Philosophy of Engineering and Technology 3

Joseph C. Pitt

Doing Philosophy of Technology

Essays in a Pragmatist Spirit

Springer

Doing Philosophy of Technology

Philosophy of Engineering and Technology

VOLUME 3

For further volumes:
http://www.springer.com/series/8657

Joseph C. Pitt

Doing Philosophy of Technology

Essays in a Pragmatist Spirit

Springer

Joseph C. Pitt
Department of Philosophy
Virginia Tech
24061-0126 Blacksburg
VA, USA
jcpitt@vt.edu

ISSN 1879-7202 e-ISSN 1879-7210
ISBN 978-94-007-0819-8 e-ISBN 978-94-007-0820-4
DOI 10.1007/978-94-007-0820-4
Springer Dordrecht Heidelberg London New York

Library of Congress Control Number: 2011922764

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Dedicated to My Mother, Miriam Pitt

Introduction

This book does not run a straight course from beginning to end. It hunts; and in the hunting, it sometimes worries the same raccoon in different trees, or different raccoons in the same tree, or even what turns out to be no raccoon in any tree. It finds itself balking more than once at the same barrier and taking off on other trails. It drinks often from the same streams, and stumbles over some cruel country. And it counts not the kill but what is learned of the territory explored

Nelson Goodman, *Ways of Worldmaking*. Indianapolis: Hackett Publishing Company, 1978, p. ix.

Doing Philosophy of Technology

The title of this collection raises several questions: (1) what is philosophy? (2) What is philosophy of technology? (3) What is it to do philosophy of technology? Let's look at them in order.

If the answer to the first is "a body of work – a bunch of books on the shelf", then we may be asking the wrong sort of question. Yes, there are hundreds, nay thousands of philosophical treatises in which writers over a period of at least 2500 years have discoursed on a variety of topics. However, I want to argue that is not Philosophy. That may be the product of having philosophized, but philosophy is not the bound result of over two millennia of thinking about certain kinds of things. There are several reasons for making this claim, one has to do with the kinds of things with which philosophers are allegedly to be concerned, i.e., the questions of the perennial philosophy. More on this below. But more importantly, conceiving of philosophy as a set of books, or essays, is to see it ossified.

Another reason to reject the idea that philosophy is some set of results to be codified and put on a shelf has to do with whom we baptize as philosophers. When the year 2000 was approaching, everyone was creating lists of the greatest this and that of the twentieth century, so I thought I would try to find out who my colleagues thought were the greatest philosophers of the twentieth century. The usual suspects surfaced (all restricted, interestingly enough, to the Western philosophical traditions): Russell, Wittgenstein, Heidegger, Dewey, Quine. And I suspect the criteria used in selecting these individuals had something to do with who it is assumed would be footnotes when the history of philosophy is written 200 years from now. The list is interesting for several reasons, for who is on it and who is not. Why is Quine there?[1] He will certainly fade from view by midcentury for his contributions are in retrospect miniscule. By the way, contributions to what? My guess here is: contributions to solving the questions of the perennial philosophy – but this is suspect as a starting point from the beginning – more anon – if you feel pressed for an answer right now see "Against the Perennial" in this volume.

Of the philosophers cited above there is only one who actually made some kind of difference in the way we live: John Dewey. For better or worse his ideas had a major impact on educational thought and practice in the United States and how we teach our children. None of the others in that list actually had an impact on how we live.

And why is having an impact on how we live important to philosophy? Well, philosophy is by definition "love of wisdom." This doesn't mean a philosopher is wise, just that he or she loves wisdom. So what is wisdom? I don't know. But I do know that it is different from truth. We do science to arrive at the truth – for truth has something to do with what is out there and how it works. But wisdom, that which philosophers love, is not exclusively about what is out there. Wisdom is about how to live in the out there. But if philosophy really is the love of wisdom, then if a philosopher is a lover of wisdom and wisdom has to do with living and living is an activity, then what philosophers do as philosophers is an activity. And, I suggest, that activity is the seeking of wisdom. To put it more forcefully, seeking wisdom ought to be the regulative ideal by which philosophers should determine the appropriateness of their activities.

