On the one hand a country, France, which gets most of its electricity from nuclear power, and on the other a country that has

[223] Note, however, that they too need to start from a moral base. All the utilitarians are doing is to reduce any questions of morals to a question of weighing the pain. Obviously, for them pain is to be avoided. And that is a moral principle.

experimented with the technology and then decided more or less against it.[224] But to this day opinions are divided as to why this happened the way it did. The suspicion is that the reasons are not scientific. The French society is not so much different from the German society to explain that either. Could it just be something more elementary such as the fact that Germany has (or has had) plenty of coal? The suspicion is that the reason the nuclear renaissance touted by the industry is not happening is not lack of will but rather a lack of economic viability.

And that is precisely what we should always remember. Even the most scientifically minded society will have to take note of nonscientific factors. Technologies may be safe, but once they are being built they are subject to constraints of all sorts. To build a reactor one needs skilled people, financial resources, and standardized materials. And to make sure all this is implemented properly, thorough oversight is needed as well. Finally, how we eliminate corruption?

Scientific questions are specific, to the point. You can investigate the nature of the steel that you need to build your reactor from. You can ask about probabilities for certain failures. But moral and religious questions are about everything at once. It is no good to pretend that we have no such questions. We ignore them at our peril. Psychologists know that. You can resist asking painful questions for a while, but hardly forever. They may come and haunt you in a different guise. Similarly, with religious questions. How do you answer these questions scientifically? What is the meaning of life in scientific terms? What is a good life? What is to be done and what is avoided and why?

There are many people who think that you cannot be both religious and do science. At least officially this seems to be a consensus. Although people are happy to claim that they believe in some sort of higher spirit they are reluctant to identify with any specific religious dogma. This means however that morality is a completely private affair. No public appeal can be made on the basis of religious dogmas. This goes for scientific publications as well. As Holm Tetens once complained[225] you can invoke all kinds of metaphysically weird entities like possible (or even impossible) worlds to argue your case. But dare to mention the word "God" and your submission to a philosophy journal will be rejected. You can be a naturalist no problem and pretend that the questions of life will eventually be reduced to a play of elementary particles and forces, and you can have your work published. But declare openly a religious affiliation or declare that the nature of the human soul is divine, and your arguments are no longer considered publishable. By contrast

[224] Sometimes with a big splash. The thorium reactor at Hamm-Uentrop got built in the 1970ies, was operational for a few months and then got shut down permanently.

[225] Tetens: *Gott denken.*

in the Middle Ages people saw no contradiction in doing logic and theology. And in modern days, no lesser people than Wittgenstein and Putnam have argued that religion is not illogical, it is simply alogical.[226] That means, you cannot refute religious attitudes by means of logic. That is so since the content of religious beliefs cannot be captured by a logical formula, so the tools cannot be applied.

Theologians have also recognised this. Georg Picht, who coined the phrase of the "educational catastrophe" to be talked about later, pointed out that science actually needs the ability to believe. Furthermore, he argued, it is not so that religion does not deal with truth. There is religious knowledge. And so, if we let go of religion because we do not want people to believe, we will also weaken science. Because if someone is truly rational and accepts nothing without proper evidence he will have to greet most of what scientists claim with utmost skepticism. The end of religion may not be the beginning of an enlightened era, it may well be the beginning of nihilism.[227]

Ignorance

We can detect changes not only in the appreciation of science by nonscientists. There is also a change inside science itself. Primarily on the question of how much science can contribute to answering our questions the views have become more pessimistic over time.[228] Max Weber held the view that though we may be ignorant about something there will eventually be an answer to every question. One generation later the optimism started to wane. Now it was claimed that we will not know everything. Our knowledge will contain gaps, so to speak; but that at least on important matters progress is the norm. Today it seems that in some areas we have reached the end of our capacity to obtain answers.

There is a common theme in the popular literature on science: our ignorance is growing. The more we know the more we don't know. On the face of it this is actually not true. The statement is self-contradictory since it acknowledges that our knowledge is growing. What is growing at the same time, however, is the number of questions that are being raised. And that

[226] See among others Hilary Putnam: *Renewing Philosophy*, Cambridge (Mass.), Harvard University Press, 1992. Putnam in his later life has put a lot of energy into the question of value judgements from a philosophical point of view, including the Humean thesis that values are orthogonal to facts.

[227] See the discussion in Hauke Ritz: *Der Kampf um die Deutung der Neuzeit. Die geschichtsphilosophische Diskussion in Deutschland vom Ersten Weltkrieg bis zum Mauerfall*, München, Wilhelm Fink, 2014.

[228] Peter Wehling: *Im Schatten des Wissens? Perspektiven der Soziologie des Nichtwissens*, Konstanz, UVK, 2006.

creates the impression of ignorance. For it is not the case that scientists are doing research on just anything that comes to their minds. Rather, research is guided by a question, which in turn is justified by previous research. This is a reality not the least since research needs money and spending money on anything must be justified beforehand. So, research is done because someone formulated a question that he and maybe some other people found interesting and worth pursuing. Typically, however, when an experiment is conducted, it raises many more questions. It rarely answers the initial question for good. Typically, the methodology is subject to criticism so more experiments need to be done, variations need to be introduced, better equipment needs to be used, more variables need to be controlled for, and so on.

So the number of questions that are waiting for an answer is indeed growing. They present a forceful case for applying more science. But they also demonstrate that the claim that we shall one day answer all of our questions is empirically false. I know of no such area of science.

However, consider this: the progress in technology not only introduces new things for us humans, they also change the environment we live in. The collection of essays by Nordmann et al. discusses whether or not science is undergoing a change and is slowly becoming what they call "technoscience".[229] The basic problem is that because humans have changed the earth to such a degree that there is hardly any untouched area left, and because most areas that humans inhabit have been more or less thoroughly invaded by technology, empirical scientists are no longer in a position to study nature. What is left for them to study is a technologically entrenched nature. This allows us to use our intimate knowledge of technology because we have introduced it in the first place. It also forces a new kind of interdisciplinarity on us. And the more technology progresses the more we are facing a surrounding of our own making. We are studying ourselves by studying nature. But that also means that we cannot study nature itself anymore. Consider genetically modified organisms. As it happens, there is no way to control them once they get released. Eventually, they will be all over the earth. And the DNA of plants and animals will be changed for good. Nobody knows what the implications of that will be. And from that point on nobody will be able to study nature as it was before that happened.

This highlights a dilemma that has been present all along. An experiment is a way to find out what is the case. It is performed because we do not know. But it is also a way to force nature into doing something it ordinarily would not do. So, in order to get a realistic picture scientists have come up with

[229] Alfred Nordmann – Hans Radder – Gregor Schiemann (eds.): *Strukturwandel der Wissenschaft. Positionen zum Epochenbruch*, Weilerswist, Velbrück Wissenschaft, 2014.

the idea of a "real world experiment".[230] This describes an experiment that takes the form of real events. For example, burning down a small part of the forest to see how this develops, releasing animals on an island to see how they thrive there, medicating people with a new drug without their knowledge, and so on, are forms of real world experiments. Unlike traditional experiment there is no illusion that the results of the experiment can be undone. The change is introduced willfully, changes are being observed and documented. The rationale is simply this: from a certain point on there is no knowing what the consequences really are. So, the best one can do is to implement them and see what happens in the hope of being able to stop it when a dangerous line has been crossed.

But such an approach has its limitations. Only very few people will risk their lives to fly in a space glider to test its performance. Few people will be desperate enough to try out new drugs when it is unknown how their body reacts to them. And that is why performing such experiments is highly regulated. Before one can perform the experiment safety precautions must be taken. For example, an airplane will not get an authorization unless it can be shown that a complete failure of the aircraft happens at most 1 out of 1 billion operation hours. That is, it must be shown that a plane crash will happen in less than 1 out of 100,000 years (per airplane).[231] The statistical calculations must be shown and must satisfy specified criteria. However, that in itself is a problem. Technology has always been risky. The problem is that the risks are growing exponentially. An airplane crash is dreadful, of course. But compared to that the disaster of Tchernobyl was of epic proportions. And nobody knows what the damage the release of some hazardous DNA will produce.

Ironically, as some have noted, the initial estimates for the likelihood of a so-called design basis accident or maximum credible accident had been estimated in the 1970ies to be around 1:20,000 to 1:10,000 years, which amounts to between 1 out of 200 Million and 1 out of 100 Million operating hours. Thus, an airplane has to be demonstrably 5 to 10 times more safe than a nuclear power plant. Given that around 400 nuclear power stations are in operation and we have had one such accident every 25 years, it seems that the second figure is pretty accurate. Moreover, as Ladkin bitterly notes, the possibility of a Fukushima type station blackout was noted in Perrow but simply ignored.[232]

[230] Matthias Groß: *Experimentelles Nichtwissen. Umweltinnovationen und die Grenzen sozial-ökologischer Resilienz*, Bielefeld, transcript, 2014.

[231] Peter Ladkin: Hazard Analysis and its Application, in Peter Ladkin – Christoph Goeker – Bernd Sieker: *The Fukushima Dai-Ichi Accident,* Wien, LiT, 2013, 33–44.

[232] Ladkin: *Hazard Analysis*; Charles Perrow: *The Next Catastrophe: Reducing Our Vulnerabilities to Natural, Industrial and Terrorist Disasters*, New Jersey, Princeton University Press, 2007.

The problem is that the prevailing attitude is that accidents are exceptions. They are not. They are part of the risks of operating a technology. It is just a matter of time until they happen. It is not a question whether a technology fails but when it does. This raises very serious questions concerning our regulatory system. In Taleb et al. it is claimed that the massive distribution of GMOs can in principle result in a fatal disturbance of the natural ecosystem worldwide.[233] Thus failure of the technology is actually *existential*.[234] Whence they argue that we must be absolutely certain what we are doing before we even consider it.

But this is not happening. At present the protection of private interests prevails. It is considered more important to let the current system continue than to risk that some companies will have to go bankrupt. Moreover, it is not that the risks some of these technologies pose are unknown. Certainly it is hard to quantify them exactly. But the bigger the risks are the more we need to err on the cautious side. So, it rather seems that currently we prefer to be willfully ignorant on the risks.