Why do we do philosophy and in what does doing philosophy consist? Shifting the topic from what is philosophy to what is it to do philosophy is a major change of emphasis. And this shift comes from my deep commitment to the belief that ideas have the power to change the way we live our lives and see ourselves. For while seeking wisdom, the philosopher's stock in trade is the set of ideas he or she proposes and defends regarding how to live well. Philosophers should not be splitting hairs to make abstract points of no interest to anyone other that the five people with whom they went to graduate school. They should be concerned with how best to make it through the muddle that is life. Those who know how to do that are wise persons – most of them are grandmothers.

Which brings me to those who are not on the canonical list of great philosophers of the twentieth Century, but who ought to be. These are the individuals who will be remembered far into the future for the impact their ideas and actions had on how we live. I offer three twentieth Century examples: Lenin, Mao and Gandhi, each because he thought he could bring about a different way of life; one failed, one succeeded, and at this point in time it is not clear how to characterize what Mao wrought. However, in each case they were actively engaged in trying to bring about a change in actuality with a vision for a better future. The implication here is that

[1] Willard van Orman Quine, 1908–2000, spent his professional career at Harvard. He was primarily a logician who also worked in the analytic mode in epistemology and philosophy of science.

philosophy is involved in creating ideas that motivate people to do something and that it is fundamentally normative. Ideas are the most important and powerful things we create. They have the power to bring down kingdoms, to mobilize millions, to give hope to the downtrodden and to inspire the rest of us. Great ideas move people to act. We have no better example than our own Declaration of Independence. But this is not a new idea, consider the following:

> The ideas of economists and political philosophers, both when they are right and when they are wrong, are more powerful than is commonly understood. Indeed the world is ruled by little else. Practical men, who believe themselves to be quite exempt from any intellectual influence, are usually the slaves of some defunct economist. Madmen in authority, who hear voices in the air, are distilling their frenzy from some academic scribbler of a few years back. I am sure that the power of vested interests is vastly exaggerated compared with the gradual encroachment of ideas. Not, indeed, immediately but after a certain interval; for in the field of economic and political philosophy there are not many who are influenced by new theories after they are twenty-five or thirty years of age, so that the ideas which civil servants and politicians and even agitators apply to current events are not likely to be the newest. But soon or late, it is ideas, not vested interest, which are dangerous for good or evil. (From John Maynard Keynes. *General Theory of Employment, Interest, and Money.* Chapter 24, Section V, Concluding Notes.[2])

But not all ideas have this power. I seriously doubt if Quine's "to be is to be the value of a bound variable" is a dangerous idea. But the idea that "all men are created equal" is. So is the idea that all men and women are entitled to life, liberty and the pursuit of happiness. Dangerous ideas threaten the *status quo*. But let us be clear how the do this: they instill doubt – they force us to question whether the values of the *status quo* really capture the best ideals for how I should live my life; they challenge what we have been taught; worse yet, they can force us to think about doing something.

Returning to what philosophers do (rather than what philosophy is). Given the direction of the above remarks, it is reasonable to assume that the conclusion should be that somehow philosophers have some kind of priority on the creation of dangerous ideas. But no – if that were to be the conclusion, then the moral would be to eliminate philosophers – for disrupting the *status quo* with a constant barrage of new, threatening ideas would surely destabilize society. And while that may be a good thing every once in a while, it cannot be a permanent state of affairs – witness the chaos resulting from Chairman Mao's continuous revolution. On the other hand, if challenging the *status quo* every once in a while is a good thing, how are we to insure that this can happen? For one thing, you don't do it by eliminating philosophers. Actually, the very notion that philosophers as we know them, are dangerous is amusing. Consider the image of philosophers in popular culture: a bunch of fuddy-duddies who will argue on and on about anything, but surely not anything of importance.

At this point it might be worth switching emphasis one more time. We really don't know what philosophy is and we really don't believe the vast majority of philosophers create dangerous ideas, at least dangerous ideas are not their intended

[2] I thank Daryl Farber for the reference.

stock in trade. And it is even the case that those who are thought of as having created dangerous ideas, e.g., Martin Luther King, jr., are not thought of primarily as philosophers; they are seen as activists, rabble-rousers, perhaps, but *philosophers*?

Today there are some grounds for thinking of philosophers as members of a profession – and the one thing the members of this profession have in common is that they almost universally are or have been teachers, primarily in colleges and universities, and they teach in one way or another the history of philosophy – that 2500 year long conversation about a bunch of ideas, first articulated by Greek thinkers roughly 2500 years ago, that some have maintained are the topics of perennial questions about the human condition. Why the ancient Greeks should be consecrated with having had this insight is not clear and why we continue to think these ancient thinkers had some innate grasp on the fundamental questions of life and reality is even less clear. More importantly, to think that the questions asked 2500 years ago are the same questions we ask today, even though their formulations look similar, is to ignore the context in which such questions are asked. The citizens of fifth century B.C.E. Athens had very different expectations of what kind of answer was appropriate to the question "what is the good life" than we do today. And if the answer to the question will be expected to be different, the question must also be different, even if it looks the same on the surface.

Even those who think they are teaching so-called contemporary philosophy teach the history of philosophy, for the ideas they are discussing, the ideas in general of other teachers of philosophy, are themselves formulated as responses to claims made in that tradition of thinking about a certain limited domain of ideas first articulated by the ancient Greeks. And to the extent that these topics, ideas, domains, etc., have no bearing on how we live our lives today, then our students are correct in asking why they should bother trying to understand them. Even in the domain of ethics, the theories discussed have nothing to do with how people actually think about what they should do – instead philosophy teachers have abstract discussions about, for example, the greatest good for the greatest number – really gut-wrenching stuff.

But teachers of philosophy also do something else beyond wallowing in the restricted domain of the past – sometimes they don't even know they do it. It is what makes it possible for someone to challenge the *status quo* when that is needed. One thing we philosophy teachers do, or should do, is to teach our students to challenge assumptions. We keep asking them "why do you think that?" or "how do you know that?" In asking these questions we are forcing them to consider their beliefs, their reasons for their beliefs and the reasonableness of both reasons and beliefs. From challenging their beliefs it is but a short step to challenging the beliefs and assumptions of others. I jokingly warn by students in my introductory course that they will soon find their friends avoiding them because if they have caught the bug, they will be asking "why?" of everyone and everything. That is how you create revolutionaries. Idea-generated change comes from challenging the *status quo*. But you can't do that unless you are in a position to recognize that that idea or belief can be challenged.

Having arrived at the view that philosophy is the love of wisdom and that the love of wisdom requires we do something and the something we are required to do is challenge assumptions and the *status quo*, it is time to turn to our second question: what is the philosophy of technology? If we continue the line of thought developed above, then the philosophy of technology should be understood as seeking wisdom about how to deal with the technologies we create and the world they help us make. This is probably not the place to entertain a prolonged discussion of the nature of technology, but something does need to be said. In *Thinking About Technology* I argued that we ought to think of technology as *humanity at work (p. 11)*. The argument that led to that conclusion can be found there so I won't rehearse it here. The result of that argument is that the philosophy of technology should be understood as that activity in which we seek wisdom about how and why people do what they do when working. What are the assumptions people employ when they act and make things and use processes and systems and tools and are these assumptions justified? We also need to consider the consequences of those actions and what they reveal about what we think we know and our values and goals.