However, there is a fundamental dilemma. We need to make a choice on whether or not to proceed with GMOs or other high risk technologies. If we decide that they are too dangerous and should not be pursued then we decide that some of our questions will never receive an answer because we think that providing them will release the genie out of the bottle. Some think that human curiosity is so strong that it is not going to happen. It would require a coordinated effort to stop that research. Because any country that decides to ban research puts others that do not ban it at a momentary advantage.[235] It is hard to imagine a ban on research at this point. But if we continue on the current path then most likely the scale of the catastrophes will get much larger. And it is quite likely that at some point the scientists will take all the blame for it. If there is a sustained crop failure, say, this will be such a disaster that no one can defend themselves saying they did not know it was such a dangerous idea. If people go hungry they will not care much about subtleties of scientific knowledge.

One may say that painting such a bleak picture is unjustified. There is no proof that GMOs can destroy the natural ecosystem. True, but that is not the point. The point is not whether or not it will actually be that bad. The point

[233] Nissim Nicholas Taleb et al.: *The Precautionary Principle (with Application to the Genetic Modification of Organisms)*, Tech. rep., http://www.fooledbyrandomness.com/pp2.pdf, 2014.

[234] Nick Bostrom is one of the leading figures in the movement of identifying existential risks. It is a pretty frightening prospect that many of them are actually man made.

[235] This argument has been used quite often in the past. It presupposes however that the research will provide a real advantage in place of just a symbolic one. Because if we are just competitors in an Olympic race then it does not matter really what discipline we choose. It could be social psychology or mathematics in place of genetic engineering.

is whether there is an uncertainty as to whether it could be that bad. Given that no experiment to date has made us fully certain we know exactly what is going on, given that we have been remarkably unable to predict the future, any indication that technology could be fatally flawed is to be taken seriously. For if we believe in scientific progress we also ought to give credit to those that point at the risks. The more we understand natural ecosystems the more we know what could go wrong. The more we scour the earth and count life on earth the more we know what damage we are currently inflicting on it. There is no denying that we also do know more about the risks we are taking. Sure, we have always been racing into the unkown. But the speed seems to be accelerating.

Blackout

While the previous discussion focussed on internal problems of scientific progress, I will now return to the external constraints under which it is operating. The power of the technical civilization is currently at its peak and it threatens to get weaker over time. This casts a big shadow over the world. Consider only the problem of incurred risks. While you can run a nuclear reactor safely by adding all possible safety precautions, shutting it down presents a problem. A nuclear reactor needs to be cooled for several years after being shut down. During that period approximately 5 % of its name plate power is needed for that purpose. As long as the reactor is operational this is no problem: it produces that energy and much more. But when it is shut down the energy must come from elsewhere. If and when France will decide to shut down its fleet of nuclear reactors it will not only have to come up with energy to replace them. It will also have to provide for their cooling after service. And that requires an electricity net that functions without interruptions. If the net is down for 48 hours the reactor will explode. Imagine this however. The decommissioning will not happen because some better technology is suddenly available. I know of no such technology.[236] Rather it will happen because it has to, because the fuel runs out. Furthermore, it seems that when uranium runs out, so does oil, and coal. Depending on statistics, we have a few decades left until the shutdown will have to start. To have to rely on a grid powered by renewables is a tall order. We do not know at present how this transition will work. We only know that it has to.[237]

[236] Fusion will probably never work for us.

[237] Ugo Bardi – Sgouris Sgouridis: In Support of Physics-Based Energy Transition Planning: Sowing Our Future Energy Needs, *BioPhysical Economics and Resource Quality* 2.4 (2017), have estimated the costs of transitioning to renewables while that is still possible with our fossil infrastructure in place. Bottom line is that the present investments are no way near what is required.

Now take a step back and consider how much our society has become dependent on the continuous supply of energy, especially electrical energy. Knowledge, its acquisition, its preservation, the access to it as well as education in general, everything is predicated on a supply of raw materials and energy. The technological advances have not been used to consolidate the process but instead they have been used to fire up the progress even more. The easier it got to store information the more information got stored.[238] And more and more services and technologies have been invented and are now populating our lives.

However, that process has come to certain standstill. Total energy consumption is not growing.[239] The so-called middle scenario of Meadows et al. has been pretty accurate.[240] By way of example consider oil. The peak of crude oil production occurred in 2010 plus or minus a few years. Most of the oil that is being added to the market is not crude oil in the traditional sense. Its contribution to the oil supply is questionable, especially in the case of fracking. Discoveries are dismal. Other fossil fuels are not much different. We are living on borrowed time. The horizon until they technically run out is measured in decades for all of them. It used to be centuries not long ago.[241] It is not at all clear that renewable energies can make up for the shortfall. While it is not impossible to switch to renewables the investments must be much larger than they currently are.[242] And at the moment the investments

[238] This is known as Jevons-Paradox or the rebound effect.

[239] See Bardi – Yaxley: *A Lotka-Volterra Model*; Robert L. Hirsch – Roger H. Bezdek – Robert M. Wendling: *The Impending World Energy Mess. What It Is And What It Means To YOU!*, Ontario, Apogee Prime, 2010; Tariel Mórrígan: *Peak Energy, Climate Change, and the Collapse of Global Civilization: The Current Peak Oil Crisis*, http://www.global.ucsb.edu/climateproject/papers/pdf/Morrigan_2010_PECC2.pdf, University of California, Santa Barbara, 2010. The estimates vary but it is pretty clear that we are close if not at the peak.

[240] Donella H. Meadows – Jørgen Randers – Dennis L. Meadows: *Limits to Growth. The 30-Year Update*, Chelsea, Chelsea Green, 2013. See also Ugo Bardi: *The Limits to Growth Revisited*, Springer, Springer Briefs in Energy, 2011; Graham R. Turner: A Comparison of The Limits to Growth with 30 Years of Reality, *Global Environmental Change* 18 (2008), 397–411. The authors of the original study have also produced a rerun 30 years later with similar results.

[241] At the 2016 conference run by the ASPO Martin Jänicke spoke of his role in advising the Chinese government in environmental affairs. He said that the bosses of the coal mines were treated like ministers; they got every support they needed. The polit bureau used to say that China had coal for a few centuries. Until it was pointed out to them that with growth rates of 10 percent coal would be gone within decades, not centuries. To their credit, they completely changed the fundamentals of their energy policy since.

[242] See Bardi – Sgouridis: *In Support*. The amount needed for investment is given in percent of the current energy production. It depends crucially on the energy pay back time for the installations.

are not going up in the OECD countries. The problems are piling up. If climate change continues, summers might get so hot that power plants might have to be shut down because there is not enough water to cool them.[243]

Blackout is the ghost that haunts our civilization. Without electrical power nothing works. When the region of Münsterland experienced an outage of two weeks the results where so devastating that the German government produced a report showing just how bad the situation can get.[244] Apart from the obvious services there are a lot of other vital ones that cease to function almost immediately such as water supply, gas pumps, and even jails, as their security systems all work on electrical power. Hospitals have backup generators but fuel for a mere 48 hours.[245] Solar power systems will not work if all they do is feed the grid; telephones will not work, only portable ones, until the battery is empty. Now imagine we would all drive electric cars. Then unless power systems are local and independent of the larger grid even those would be useless. There is no denying that the time has come to an end where we do not have to worry.

In addition to all that communication breaks down as well. You will not be able to pay at the store, you will not be able to withdraw cash, and the administrations will be closed down entirely. All production in the factories will come to a halt, not only because the supply of materials will eventually end but also because literally everything is run electronically.

However in the long run there are additional problems. The raw materials on which the electrical system depends, are limited in supply. Ironically, it is not the rare earth metals that are most critical, it is the common ones such as lithium and copper. It has been calculated that it is not possible to switch to renewable energy and at the same time electrify the vehicle fleet.[246]

[243] See http://www.thebulletin.org/web-edition/columnists/dawn-stover/treading-water and http://oilprice.com/Energy/Energy-General/The-U.S.-Drought-and-Electricity-Generation.html. Thermoelectrical power generation was responsible for 41 percent of fresh water consumption in 2005. Nuclear power plants need 2750 l water per Megawatthour, coal powered stations 1890 and gas powered plants only 720. Hugh Byrd and Steve Matthewman: Exergy and the City: The Technology and Sociology of Power (Failures), *Journal of Urban Technology* 21 (2014), 85–102, claim that, at least in Kenya, India, Tanzania and Venezuela lack of rain caused some of the blackouts.

[244] The report 17/5672 of 24. April 2011, *TA-Projekt: Gefährdung und Verletzbarkeit der modernen Gesellschaften am Beispiel eines großräumigen und langandauernden Ausfalls der Stromversorgung* (Endangeredness and vulnerability of modern societies exemplified by a long lasting power outage in a larger region), available at http://dipbt.bundest- ag.de/dip21/btd/17/056/1705672.pdf.

[245] Norbert Rost, p.c. He did a survey near Dresden to find out about preparedness levels of communities.

[246] See Ernst Schriefl – Martin Bruckner: Bedarf an Metallen für eine globale Energiewende bis 2050 – Diskussion möglicher Versorgungsgrenzen, in Andreas Exner – Martin Held – Klaus Kümmerer, *Kritische Metalle in der Großen Transformation*, Springer, Springer Spektrum, 2016, 217–233, for various scenarios.

The following is therefore a likely scenario. When mining the minerals gets more expensive this percolates through the system and makes all other products more expensive. Given a likely trend of rising food prices relative to income, there will be less disposable income, and as a result less economic power to buy electronic equipment. Already the trend is not to buy any PCs or even laptops. Increasingly people just have smart phones, since they can do practically everything they need to do using these phones. Computers are not needed. Their use will be reduced to administrations, businesses and universities.

Theoretically we could deal with these problems. If the energy needed to store information is going down there is no need to worry about the share the information society needs for itself. We can then dispense with printed books and move to electronic versions (ebooks, PDFs and so on). However, the list of conditions under which this is a workable proposition is growing. The speed of computer chips is not increasing. Likewise, the density of information (cf. Footnote 170). So, it is not likely that Moore's law will apply in the future. But even if it did, the more sophisticated the equipment gets the taller the pyramid of factories that do the preprocessing. Materials of utmost precision need to be fabricated to make all this possible.[247] It is not useful to consider the case of chip manufacturing alone. They could not make the chip if they could not in turn rely on a supply of standardized high precision equipment. They too consume a lot of energy and raw materials.

Most energy is anyway not needed for information processing. It is needed to produce buildings and heat them, and to manufacture and move around goods. So, the energy hungry information sector is in direct competition with the other sectors that arguably provide more essential services. It is not the case as is sometimes claimed that these services are no longer needed once we have transitioned to the information society. It is not the case that there is no Peak Oil, only Peak Demand, as if the demand in energy will magically disappear. The fact the oil supply is not growing is that it cannot be expanded, and exploration costs are skyrocketing. A while ago Steve Kopits has given a presentation debunking this idea.[248] Specifically, he notes that any conservation of energy that the OECD countries manage to implement leads to an increase in energy consumption in emerging countries, mainly India and China.

[247] See Vaclav Smil: *Making of the Modern World. Materials and Dematerialization*, New Jersey, Wiley and Sons, 2014, on materials for electronics. The purity of electronic grade silicon is 9 – 11 N, which means defects may occur at most $1 \cdot 10^{-11}$ to $1 \cdot 10^{-9}$ times, a thousand times purer than silicon for photovoltaic cells.

[248] A summary can be found at http://ourfiniteworld.com/2014/02/25/beginning-of-the-end-oilcompanies-cut-back-on-spending.

Given this we should ask ourselves whether people will save heating to pay for their energy bills for information processing, and whether they will agree to pay more for their informational needs and thus consume less food. Given that food production is a rather inefficient affair—it costs a lot of energy to make food rather than providing us with energy[249]—we should be careful with our increased reliance on energy for just about anything that we are doing.[250]

The Education Catastrophe

As I have argued, knowledge needs energy. This energy input is however mostly of an indirect nature. It is usually paid for using money. One may be excused for thinking that energy problems can be solved by intelligent technology. But the educational sector more than any other is facing stiff financial problems. Already, the idea of an ever growing wealth is an illusion. James Kenneth Galbraith once wrote the that "[t]he financial crisis (and the world economic crisis it engendered) [...] represented not so much the natural outgrowth of natural inequality as a further phase; it was a consequence of the deliberate effort to sustain a model of economic growth based on inequality that had, in the year 2000, already ended. By pressing the system beyond all legal and ethical limits the United States succeeded in prolonging an "era of good feeling," and in ensuring that when the collapse came, it would utterly destroy the financial sector." Immanuel Wallerstein thinks that the falling profit margins will spell an end to capitalist system. Even if one does not share this pessimistic view, it is obvious that the rising welfare costs are eating the budgets of municipalities and national governments alike. And thus they are in direct competition with another growth industry: education.

The success of the modern technology can be traced to the end of the Second World War. Vannevar Bush, who had been a highly influential figure during the war, wrote a report "Science, The Endless Frontier" pushing for a continuation of the research efforts even after the war, saying that science was the pacemaker of technological progress. This, in conjunction with the so-called Sputnik shock of 1957, led to a massive increase in spending on education. It was however not only the technological advances that were used to argue for more education. Several influential figures in Germany, among

[249] David Pimentel – Marcia Pimentel: *Food, Energy and Society*, London, Edward Arnold, 1979.

[250] See the article by Carey W. King, https://www.americanscientist.org/article/the-rising-cost-of- resources-and-global-indicators-of-change. King measures standard of living by comparing the cost of food and energy to income. According to his data, this figure was the lowest around 2000 and is slowly rising ever since. This induces a change in the global structure of industry.

which Ralf Dahrendorf, pointed out that education is an essential necessity for a democratic society. In the context of the cold war, a race between two ideologies, communism and democracy, it was thus argued that governments had to take every possible measure to assure equal access to education. Thus, investment in education became the equivalent of the social policies of the 19th century.

However, as much as the early investment allowed a large share of the population to move up, that process has soon stalled.[251] Indeed, it had been noted that the original hope of providing better chances for the less well-off had not materialized. It is difficult to understand why that is so, but I would suggest that one factor is that there are limits as to how many people the market can absorb in particular sectors.[252] So the number of people in higher paid jobs has stalled. The only thing that keeps rising is costs. The era after the Second World War was characterized by a sustained period of growth—the French call it the glorious thirty years ("Trente Glorieuses")—which ended in the sixties. The expansion of the education sector ended abruptly. However, the people had been hired, and needed to be paid for. Moreover, as the race with the Eastern Block was still in full force, there was no turning back. Thus, began the phase of rising quotas of students in higher education, a trend that the OECD has been welcoming until quite recently. However, the increasing costs for higher education began to take a toll on the budgets.

The strategies for balancing this trend have been manifold. In Germany the attempts to make people pay for universities have been thwarted. After a brief period where students had to pay a modest amount (a thousand euros per year) the protests have led to an abandonment of that idea. Other countries in Europe are however charging students. In the US, where universities have been prized institutions that are proud of leading the ranks in practically all education surveys, the rising costs have been charged on the students instead. Within the last decades even the health care costs did not rise as much as the costs of education. And this is not just affecting private universities. The University of California Los Angeles, once free of charge for residents, now charges 13,000 USD per year just for tuition (2014), up from a mere 4,500 in 2002. It is hard to justify these enormous costs when you consider that my own university needs approximately 10,000 Euros per student and year to break even.[253]

[251] Tanjev Schultz – Klaus Hurrelmann (eds.): *Die Akademiker-Gesellschaft. Müssen in Zukunft alle studieren?*, Beltz Juventa, 2013.

[252] It does also seem that the distribution of intelligence plays into it. One cannot turn everyone into an academic or a high performing manager.

[253] This is an approximate calculation as the figures are hard to establish. The annual budget is 200 Million, all paid for by the government, with a total of 20,000 students, which amounts to 10,000 per year. For comparison, the annual budget of UCLA (42,000 students) was 5,7 billion, of which 45 % is the hospitals and 16 % is third party research. This means that

It used to be claimed that investments in education pay off. But that seems to be no longer true. First, the educational sector has expanded, so more people compete for the highly paid jobs. Second, while the incomes of most people have remained level, the costs of getting a degree have soared. In 2016 the average debt for a graduate was above 37,000 USD, up six percent from the previous year.[254] It does not seem likely that the investment will get the people a higher paid job that will somehow enable them to pay it back. Increasingly it is the fear that they will not get a job at all that makes students enter university. But there is a limit on how much they will be able to pay. Moreover, as Gordon puts it, the education sector is less and less efficient. The degrees awarded are less and less worth the money.[255]

Peak Knowledge

The 20th century was a century of startling success for the sciences. It has given many bright people an opportunity to pursue a scientific career. Yet, as much as progress seemed endless, even the scientists themselves started to ask whether that is just an illusion. Horgan collects the view of dozens of scientists, many of them Nobel Laureates, on this question.[256] Interestingly, no one mentions the fact that our scientific societies are burning fossil fuels at an unsustainable rate, and how we are to science when the party ends. There is a childlike belief that everything turns out well. As if "Limits to Growth" had never been published they exude in esoteric fantasies of final answers to everything. Especially worrying are people working in artificial intelligence, who dream of new super-humans equipped with arbitrary amounts of computing power.[257] The frightening aspect of it all is that these people never ask whether anyone wants to live like that.

But maybe they don't have to. As I outlined in this book, I do not share the assumption that the technological society stands a chance to continue for much longer. There are inherent reasons for that, as I have hopefully demonstrated. The underlying problem is lack of energy, but on the surface

---

the budget proper is 2,2 billion, or 56,000 per student, roughly the charge for out-of-state tuition (40,000 USD per year). Approximately 1/6 of the budget is covered by the State of California.

254 http://www.rollingstone.com/politics/features/taibbi-the-great-college-loan-swindle-w510880.

255 Robert J. Gordon: *Is US Economic Growth Over? Faltering Innovation Confronts the Six Headwinds*, Tech. rep. 63, www.cepr.org, Center for Economic Policy Research, 2012.

256 Horgan: *End of Science*.

257 Many chapters are not like that, and I have to say that the older a discipline the more sane the comments by the people. AI seems to be too young to understand its own limitations. And they don't seem to understand that the human predicament has been understood centuries before.

it can manifest itself in many ways. Mostly, it ends being a financial issue. Unfortunately, there are not few people who think that it is *only* a financial problem. If so, it could be solved by printing more money. But can it really? Recall that anything you do at any given moment means endless things you do not do at that moment. There is no way to change that. Any project that the government finances involves people and resources now working on *that* project as opposed to another. Maybe we should worry about transitioning our society to renewable energy instead of dreaming how to spend the energy we soon do not have.

So there is only so much we can do, and there is a limit on what we can pay for. Knowledge must be paid for, as students in the US and elsewhere are painfully aware. But they are not alone. Physicists are feeling the costs as well. At this point of time the largest device for studying subatomic particles is the Large Hadron Collider, housed in a ring of 27 km length close to Geneva. The costs of this machine as well as its energy consumption are colossal. It consumes 170 MW of electrical energy, the engine without detectors cost 3 billion Euros. There is no hope that a larger machine will be built.[258] The same holds for telescopes or space missions. Although there is talk of missions to Mars, so far people are just being promised one way tickets. This does not sound much like progress.

Nicholas Rescher has been one of the early figures in pointing out that basically all areas of science suffer from diminishing returns.[259] This law, also called Gossen's Law in economic circles, says that the benefit of each additional unit decreases. Translated into science this says that the value added by each patent, each instrument, each university or even each professor decreases as the numbers grow. But the costs remain the same, or in some cases, increase even. It is evident that there can be no infinite growth. Likewise, if you look at knowledge in general, the costs of creating new theories is increasing. And that is not only because newer theories are expected to be better. It is also because an expanding knowledge base in itself costs more resources to maintain.

We are thus approaching *peak knowledge*. This is the time at which the sum of all knowledge societies have at their disposal is at its maximum. This is similar to peak oil, the peak of oil production; or peak coal, or what have

[258] Indeed, in the 1990ies the US had started to build a particle collider in Texas but abandoned it after spending 2 billion dollars. Another project of that sort never even got past the planning stage. The fusion reactor ITER is currently estimated to cost well over 10 billion Euros and no one knows when it be finished.

[259] Nicholas Rescher: *Unpopular Essays on Technological Progress*, Pittsburgh, University of Pittsburgh Press, 1980.

you.[260] Of course, I am not saying that knowledge is just like oil. You do not extract it out of the ground. But as knowledge needs energy, of which a part is oil, we should have reason to think that knowledge is going to decline.