Which brings us to *Doing Philosophy of Technology.* The essays in this volume grew out of a deep dissatisfaction with the *status quo* regarding our thinking about technology. Beginning in the late nineteenth century there was a popular movement that itself grew out of what appeared to be, in the Western world, a reaction to what was seen as an unabashed embracing of all things scientific and technological and the challenges the brave new world science and technology promised posed for traditional values. This reaction can be seen in many forms but to mention just two we see it in Mumford's work[3] and in such films as Fritz Lang's 1927 *Metropolis.*[4] Likewise Mary Shelley's *Frankenstein* helped represent, if not create, a mindset of fear regarding the results of science. It is by way of response to this blind negative reaction to the fruits of scientific and technological work that started me down the path that lead to the present work. I have told the story before, but it bears repeating in this context. My interest in the philosophy of technology began when I was following a pickup truck around the drill field on the Virginia Tech campus. It sported a bumper sticker that read "Guns don't kill, people kill". This was supposed to express a sentiment favoring less government control over gun ownership and support of a particular reading of the 2nd Amendment to the Constitution of the United States (one with which I disagree). However, there is a sense in which what the sticker claimed is true and another in which it is sort of not true. And it was in an attempt to sort out this ambiguity that I turned to the work of the then contemporary practitioners of the philosophy of technology, only to discover, much to my dismay, that much of that work was a continuation of the sentiments expressed earlier in the century: science is to be feared and technology is harmful. Clearly, putting it in is this stark form is an exaggeration, but it helps to set the stage for the second set of developments that led me down this path.

---

[3]The Myth of the Machine by Lewis Mumford. New York, Harcourt, Brace & World 1967–1970. [1st ed.]

[4]This film can be read on many levels – it has a clear political message as well as presenting a frightening pictures of the world controlled by men who control machines. It is the latter aspect that interests me here.

The university where I have taught for the past 40 years, Virginia Tech, has a strong technological tradition, from agriculture to engineering to architecture. As a teacher, I am always trying to find a way to make philosophy relevant or even interesting to my more technologically and scientifically oriented students. This can sometimes be frustrating. It was clear that telling these students that the areas in which they were majoring were responsible for all the ills that beset society (even if it were true, which it is not, more anon) was not the way to win their hearts. But it was equally clear that showing them that there were philosophically interesting issues associated with the topics that captured their imaginations was key to having them come back for that second course, since all teachers know that capitalizing on what interests a student is essential to getting them excited about learning.

What I discovered fairly quickly was that the students were a lot smarter than the philosophers. Philosophers of technology at that time (1970s and 1980s) tended to make general claims about the evils of technology and then focus on a case study or two, concluding that the generalization is true (See “The dilemma of case studies”). The students saw the flaw in this form of reasoning quickly and suggested we start with specific examples of a technology and see if we could assess it without prior assumptions of good or evil. What we found out together was that you need to know a lot about a technology before you can begin forming normative judgments and even then those judgments have to be tempered because technologies don’t operate in a vacuum, they are always socially, geographically, and historically contextualized and what the context is makes a huge difference on which normative conclusions we draw. But it is not enough to say you need to know a lot. The very concept of knowledge here becomes problematic. For it is not clear what kind of knowledge we need to have.

So, finally, we have come to something of interest: is technological knowledge *sui generis* or is it a special form of some other kind of knowledge? In other words, there are other philosophical issues to talk about than the normative ones. The epistemological questions naturally lead us to metaphysical questions regarding the ontological character of technological artifacts, a topic addressed by Heidegger, but in such a general and obtuse fashion that it is hard to get a grip on the issue. But the situation changes when you ask the straightforward question: do technical artifacts exist in the same sense as, e.g., trees? One clear difference is that artifacts are designed – to argue that trees are also designed is to interject a theological element into the discussion and essentially change the topic. Thus, we can conclude that the philosophy of technology is about more than normative issues, there are also epistemological and metaphysical questions.

Which brings us to the final of our three original questions, “What does it mean to do philosophy of technology?” If doing philosophy is seeking wisdom, and if the philosophy of technology concerns the normative, epistemological and metaphysical issues raised by the world we have built, then doing philosophy of technology should mean something like: seeking wisdom with respect to the world we have built, by seeking out all the accompanying epistemological, normative and metaphysical questions that world and what we do to create it and act in it raises. Wow, that’s a mouthful! But it can be simplified. Doing philosophy of technology involves trying to figure out how to live in the world we have created. It will involve examining normative claims about what we ought to do and not do. It will involve figuring out what it means to act on our knowledge and in what that knowledge consists, recognizing that what we mean by "knowledge" changes as we create more and different devices to help us explore and exploit nature and the universe. It will involve wrestling with the contrast between the made world and the "natural" world. And much more. But at its heart it is to think about what we ought to do today in light of what we have done in the past and how that bears on what we ought to do about tomorrow.