That is not good. However, as has been emphasized in connection with energy, it is not the case that we run out of it immediately nor does it mean that none will be available in the future. However, there will be far less. The parallel with knowledge is striking. For knowledge will certainly not disappear. For there is one thing that does not change: our own brain. People are more or less the same they used to be for thousands of years, and there is every indication that they will remain the same. Their capacity for learning and remembering is basically undiminished. What has changed during the last centuries is the addition of data storage, administration, and scientific knowledge production. Specifically the last cannot be expanded any more. This is to say that scientific progress is more or less at a standstill.

One should not confuse that with the more visible growth in the number of publications. Most publications are not really new. My own impression is that the rise in the number is due to easier methods of publications, a growing number of scientists and a global access. The impact of the individual publications is actually diminishing rapidly. It has been claimed that most publications in the humanities are read by just one person. This is no accident. Given that at a certain moment every workshop ended in the question of whether the participants would like to write up their contribution into a paper for a volume little time was left for people to read if they wanted to attend their usual share of conferences. There is only so much time one can spent on reading.

However, even if the written output would be curtailed, there still is a problem. If the newer and better physical theory is not only more expensive but also more difficult to understand there will eventually be a limit to our understanding of the world. While Newtonian physics can be taught in reasonable time, already relativity theory and quantum physics are beyond the conceptual reach of most people. Even if all mathematics is left out the conceptual reframing needed to understand modern physics is very difficult to do. In the end most people simply accept its claims. String theory is even worse. It is hard to see how it could at all be an explanation of things in this world. Mathematics is no different in this regard. Some of its theorems simply take years to understand. And if contrary to expectation the ultimate unified theory of physics was much simpler—what difference would that make? Newtonian physics is simple enough, but the application to real problems can still be dauntingly difficult.

[260] See Heinberg: *Peak Everything*, for the idea that basically all human activity has reached a limit.

I can understand that scientists are not willing to accept that the thing they have devoted their life's efforts to is going to belly up. Especially when in the case of Nobel Laureates they have acquired a lot of praise and fame doing that. But if their mission is to find truth, they should be willing accept the corollary that progress is over. Alternatively, they should find a compelling reason why that is not so. In the light of our ecological and energy predicament it is a waste of time to continue baseless growth fantasies.

The question then becomes the following. Suppose that the finances of universities and research institutes are under pressure. Then what are we supposed to do? Cut the salaries of all professors or reduce their number? If we reduce their number, which of the fields should take a hit? One could argue that lots of fields are of questionable value anyway. We do not really know what contribution the humanities make to the GDP other than turning out teachers. But what exactly is the added benefit of astronomy or high energy physics? Moreover, it is unlikely that the economic contribution of many sciences can be upheld if the research findings cannot be implemented.

This is not the point of making any recommendations. All that I am saying here is that these questions will sooner or later arise. And it is at this point not clear how they will be answered. As much as high energy physics is valued, if there are no accelerators then pursuing a career as a high energy physicist is a bit pointless. And similarly for the very complex and costly areas of medicine. In proportion to the financial decline there will be less medical treatments. And they will remove the incentives for certain domains.

Take this in conjunction with an increasingly ephemeral storage and we have a perfect mixture that ensures that once certain fields will be abandoned the knowledge about them will disappear in no time. Online journals whose subscription is discontinued for whatever reason or hard drive storage that goes offline for lack of maintenance, all these add to the problems that we already have.

What Are We Losing, Really?

The prospect of a sudden and prolonged blackout is pretty daunting. When the entire communication system breaks down, no reasonable decision can be made. Stock traders are looking into black screens, administrators cannot make decisions, and scientists are without any means of gathering data. Because everything is computerized, nothing works when the computer is out of order.

Imagine instead of a sudden blackout a prolonged emergency. Imagine that computers are not leaving us overnight but instead very slowly, maybe in the course of a few decades. Then, no matter how slow that process is,

we are witnessing a train wreck in slow motion. For the society is still used to expand. It is habituated to a situation where everything is getting better, regulations are becoming more and more sophisticated and science is solving more and more difficult problems. Dennis Meadows calls this *upward pressure*. It means that despite the fact that there is no growth, growth is still expected and societies are preparing themselves accordingly. They expect their children to go to college, to get better degrees than their parents, they expect doctors to cure even more diseases, they expect that houses be bigger and better at the same time. They expect that supply in goods increases while the ecological footprint diminishes with the help of intelligent technology.

However, these expectations will sooner or later be disappointed. Cracks are already appearing in the system. The bureaucracy is often incapable of effectively keeping the records in order. Scientists have to perform their job with less equipment. When the money runs out, the pressure is mounting to justify everything. Why run a synchrotron that costs millions every year? Why work with huge telescopes? And once that is being questioned it is clear that the next generation of students will not enroll in astronomy or high energy physics since they know it has no future. The next step is that the knowledge will not be passed on to the next generation. And then it will be lost forever.

Make no mistake: there are disciplines that are so difficult that you cannot really learn them from a book. They need an institutional framework for survival. And I suspect there are many of them. Some mathematical disciplines are like that: algebraic geometry or set theory, for example. Or consider molecular biology and genetic engineering. All these disciplines generally take a lot of effort and allow only the brightest to continue. Add to that a certain trend to make the teaching cheaper, like online courses. While officially this is about giving access to universities for more people, they also create the temptation to replace face-to-face conversations and teaching by something much cheaper but far less effective.

The effects will be dramatic. When the sciences at the frontier are being abandoned because of lack of money or incentives that will mean that sooner or later the institutes will be closed down. At first that will only affect some fringe technologies. But consider what happens when the expansion of computer technology is over. Currently there are so many universities that not only have faculties of computer science but are also heavily investing in big data science. They are all geared towards more and more data. Nobody at this point is asking how we can manage with less data. That would mean asking why they have their job in the first place.

Or consider what happens when insurance companies start refusing to pay expensive treatments. That will change the business model of pharmaceutical companies or makers of expensive medical equipment. It will

lower the prospects for many doctors. Eventually it will lead to a change in the medical teaching curriculum. Again, once the subject is no longer taught, the knowledge will probably be gone forever.

Alan Weisman in his book "The World Without Us" has us consider a world where humans have disappeared from one moment to the next. It is a thought experiment. But it tells a deeper story that is worth getting into. For it lays bare just how many risks we humans have piled up for ourselves and nature that we need to constantly monitor. The most harmless ones are skyscrapers that may eventually collapse. More concerning are the basins and lakes full of chemicals that will eventually be released, the dams that break, and the nuclear reactors that will explode in no time when left to themselves. What will happen when society lacks not the will but the capacity to keep these dangers at bay? Who will command all the knowledge to be able to successfully control all this? What will happen when the many chemicals that we have released start to react in uncontrollable ways? Who will be there to tell us how to contain them?

At a certain point it will start to look pretty much like the Red Queen Syndrome: "It takes all the running you can do to keep in the same place." says the Red Queen to Alice. Most critics of our technological society only talk about the dangers of an existing technology. They hardly argue from the point of view of a systematic decline. Some people say that it is all very bad but we are doomed to make progress. We simply have to continue the research and the expansion. That may be so. But the question remains what is awaiting us at the other side of the mountain. It may be more enticing to study what happens on the way up, but that is no excuse. For as I have demonstrated, there is ample reason to suspect that we will face an era of involuntary degrowth.

On the one hand nobody likes to lose, the loss of knowledge makes no exception. It is dramatic to face the prospect of forgetting. We all know that. But to those who have already forgotten it may feel as if they had never known. And it already appears to me as if the collective mind is already taking note. It does not appear to me as if the discovery of the Higgs-Boson at CERN was greeted with the same enthusiasm as similar events several decades ago. It does not seem that people are as anxious to await the landing of humans on Mars as they were when sending people into space was still a dream. Somehow, maybe, it is that all this has become normal. But I also notice that no one talks anymore enthusiastically about the new PC that he bought that has faster processor or better graphics. Because there is no more PC of that sort. Now it is new apps that people talk about.

Indeed, even if that sounds unlikely, we will adapt to the new situation, whatever it may be. Much of what we so desperately like to know will at some point be irrelevant to us. That happened in the past as well. Once upon a time

historical linguistics was an important subject. Linguists wanted to know how Indo-European looked like and where the people came from that spoke this language. Nowadays the institutes for historical linguistics are being closed and nobody cares. It could happen to high energy physics and astronomy as well.

Many people are already pretty bored by the current way of doing science. Increasingly, it consists in filling out forms to ask for money, in massaging one's CV so as to look more formidable, and to promote the research that one has no time to do. Society anyway does not consider scientists worth their money if all they do is sit in an armchair and think. If we add the fact that the entire system is predicated on a growing market of conferences on all continents or far out locations and that traveling there is a waste of resources, scientists may no longer enjoy the image of promoters of societal wellbeing. For it is one thing to study the climate and another to save it. Politicians are mostly happy if scientists stick to the former, and as a compensation they will continue to pay for their conferences.

On the other hand, a lot of justification of science does appear to be self-congratulatory. The benefit of it is less clear than it appears at first sight. So, far, the natural sciences have had the upper hand, they could demonstrate without doubt that they could deliver on their promise. But that was predicated on a massive use of raw materials and energy. Take that away and the advantage of doing science is not so clear anymore. There is a lot of science that is purely defensive in nature. What exactly is the use of counting how many birds there are on earth and how they are migrating? Of course for a biologist this an interesting question and they may say that it helps to understand how to prevent damage. But who knows whether it does anything at all to change the destruction of wildlife? Who can guarantee that the knowledge is not also used to exploit nature even better? Not that I deny the value of the knowledge, I am asking about the real benefits.[261] Seen this way it is not clear why all the research into biology, chemistry or physics is of larger value than editing the complete works of some philosopher. Consider by way of example the climate conferences. It almost seems as if the mere fact that agreements have been signed makes the pollution of our atmosphere disappear. Sure, if the Paris agreement were just a piece paper, the countries would not have fought so bitterly for their interests. But still it does seem that despite best interests of the parties involved to reverse the ongoing

[261] Rex Weyler, one of the founders of Greenpeace, recently shared that sentiment, see http://www.energybulletin.net/stories/2012-09-05/nature-system-systems. His conclusion: as long as society does not change, all research is in vain. The idea promoted by politicians that researchers are supposed to find solutions to our problems, is pure laziness. It allows them to defer action. The notion that we do not know enough to make decisions is wrong. We know essentially what we need to know. We still have to do it, though.

development is a tall order.[262] And that means that no amount of scientific modelling will do the trick. The basic correlation between economic growth and destruction of nature and the atmosphere is known. What is not known is how science can effectively help to change it.

In the end, the energy predicament that is upon us may threaten our knowledge. But it may also end the destruction that our knowledge has enabled. As much as scientists have been able to impress societies with their technological sophistication, they are complicit in the great game. Science is the motor of civilisation. When the curtain finally falls, the verdict will be harsh. It is time to remember the Latin proverb: "Quidquid agis, prudenter agas et respice finem." (Whatever you do, do it diligently and watch out for the end.)

[262] Germany has not managed to do it. Basically, the only time CO_2 emissions went down was after the collapse of East Germany and the ensuing decommissioning of a lot of dirty coal plants. The continuing rise of GDP is however incompatible with lowering emissions. Germany just like any other nation cannot both grow the economy and restrain emissions.

BIBLIOGRAPHY

Achourioti, Theodora – van Lambalgen, Michiel: A Formalisation of Kant's Transcendental Logic, *The Review of Symbolic Logic* 4 (2011), 254–289.

Aczel, Amir D.: *The Artist and the Mathematician*, New York, Basic Books, 2007.

Armaroli, Nicola – Balzani, Vincenzo: *Energy for a Sustainable World – From the Oil Age to a Sun-Powered Future,* Weinheim, Wiley-VCH, 2011.

Bardi, Ugo: *The Limits to Growth Revisited*, Heidelberg, Springer Briefs in Energy, Springer, 2011.

Bardi, Ugo: Mind Sized World Models, *Sustainability* 5 (2013), 896–911.

Bardi, Ugo – Sgouridis, Sgouris, In Support of Physics-Based Energy Transition Planning: Sowing Our Future Energy Needs, *BioPhysical Economics and Resource Quality* 2.4 (2017).

Bardi, Ugo – Yaxley, Leigh: *A Lotka-Volterra Model of Worldwide Energy Consumption – ASPO-5*, 2006.

Barthes, Roland: *Mythologies*, Paris, Editions du Seuil, 1957.

Barwise, Jon – Perry, John: *Situations and Attitudes* (4th ed.), Cambridge (Mass.), MIT Press, 1986.

Beniger, James R.: *The Control Revolution. Technological and Economic Origins of the Information Society*, Cambridge (Mass.), Harvard University Press, 1986.

Ben-Naim, Arieh: *Entropy and the Second Law. Interpretations and Mis-Interpretations*, Singapore, World Scientific, 2012.

Ben-Naim, Arieh: *Entropy Demystified. The Second Law Reduced to Plain Common Sense* (2nd ed.), Singapore, World Scientific, 2008.

Bischof, Norbert: *Struktur und Bedeutung. Eine Einführung in die Systemtheorie* (2nd ed.), Bern, Hans Huber, 1998.

Burke, Peter: *A Social History of Knowledge II: From the Encyclopaedia to Wikipedia*, Cambridge, Polity, 2012.

Byrd, Hugh – Matthewman, Steve: Exergy and the City: The Technology and Sociology of Power (Failures), *Journal of Urban Technology* 21 (2014), 85–102.

Carsetti, Arturo: *Causality, Meaningful Complexity and Embodied Cognition*, Heidelberg, Springer, 2010.

Cassirer, Ernst: *Philosophie der symbolischen Formen Erster Teil. Die Sprache* (2nd ed.), Darmstadt, Wissenschaftliche Buchgesellschaft, 1953.

Chomsky, Noam: *Rules and Representations*, New York, Columbia University Press, 1980.

Coleman, James S.: *The Asymmetric Society*, Syracuse, Syracuse University Press, 1982.

Coulmas, Florian: *Writing Systems. An Introduction to Their Linguistic Analysis,* Cambridge, Cambridge University Press, 2002.

Damerow, Peter – Nissen, Hans J. – Englund, Robert K. (eds.): *Archaic Bookkeeping: Early Writing and Techniques of Economic Administration in the Ancient Near East*, Chicago, University of Chicago Press, 1993.

Diamond, Jared M.: *The Third Chimpanzee: The Evolution and Future of the Human Animal*, London, Hutchinson Radius, 1991.

Diamond, Jared M.: *The World Until Yesterday: What Can We Learn from Traditional Societies*, New York, Viking Press, 2012.

Druker, Steven M.: *Altered Genes, Twisted Truth: How the Venture to Genetically Engineer Our Food Has Subverted Science, Corrupted Government, and Systematically Deceived the Public*, Salt Lake City, Clear River Press, 2015.

Earman, John: *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*, Cambridge (Mass.), MIT Press, 1992.

Eisenstein, Elizabeth I.: *The Printing Press as an Agent of Change*, Cambridge, Cambridge University Press, 1983.

Evans, Nicholas D.: *Dying Words. Endangered Languages and What They Have To Tell Us,* Oxford, Wiley-Blackwell, 2010.

Fara, Martha: *Visual Agnosia. Disorders of Object Recognition and What They Tell Us*, Oxford, Wiley-Blackwell, 1990.

Finke, Peter: *Ökologie des Wissens. Exkursionen in eine gefährdete Landschaft*, K. Alber, 2005.

Fleck, Ludwik: *Entstehung und Entwicklung einer wissenschaftlichen Tatsache: Einführung in die Lehre vom Denkstil und Denkkollektiv*, Frankfurt am Main, Suhrkamp, 1980.

Frenkel, Edward: *Love and Math. The Heart of Hidden Reality*, New York, Basic Books, 2014.

Frey, Carl Benedikt — Osborne, Michael A.: *The Future of Employment: How Susceptible Are Jobs to Computerisation?*, Oxford, Tech. rep., Oxford University, 2013.

Fukuoka, Masanobu: *The Natural Way of Farming: Theory and Practice of Green Philosophy,* Madras, bookventure, 1993.

Gaifman, Haim: On Ontology and Realism in Mathematics, *The Review of Symbolic Logic* 3 (2012), 480–512.

Gettier, Edmund: Is Justified True Belief Knowledge?, *Analysis* 23 (1963), 121–123.

Gibson, James J.: *The Senses Considered as Perceptual Systems*, Boston, Houghton Mifflin, 1966.

Gordon, Robert J.: *Is US Economic Growth Over? Faltering Innovation Confronts the Six Headwinds*, Tech. rep. 63, www.cepr.org, Center for Economic Policy Research, 2012.

Gordon, Robert J.: *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War*, New Jersey, University of Princeton Press, 2016.

Gross, Matthias: *Experimentelles Nichtwissen. Umweltinnovationen und die Grenzen sozial-ökologischer Resilienz*, Bielefeld, transcript, 2014.

Heinberg, Richard: *Peak Everything. Waking Up to the Century of Declines* (2nd ed.), British Columbia, New Society Publishers, 2010.

Hilbert, Martin – López, Priscila: The World's Technological Capacity to Store, Communicate, and Compute Information, *Science* 332 (2011), 60–65.

Hirsch, Robert L. – Bezdek, Roger H. – Wendling, Robert M.: *The Impending World Energy Mess. What It Is And What It Means To YOU!*, Ontario, Apogee Prime, 2010

Horgan, John: *The End of Science. Facing the Limits of Knowledge in the Twilight of the Scientific Age,* New York, Broadway Books, 1996.

Hudson, Michael: *Killing the Host: How Financial Parasites and Debt Bondage Destroy the Global Economy*, London, ISLET, 2015.

Jackendoff, Ray: *Semantics and Cognition*, Cambridge (Mass.), MIT Press, 1983.

Janich, Peter: *Handwerk und Mundwerk. Über das Herstellen von Wissen*, München, C. H. Beck, 2015.

Jussim, Lee: *Social Perception and Social Reality. Why Accuracy Dominates Bias and Self-Perception*, Oxford, Oxford University Press, 2012.

Kant, Immanuel: *Kritik der reinen Vernunft*, Frankfurt am Main, Suhrkamp, 1974.

Keen, Steven: *Debunking Economics. The Naked Emperor Dethroned*, London, Zed Books, 2011.

Klein, Ursula: *Nützliches Wissen. Von der Erfindung der Technikwissenschaft*, Göttingen, Wallstein, 2016.

Klimpel, Paul – Keiper, Jürgen (eds.): *Was bleibt? Nachhaltigkeit der Kultur in der digitalen Welt,* iRights.Media, 2013.

Kogge, Werner: *Experimentelle Begriffsforschung. Philosophische Interventionen am Beispiel von Code, Information und Skript in der Molekularbiologie*, Berlin, Velbrück Wissenschaft, 2017.

Konrad, Franz-Michael: *Von der Geschichte der Schule. Von der Antike bis zur Gegenwart*, München, C.H. Beck, 2007.

Kornai, András: *Mathematical Linguistics. Advanced Information and Knowledge Processing*, Berlin, Springer, 2007.

Kracht, Marcus: Gnosis, *Journal of Philosophical Logic* 39 (2010), 397–420.

Kracht, Marcus – Klein, Udo: Notes on Disagreement, in Daniel Gutzman – Jan Köpping – Cécile Meier (eds.): *Evaluations – Denotations – Entities. Studies on Context, Content and the Foundation of Semantics,* Amsterdam, John Benjamin's, 2014, 276–305.

Ladkin, Peter: Hazard Analysis and its Application, in Peter Ladkin – Christoph Goeker – Bernd Sieker, *The Fukushima Dai-Ichi Accident,* Wien, LiT, 2013, 33–44.

Lakatos, Imre: *Proofs and Refutations. The Logic of Mathematical Discovery*, Cambridge, Cambridge University Press, 1976.

Lee, Penny: *The Whorf Theory Complex. A Critical Reconstruction*, Amsterdam, John Benjamin's, 1996.

Leibniz, Gottfried Wilhelm: *Die Grundlagen des logischen Kalküls* (Lateinisch – Deutsch), *The Foundations of the Logical Calculus* (Latin – German), Philosophische Bibliothek 525, Hamburg, Meiner Verlag, 2000.