The essays in this volume represent a spectrum of my writings in the philosophy of technology over several decades. The more I have toiled to understand complex technologies, how they work, what they actually do and our expectations of them, the more convinced I am that much of contemporary philosophy doesn't address the hard core questions of living, knowing and being (pace Heidegger). My earlier not-so-positive comments about Quine and the non-threatening nature of his criterion for existence was not just my being snide. I truly believe philosophy must make a difference in the way we live or it will fade into irrelevance. Some of the questions and mechanisms philosophers have asked and devised are really clever and exhibit the kind of precision and focus we have come to associate with really clever people. But in the end, what difference does it make? And to quote Williams James.

> There can be no difference anywhere that doesn't make a difference elsewhere – no difference in abstract truth that doesn't express itself in a difference in concrete fact and in conduct consequent upon that fact, imposed on somebody, somehow, somewhere, and somewhen. (James 1907, 1955, p. 45.)

And so I put forth in an open fashion my pragmatist stance. Pragmatism, with its twin emphases on consequences and action is the only philosophical position I know that forces philosophers to own up to the real world affects of their rarified speculations, their logic chopping, and their hair-splitting avoidance of the real world and its problems. The fact that pragmatists insist on there being a difference in conduct resulting from philosophical ruminations does not imply approval of those results. No, once we have thought about how these idle thoughts would play out in the world of human activity the hard work begins, i.e., trying to figure out what we should do, now that we understand this or that. And, here again, pragmatism, with its emphasis on concrete consequences and action is the only philosophical position that forces us to be actors and face up to the results of our armchair musings. For ideas have power and can be dangerous only we if do something about them.

The book is divided into four parts – but they are somewhat arbitrary. For example, Part I is about ethics and society and values, but contains a piece about nanotechnology and nanotechnology has its own area in Part IV. Nevertheless, I think there is an order to things and what I want to do is start with the issues I have raised in this Introduction about the need for philosophy to be concerned about how we live in the world we have made and how we deal with nanotechnology is now part of that problem. Part II deals with methodological issues. This is a bit more tricky since "methodology" can mean a variety of things. Here I cover a variety of methodological issues from why we, as philosophers, feel the need to be stuck in a 2500 year old box to what constitutes a technological explanation. Part III concerns a set of questions raised by engineering and the problems besetting design. There is a kind of agreement among those of us who work in the area that design is at the heart of the philosophical problems surrounding engineering and that the disciplines that constitute engineering are fundamental to our understanding of technology. Finally, I end with some papers on nanotechnology. Nano is the new frontier in science and technology and there are some really interesting philosophical issues that it raises. I do not pretend to deal with all of them, but the three I deal with in this section are, in my mind, fundamental to our understanding of the kinds of problems the future holds for us.

There are two themes that continuously appear and reappear in these essays. The first is Wilfrid Sellars' conception of the aim of philosophy; briefly: to see how things in the broadest sense hang together in the broadest sense. It is a prescription that has influenced me in more ways than I am sure I know. Second, There is the conception of knowledge I have relied on, knowingly and unknowingly, over the years. It involves the idea of a feedback loop and comes out of much earlier attempts to make sense of the notion that science is self-correcting. As I worked with it and refined it I found that it applied to much more.

The debt I owe to so many people cannot be fully explained in these short pages. But there are six that are especially important and need to be publically thanked. First, without my Donna beside me these many decades, lending her quiet support and encouragement none of this would have transpired. Second, despite his grumbly manner and incessant needling or maybe because of it, Paul Durbin may be the other person to whom I am most indebted. His presence as the founder of modern American philosophy of technology and his continuous support of my work despite disapproving of most of it, has made these efforts of mine personally rewarding. Third, I must thank my students, they serve not only as inspiration but as my best and most forceful critics. Fourth, David Gordon for his editorial assistance and fifth, Nick Perich for putting the whole package together in a coherent form. Finally, although I constantly rail against those who make abstract concepts into real things, I must thank Virginia Tech. The use of the name of the university is merely shorthand for the many Hokies who have supported my efforts and made my life here worthwhile.