Levinson, Stephen C.: *Space in Language and Cognition. Explorations in Cognitive Diversity*. Language, Culture and Cognition 5, Cambridge, Cambridge University Press, 2003.

Leyton, Michael: *Symmetry, Causality, Mind*, Cambridge (Mass.), MIT Press, 1992.

Locke, John: *An Essay Concerning Human Understanding*, 1690.

Lyons, Martin: *Books – A Living History*, London, Thames and Hudson, 2011.

Lyre, Holger: *Informationstheorie. Eine philosophisch-naturwissenschaftliche Einführung*, München, Wilhelm Fink, 2002.

MacKenzie, Donald: *Mechanising Proof. Computing, Risk and Trust*, Cambridge (Mass.), MIT Press, 2001.

Mackie, John Leslie: *The Cement of the Universe: A Study of Causation*, Cambridge (Mass.), MIT Press, 1980.

Mancosu, Paolo: Measuring the Size of Infinite Collections of Natural Numbers: Was Cantor's Theory of Infinite Numbers Inevitable?, *Review of Symbolic Logic* 2 (2009), 612–646.

Mates, Benson: *Stoic Logic*, California, University of California Press, 1973.

McCulloch, Warren S. – Pitts, Walter: A Logical Calculus of the Ideas Immanent in Nervous Activity, *Bulletin of Mathematical Biophysics,* Vol. 5 (1943), 115–133.b

Meadows, Donella H. – Randers, Jørgen – Meadows, Dennis L.: *Limits to Growth. The 30-Year Update*, Chelsea, Chelsea Green, 2013.

Mórrígan, Tariel: *Peak Energy, Climate Change, and the Collapse of Global Civilization: The Current Peak Oil Crisis,* http://www.global.ucsb.edu/climateproject/papers/pdf/Morrigan_2010_PECC2.pdf, University of California, Santa Barbara, 2010.

Mortley, Raoul: *From Word to Silence I. The Rise and Fall of Logos*, Bonn, Peter Hanstein, 1986.

Müller, Lothar: *Weiße Magie. Die Epoche des Papiers*, München, Carl Hanser Verlag, 2012.

Nordmann, Alfred – Radder, Hans – Schiemann, Gregor (eds.): *Strukturwandel der Wissenschaft. Positionen zum Epochenbruch.* Weilerswist: Velbrück Wissenschaft, 2014.

Oreskes, Naomi – Conway, Erik M.: *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*, London, Bloomsbury Press, 2010.

Pagel, Lienhard: *Information ist Energie. Definition eines physikalisch begründeten Informationsbegriffs*, Wiesbaden, Springer Vieweg, 2013.

Parsons, Terence: *Articulating Medieval Logic*, Oxford University Press, 2014.

Perrow, Charles: The Next Catastrophe: Reducing Our Vulnerabilities to Natural, I*ndustrial and Terrorist Disasters,* Princeton University Press, 2007.

Peterson, Jordan B.: *12 Rules for Life. An Antidote for Chaos*, Random House Canada, 2018.

Peterson, Jordan B.: *Maps of Meaning. The Architecture of Belief*, London, Routledge, 1999.

Pimentel, David – Pimentel, Marcia: *Food, Energy and Society*, London, Edward Arnold, 1979.

Pinker, Steven: *The Language Instinct: The New Science of Language and Mind*, Penguin, 1994.

Premack, D. G. – Woodruff, D.: Does the Chimpanzee have a Theory of Mind?, *Behaviour and the Brain Sciences* 1 (1978), 551–526.

Pullum, Geoffrey K.: *The Great Eskimo Vocabulary Hoax and Other Irreverent Essays on the Study of Language,* The University of Chicago Press, 1985.

Putnam, Hilary: *Renewing Philosophy*. Cambridge (Mass.), Harvard University Press, 1992.

Putnam, Hilary: The Meaning of "Meaning", in *Mind, Language and Reality*, Cambridge, Cambridge University Press, 1975, 215–271.

Raghawan, Barath – Ma, Justin: The Energy and Emergy of the Internet, in *ACM SIGCOMM Workshop on Hot Topics in Networks (HotNets-X),* 2011.

Rescher, Nicholas: *Unpopular Essays on Technological Progress,* Pittsburgh, University of Pittsburgh Press, 1980.

Ritz, Hauke: *Der Kampf um die Deutung der Neuzeit. Die geschichtsphilosophische Diskussion in Deutschland vom Ersten Weltkrieg bis zum Mauerfall*, München, Wilhelm Fink, 2014.

Sacks, Oliver: *The Man Who Mistook His Wife for a Hat and Other Clinical Tales*, Duckworth, 1985.

Schriefl, Ernst – Bruckner, Martin: Bedarf an Metallen für eine globale Energiewende bis 2050 – Diskussion möglicher Versorgungsgrenzen, in Andreas Exner – Martin Held – Klaus Kümmerer, *Kritische Metalle in der Großen Transformation*, Springer Spektrum, 2016, 217–233.

Schultz, Tanjev – Hurrelmann, Klaus (eds.): *Die Akademiker-Gesellschaft. Müssen in Zukunft alle studieren?*, Beltz Juventa, 2013.

Schurz, Gerhard: *The Is-Ought Problem: An Investigation in Philosophical Logic,* Vol. 1. Trends in Logic, Dordrecht, Kluwer Academic Publishers, 1997.

Searle, John R.: *The Construction of Social Reality*, Free Press, 1997.

Searle, John R.: *Mind, Language and Society: Philosophy in the Real World*, New York, Basic Books, 1997.

Smil, Vaclav: *Making of the Modern World. Materials and Dematerialization*, New Jersey, Wiley and Sons, 2014.

Stehr, Nico: *Arbeit, Eigentum und Wissen: Zur Theorie der Wissensgesellschaft,* Frankfurt am Main, Suhrkamp, 1994.

Stehr, Nico: *Die Zerbrechlichkeit der modernen Gesellschaft,* Weilerswist, Velbrück Wissenschaft, 2000.

Stiegler, Bernard: *Die Logik der Sorge. Verlust der Aufklärung durch Technik und Medien.* Frankfurt am Main: Suhrkamp, 2008.

Taleb, Nissim Nicholas et al.: *The Precautionary Principle (with Application to the Genetic Modification of Organisms),* Tech. rep., http://www.fooledbyrandomness.com/pp2.pdf, 2014.

Tarski, Alfred: The Concept of Truth in Formalized Languages, in J. Corcoran (ed.), *Logic, Semantics, Metamathematics*, Indianapolis: Hackett Publishing, 1983, 152–178.

Tetens, Holm: *Gott denken. Ein Versuch über rationale Theologie*, Reclam Verlag, 2015.

Trabant, Jürgen: *Weltansichten. Wilhelm von Humboldts Sprachprojekt*, München, C. H. Beck, 2012.

Turner, Graham R.: A Comparison of The Limits to Growth with 30 Years of Reality, *Global Environmental Change* 18 (2008), 397–411.

von Uexküll, Jakob: *Kompositionslehre der Natur. Biologie als undogmatische Naturwissenschaft*, Berlin, Propyläen, 1980.

Voit, Peter: *Not Even Wrong. The Failure of String Theory and the Search for Unity in Physical Law*, New York, Basic Books, 2006.

Weber, Max: *Wirtschaft und Gesellschaft* (5th ed), Tübingen, J.C.B. Mohr (Siebeck), 1985.

Wehling, Peter: *Im Schatten des Wissens? Perspektiven der Soziologie des Nichtwissens*, Konstanz, UVK, 2006.

von Weizsäcker, Carl-Friedrich: *Die Einheit der Natur*, München, Carl Hanser Verlag, 1971.

von Weizsäcker, Carl-Friedrich: *Der Aufbau der Physik*, München, Carl Hanser Verlag, 1985.

Williamson, Timothy: *Knowledge and Its Limits*, Oxford, Oxford University Press, 2000.

Winterling, Alois: *Caligula: Eine Biographie,* München, Pantheon Verlag, 2007.

Yates, Francis: *The Art of Memory,* Francis and Kegan Paul, 1996.

A SOROZATBAN EDDIG MEGJELENT

Monográfia

Lovász Irén: *Szakrális kommunikáció*

Galsi Árpád: *Jakab, az Úr testvére*

Pap Ferenc: *Templom mint teológia. Kulcsok az Ezékiel 40-48. értelmezéséhez*

Tóth Sára: *A képzelet másik oldala. Irodalom és vallás Northrop Frye munkásságában*

Tari János: *A néprajzi és az antropológiai filmkészítés. Történeti, elméleti és gyakorlati példák*

Buda Béla: *Empátia. A beleélés lélektana*

Németh Dávid: *Pasztorálantropológia*

Szenczi Árpád: *Az ember természete – természetes(en) nevelés. A reformpedagógia egy lehetséges reformja*

Váradi Ferenc: *Vázlatok az óvodai anyanyelvi-irodalmi nevelésről*

Semsey Viktória (szerk.): *Latin-Amerika 1750–1840. A gyarmati rendszer felbomlásától a független államok megalakulásáig*

Bagdy Emőke: *Pszichofitness. Kacagás – kocogás – lazítás*

Balázs Siba: *Life Story and Christian Metanarration. The importance of the research results of narrative identity to practical theology*

Kocziszky Éva: *Antifilozófusok huszonöt időszerű kérdése a kereszténységhez*

Honti László: *Magyar nyelvtörténeti tanulmányok*

Szummer Csaba: *Freud, avagy a modernitás mítosza*

Kun Attila: *A munkajogi megfelelés ösztönzésének újszerű jogi eszközei*

Dr. Lenkeyné dr. Semsey Klára: *Életútinterjú*

Görözdi Zsolt: *Protestáns egyházértelmezés a reformáció századában a jelentősebb egyházi rendtartásokban*

Békési Sándor: *Sisi személyének teológiai portréja*

Török Emőke: *Munka és társadalom. A munka jelentésváltozásai a bérmunkán innen és túl*

Sepsi Enikő: *Pilinszky János mozdulatlan színháza*
Mallarmé, Simone Weil és Robert Wilson műveinek tükrében

Erdélyi Ildikó: *A lélek színháza. A pszichodráma és az önismeret útjai*

Szummer Csaba: *Pszichedelikumok és spiritualitás*

Papp Sándor: *Török szövetség – Habsburg kiegyezés.*
A Bocskai-felkelés történetéhez

Bolyki János: *Teológia a szószéken és a katedrán*

Balogh Eszter: *Túlélési stratégiák a magyar gazdaságban.*
Esettanulmányok a 2000-es évek elejéről.

Falussy Lilla: *Trendek a kortárs olasz drámában*

Bekő István Márton: *Jézus csodáiról szóló elbeszélések Márk evangéliumában*

Kovai Melinda: *Lélektan és politika. Pszichotudományok a magyarországi államszocializmusban 1945–1970*

Spannraft Marcellina: *Költészet és szakralitás*

Németh István: *Műtárgyak a boncteremben. Tanulmányok az orvoslás és a képzőművészet tárgyköréből*

Homicskó Árpád Olivér: *A magyar társadalombiztosítási*
és szociális ellátások rendszere

Fabiny Tibor: *Az eljövendő árnyékai. A figurális tipológiai olvasás*

Szetey Szabolcs: *Adatok a magyar református prédikációs gyakorlat újraértékeléséhez 1784–1878 között*

Balogh Tamás: *Huizinga Noster.*
Filológiai tanulmányok J. Huizinga magyar recepciójáról

Dávid Gyula: *Angol fogalmi idióma szótár.*
Angol idiómák és magyar megfelelőik

Beke Albert: *Gyulai Pál személyisége és esztétikája*

Németh Dávid: *Vallásdidaktika.*
A hit- és erkölcstan tanítása az 5–12. osztályban

Papp Ágnes Klára: *A tér poétikája – a poétika tere. A századfordulós kisvárostól az ezredfordulós terekig a magyar irodalomban*

Lányi Gábor János: *A kálvinizmus nyitánya. Berni zwingliánusok és francia-svájci kálvinisták vitája az egyházfegyelem gyakorlásáról*

Kovács Dávid: *Nemzetfelfogás és történelemszemlélet a 20. századi Magyarországon*

Pap Ferenc: *Az egyházi év*

Farkas Ildikó: *A japán modernizáció ideológiája*

Horváth Csaba: *Megtalált szavak*

Csízy Katalin: *Az ideális vezető politikus alakja a görög-római hagyományban*

Siba Balázs: *Pasztoralteológia*

Indries Krisztián – Szummer Csaba: *Freud és a felkelő nap országa*

Bognárné Kocsis Judit: *A vallás mint forrás Karácsony Sándor nevelési koncepciójában*

Vető Miklós: *Jonathan Edwards gondolatvilága*

Kodácsy Tamás: *A homiletikai aptum – Hogyan lesz a szószéki beszédből igehirdetés?*

Simon Gabriella: *A mentorálás elmélete és gyakorlata. A mentori kompetenciák attitűdelemei és a kezdő pedagógusok mentormodellek szerinti támogatása*

Legéndy Kristóf: *A természettudományos és a teológiai megismerés modelljei*

Siba Balázs: *Pasztorálteológia* (Második, javított kiadás)

Tanulmánykötet

Császár-Nagy Noémi, Demetrovics Zsolt, Vargha András (szerk.): *A klinikai pszichológia horizontja. Prof. dr. Bagdy Emőke 70. születésnapjára készített emlékkötet*

Szávay László (szerk.): *„Vidimus enim stellam eius…"*

Petrőczi Éva, Szabó András (szerk.): *A zsoltár a régi magyar irodalomban*

Horváth Erzsébet, Literáty Zoltán (szerk.): *Történelmet írunk. Tanulmánykötet Ladányi Sándor tiszteletére 75. születésnapja alkalmából*

Czeglédy Anita, Fülöp József, Ritz Szilvia (szerk.): *Inspirationen. Künste im Wechselspiel*

Gudor Botond, Kurucz György, Sepsi Enikő (szerk.):
Egyház, társadalom és művelődés Bod Péter korában

Péti Miklós, Ittzés Gábor (szerk.): *Milton Through the Centuries*

Kendeffy Gábor, Kopeczky Rita (szerk.): *Vallásfogalmak sokfélesége*

Zsengellér József, Trajtler Dóra Ágnes (szerk.): *„A Szentek megismerése ad értelmet." Conferentia Rerum Divinarum 1–2.*

Trajtler Dóra Ágnes (szerk.): *Tan és módszertan. Conferentia Rerum Divinarum 3.*

Pap Ferenc, Szetey Szabolcs (szerk.): *Illés lelkével. Tanulmányok Báthori Gábor és Dobos János lelkipásztori működéséről*

Somodi Ildikó (szerk.): *A mindennapos művészeti nevelés megvalósulásának lehetőségei. Értékközvetítés a művészeti nevelésben*

Dávid István (szerk.): *Merre tovább kántorképzés? Gondolatok egy konferencián – Nagykőrös, 2012. október 5.*

Hansági Ágnes, Hermann Zoltán (szerk.): *Jókai & Jókai*

Dringó-Horváth Ida, N. Császi Ildikó (szerk.): *Digitális tananyagok – oktatásinformatikai kompetencia a tanárképzésben*

Erdélyi Ágnes, Yannick François (szerk.):
Pszichoanalitikus a társadalomban

Fülöp József, Mirnics Zsuzsanna, Vassányi Miklós (szerk.): *Kapcsolatban – Istennel és emberrel. Pszichológiai és bölcsészeti tanulmányok*

Pap Ferenc (szerk.):
Dicsőség tükre. Művészeti és teológiai tanulmányok

Tóth Sára, Fabiny Tibor, Kenyeres János, Pásztor Péter (szerk.):
Northrop Frye 100: A Danubian Perspective

Spannraft Marcellina, Sepsi Enikő, Bagdy Emőke, Komlósi Piroska, Grezsa Ferenc (szerk.): *Ki látott engem? Buda Béla 75*

Komlósi Piroska (szerk.):
Családi életre és kapcsolati kultúrára felkészítés

Dringó-Horváth Ida, Fülöp József, Hollós Zita, Szatmári Petra, Czeglédy Anita, Zakariás Emese (szerk.): *Das Wort – ein weites Feld*

Fülöp József (szerk.): *A zenei hallás*

József Fülöp, Szilvia Ritz (Hg.): *Inspirationen II*

Tóth Sára, Kókai Nagy Viktor, Marjai Éva, Mudriczki Judit, Turi Zita, Arday-Janka Judit (szerk.): *Szólító szavak. The Power of Words. Tanulmányok Fabiny Tibor hatvanadik születésnapjára*

Lázár Imre, Szenczi Árpád (szerk.): *A nevelés kozmológusai. Kodály Zoltán, Karácsony Sándor és Németh László megújuló öröksége*

Erdélyi Erzsébet, Szabó Attila (szerk.): *A hit erejével. Pedagógiai tanulmányok*

Makkai Béla (szerk.): *A Felvidék krónikása. Tanulmányok a 70 éves Popély Gyula tiszteletére*

Farkas Ildikó, Sági Attila (szerk.): *Kortárs Japanológia I.*

Gér András László, Jenei Péter, Zila Gábor (szerk.): *Hiszek, hogy megértsem*

Simon-Székely Attila (szerk.): *Lélekenciklopédia. A lélek szerepe az emberiség szellemi fejlődésében*

Papp Ágnes Klára, Sebők Melinda, Zsávolya Zoltán (szerk.): *Nemzet sors identitás*

Vassányi Miklós, Sepsi Enikő, Voigt Vilmos (szerk.): *A spirituális közvetítő*

Julianna Borbély, Katalin G. Kállay, Judit Nagy, Dan H. Popescu (eds.): *English Language & Literatures in English 2014*

Sepsi Enikő, Lovász Irén, Kiss Gabriella, Faludy Judit (szerk.): *Vallás és művészet*

Bubnó Hedvig, Horváth Emőke, Szeljak György (szerk.): *Mítosz, vallás és egyház Latin-Amerikában. A Boglár Lajos emlékkonferencia tanulmánykötete*

Czeglédy Anita, Sepsi Enikő, Szummer Csaba (szerk.): *Tükör által – Tanulmányok a nyelv, kultúra, identitás témaköréből*

Méhes Balázs (szerk.): *Lelki arcunk. Tanulmányok Szenczi Árpád hatvanadik születésnapja alkalmából*

Spannraft Marcellina, Korpics Márta, Németh László (szerk.): *A család és a közösség szolgálatában. Tanulmányok Komlósi Piroska tiszteletére*

Horváth Csaba, Papp Ágnes Klára, Török Lajos (szerk.): *Párhuzamok, történetek. Tanulmányok a kortárs közép-európai regényről*

Anka László, Kovács Kálmán Árpád, Ligeti Dávid, Makkai Béla, Schwarczwölder Ádám (szerk.): *Natio est semper reformanda. Tanulmányok a 70 éves Gergely András tiszteletére*

Fülöp József, Mészáros Márton, Tóth Dóra (szerk.): *A szél fúj, ahová akar. Bölcsészettudományi dolgozatok*

Zsengellér József, Kodácsy Tamás, Ablonczy Tamás (szerk.): *Felelet a mondolatra. Tanulmányok a 60 éves Bogárdi Szabó István tiszteletére*

Borgulya Ágnes, Konczosné Szombathelyi Márta (szerk.): *Vállalati kommunikációmenedzsment*

Szávay László, Gér András László, Jenei Péter (szerk.): *Hegyen épült város. Válogatás a Fiatal Kutatók és Doktoranduszok Nemzetközi Teológuskonferencián elhangzott előadások anyagából.*

Wakai Seiji, Ságai Attila (szerk.): *Kortárs Japanológia II.*

Sepsi Enikő: *Kép, jelenlét, kenózis a kortárs francia költészetben és Valère Novarina színházában*

Sepsi Enikő, Tóth Sára (szerk.): *Mellékzörej. Írások Visky András hatvanadik születésnapjára*

Erdélyi Erzsébet, Szabó Attila (szerk.): *Az üzenetjét, azt kell megbecsülni. Tanulmányok Barabás László hetvenedik születésnapja alkalmából*

Kendeffy Gábor, Vassányi Miklós (szerk.): *Istenfogalmak és istenérvek a világ filozófiai hagyományaiban*

Bárdi Nándor, Éger György (szerk.): *Magyarok Romániában 1990–2015. Tanulmányok az erdélyi magyarságról*

Zila Gábor: *Hitetek mellé tudományt*

Kiss Paszkál, Tóth Dóra (szerk.): *Ubi dubium, ibi libertas. Tudományos diákköri dolgozatok*

Julianna Borbély, Borbála Bökös, Katalin G. Kállay, Judit Nagy, Ottilia Veres, Mátyás Bánhegyi and Granville Pillar (eds.): *English language & literatures in English 2016*

Spannraft Marcellina, Tari János (szerk.): *A kultúraátörökítés médiumai*

László Emőke (szerk.): *A szolgálat ékessége. Tanulmányok Fruttus István Levente hetvenedik születésnapja alkalmából*

Spannraft Marcellina (szerk.): *Tertium datur. Tanulmányok Lázár Imre tiszteletére*

Bárdi Árpád, Gombos Norbert, Tóth Etelka (szerk.):
Kisgyermeknevelés a 21. században

Lányi Gábor János (szerk.): *A reformáció örökségében élve.*
A reformáció hatása a teológiai oktatásra

Dróth Júlia (szerk.): *Gépiesség és kreativitás a fordítási piacon*
és a fordításoktatásban

Lator László Iván (szerk.): *Sorsok és missziók a politikai változások tükrében*

Literáty Zoltán (szerk.): *Quoniam tecum est fons vitae in lumine tuo*
videbimus lumen. Teológiai témák a 65 éves Kocsev Miklós tiszteletére

Fehér Pálma Virág, Kövesdy Andrea, Szemerey Márton (szerk.):
Testképek a gyógyításban. A test mint eszköz és referenciapont

Zonda Tamás, Bozsonyi Károly, Moksony Ferenc (szerk.):
Az öngyilkosság szociológiája

Fóris Ágota, Bölcskei Andrea (szerk.): *A szabványosítás fordítási*
és terminológiai vonatkozásai

Kövi Zsuzsanna, Mirnics Zsuzsanna, Reinhardt Melinda (szerk.):
Lélek(sz)árnyak

Czentnár Simon, Nagy József, Nagy Levente (szerk.):
Beszéljétek el dicsőségét
Conferentia Rerum Divinarum 6.

Kovács Tibor (szerk.): *Tehetségek hite. Hallgatói tanulmánykötet*

Zalatnay István: *Félúton. A liturgia erőterében született írások*

Farkas Ildikó, Sági Attila (szerk.): *Kortárs Japanológia III.*

Furkó Péter (szerk.): *Az önképtől és tudásreprezentációtól*
a hatalommegosztásig. Tanulmányok a bölcsészettudomány,
a hittudomány és a jogtudomány területeiről

Olay Csaba, Schmal Dániel (szerk.): *Értelem és érzelem az európai*
gondolkodásban. Tanulmányok a 60 éves Boros Gábor tiszteletére

Fóris Ágota, Bölcskei Andrea (szerk.):
Dokumentáció, tartalomfejlesztés és szakírás

Dróth Júlia (szerk.): *Korpusz és kontrasztivitás*
a szakfordítás oktatásában és gyakorlatában

Sebők Melinda, Papp Ágnes Klára, Török Lajos (szerk.):
Kötelezők emelt szinten Balassitól Borbély Szilárdig

Béres István - Korpics Márta (szerk.):
Ha kiderül az alkony

Németh István: *El Grecóról és más mesterekről.*
Stíluskritika – ikonográfia – műgyüjtéstörténet

Karasszon István: *Az ószövetségi teológia történetéhez*

Műfordítás, forrás

Giovanni Pico della Mirandola:
Benivieni neoplatonista versének kommentárja
(Fordította: Imregh Monika)

Alice Zeniter: *Szomorú vasárnap, avagy a semmi ágán* (Fordította: Kovács Veronika, szerkesztette és a bevezetőt írta: Sepsi Enikő)

Szent Ágoston: *Írások a kegyelemről és az eleve elrendelésről* (Fordította, válogatta és a bevezetőt írta: Hamvas Endre)

Paul Claudel: *Délforduló* (Fordította: Székely Melinda)

Veerle Fraeters, Frank Willaert, Louis Peter Grijp (szerk.):
Hadewijch: Dalok
(Fordították: Daróczi Anikó, Rakovszky Zsuzsa)

Landauer Attila (szerk.): *A Kárpát-medencei cigányság és a keresztyén egyházak kapcsolatának forrásai (1567–1953)*

Yves Bonnefoy: *Hollán Sándor.*
Harminc év elmélkedései, 1985–2015 (Fordították: Gulyás Adrienn, Kovács Krisztina, Kovács Veronika, Makádi Balázs, Sepsi Enikő)

Vassányi Miklós (írta, fordította, szerkesztette): *Szellemhívók és áldozárok. Sámánság, istenképzetek, emberáldozat az inuit (eszkimó), azték és inka vallások írásos forrásaiban.*

Rauni Magga Lukkari, Inger-Mari Aikio: *Örökanyák - Világlányok. Számi versek* (Fordították: Domokos Johanna, Németh Petra)

Rajvinder Singh: *Hat szemmel. Német, angol, hindi és pandzsábi versek*
(Szerkesztette: Domokos Johanna. Fordították: Tibold Katalin, Széles Beáta, Domokos Johanna)

Günter Eich: *Katharina és más elbeszélések*
(Válogatta, szerkesztette és az utószót írta: Fülöp József)

Augusto Monterroso: *A fekete Bárány és más mesék*
(Szerkesztette, fordította és az utószót írta: Fülöp József)

Michael Galuske: *A szociális munka módszerei. Bevezetés*
(Fordította: Udvari Kerstin, szerkesztette: Bányai Emőke)

Titus Livius: *Róma története a város alapításától,*
I. kötet (I–XX. könyv) (Kis Ferencné és Muraközy Gyula fordítását és a jegyzeteket átdolgozta: Kopeczky Rita, a fordítást az eredetivel egybevetette: Kocsisné Csízy Katalin, Ligeti Dávid Ádám, Bajnok Dániel)

Titus Livius: *Róma története a város alapításától,*
II. kötet (XXI–XXX. könyv) (Muraközy Gyula fordítását és a jegyzeteket átdolgozta: Kopeczky Rita, a fordítást az eredetivel egybevetette: Czerovszki Mariann, Tamás Ábel)

Titus Livius: *Róma története a város alapításától,*
III. kötet (XXXI–XL. könyv) (Muraközy Gyula fordítását és a jegyzeteket átdolgozta: Kopeczky Rita, a fordítást az eredetivel egybevetette: Berkes Lajos, Kozák Dániel, Ligeti Dávid Ádám)

Titus Livius: *Róma története a város alapításától,*
IV. kötet (XLI–CXLII. könyv) (Muraközy Gyula fordítását és a jegyzeteket átdolgozta: Kopeczky Rita, a fordítást az eredetivel egybevetette: Ligeti Dávid Ádám, Takács Levente, Bajnok Dániel)

Tzveta Sofronieva: *A föld színe* (Szerkesztette: Domokos Johanna)

Bangha Imre: *Költők és koldusok között. Indiai napló 2001–2017*
(Szerkesztette: Frazer-Imregh Monika)

L'Harmattan France-Hongrie, Collection Károli

Anikó Ádám, Enikő Sepsi, Stéphane Kalla (szerk.):
Contempler l'infini

Tibor Fabiny, Sára Tóth (eds.): *The King James Bible (1611–2011). Prehistory and Afterlife*

Katalin G. Kállay, Mátyás Bánhegyi, Ádám Bogár,
Géza Kállay, Judit Nagy,
Balázs Szigeti (eds.): *The Arts of Attention*

Katalin G. Kállay, Nóra D. Nagy, Elizabeth Walsh, Ádám Fónai, Béla Erik Haga, Péter Káplár, Krisztina Milovszky, Gergely Molnár, Dorina Obrankovics, Tamás Szanyi (eds.): *This is Just to Say. A Collection of Creative Student-Responses*

György Kurucz (ed.): *Protestantism, Knowledge and the World of Science*

Brigitte Geißler-Piltz, Éva Nemes, Sijtze de Roos (eds.): *Inspire and Be Inspired. A Sample of Research on Supervision and Coaching in Europe*

Johanna Domokos: *Endangered literature. Essays on Translingualism, Interculturality, and Vulnerability 2018*

Tímes Kovács: *Code-Switching and Optimality*

Viktória Semsey (ed.): *National Identity and Modernity 1870–1945. Latin America, Southern Europe, East Central Europe*

Anita Czeglédy, Géza Horváth (Hg.): *Inspirationen III – Wege*

Renáta Raáb: *Austria's Schleswig Policy 1848–1852*

Anita Czeglédy, Anikó Szilágyi-Kósa (Hg.) unter Mitarbeit von József Fülöp: *Geopoetische Reisen in Mitteleuropa. Studien zur Sprache und Kultur*

Krisztina Kovács: *La relation poétique entre l'oeuvre d'Yves Bonnefoy et celle de Paul Celan*

Ádám Bethlenfalvy: *Living through extremes is process drama*

Jegyzet

Nádor Orsolya:
Tanított (anya)nyelvünk

Hegedűs Rita:
Élő magyar grammatika
– a nyelvhasználat rendszerének rövid áttekintése

Csontos Nóra, Dér Csilla Ilona:
Pragmatika a magyyar mint idegen nyelv oktatásában

H. Varga Márta (szerk.):
Segédlet nyelvi formák és grammatikai funkciók tanulmányozásához (elmélet és gyakorlat)

M. Pintér Tibor, Wéber Katalin:
Mérés és értékelés: magyar mint idegen nyelv

Holló Dorottya:
Kultúra és interkulturalitás a nyelvórán

Schmidt Ildikó:
A magyar írás és olvasás tanítása - Az alfabetizálás folyamata

Lannert Keresztély István: *Bevezetés a drámapedagógiába*

L'Harmattan France
5-7 rue de l'Ecole Polytechnique
75005 Paris
T.: 33.1.40.46.79.20
Email: diffusion.harmattan@wanadoo.fr

L'Harmattan Italia SRL
Via Degli Artisti 15
10124 TORINO
Tél: (39) 011 817 13 88 / (39) 348 39 89 198
Email: harmattan.italia@agora.it

Cover design: László Kára
Layout editor: A. Aliz Molnár
Printing: Prime Rate Ltd.
Manager: Péter Tomcsányi