References

James, W. 1907, 1955. *Pragmatism*. Cleveland, OH: The World Publishing Company.
Lang, F. 1927. *Metropolis*. Berlin: Universum Film A.G.
Mumford, L. 1967–1970. *The Myth of the Machine*. New York, NY: Harcourt, Brace & World.

Contents

Part I
Society, Ethics and Values

Chapter 1
Human Beings as Technological Artifacts

There are basically two types of people: active and passive. This assumption is introduced in order to open up a way of thinking about the relations between people and the technologies that occupy their world. The basic point is that you can go through life becoming who you will eventually be (a) by simply responding to the various impacts and stimulations you receive, or (b) by attempting to make yourself into the type of person you think you would like to become. To do the former requires no work on your part. To do the latter requires a lot of work, primarily by way of finding out the limits of your capacities and the nature and range of the options open to you, uncovering the prerequisites for achieving the steps you need to take, allowing for errors and the means to correct them, etc. There is an important sense, then, in which the active person attempts to design the person they want to become in much the same way we design an artifact, thereby becoming an artifact his or herself. To elaborate this idea, I begin by introducing some ideas of John Dewey. I then give a sketch of the design process as we find it in use in engineering. Finally, following Dewey's ideas about the nature of education, and taking the design process as a metaphor, I show how we need to educate our students better for a world of complexity unlike anything we have hitherto experienced. The motivation here is quite straightforward. We live in a technological world that at least appears to be wrapping us up in electronics and other technologies without asking for our consent. The ability to select those technologies I want to be associated with is important to who I am and who I will become. Thus it is important to be able to say, for example, I don't want a cell phone because I treasure my privacy and independence. So far these kinds of decisions have been displayed in the context of resisting technology. But to accept that construal is to accept a picture of human beings as primarily passive beings. The key question is can we be active beings in the contemporary environment?

Joseph C. Pitt, "Human Beings as Technological Artifacts", 2006, Palgrave MacMillian, reproduced with permission of Palgrave MacMillian.

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_1,

1.1 John Dewey and the Aims of Education

In a piece entitled "Education and Growth," Dewey makes the following claim:

> Education is not infrequently defined as consisting in the acquisition of those habits that effect an adjustment of an individual and his environment. The definition expresses an essential phase of growth. But it is essential that adjustment be understood in its active sense of *control* of means of achieving ends. (p. 494)

The emphasis is on the achieving of certain habits, not beliefs. This stress on habits is a key feature of pragmatism. It comes directly from C.S. Peirce, the founder of pragmatism. Thus Peirce says first,

> The essence of belief is the establishment of a habit; and different beliefs are distinguished by the different modes of action to which they give rise. (29)

He completes the thought this way:

> . . .the identity of a habit depends on how it might lead us to act, not merely under such circumstances as are likely to arise, but under such as might possibly occur, no matter how improbable they may be. For what the habit is depends on *when* and *how* it causes us to act. (30)

In his insistence of the developments of habits, Dewey is straightforwardly in the pragmatist tradition – likewise with respect to the tie of habits to action. Pragmatists are concerned with how we act in the world and why we act the way we do. Thus, according to Dewey,

> The savage is merely habituated; the civilized man has habits which transform the environment. The significance of habit is not exhausted, however, in its executive and motor phase. It means formation of intellectual and emotional disposition as well as an increase in ease, economy, and efficacy of action. Any habit marks an *inclination* – an active preference and choice for the conditions involved in its exercise. (pp. 494–495)

Finally he adds,

> Above all, the intellectual elements in a habit fixes the relation of habit to varied and elastic use, and hence to continued growth. (p. 495)

So, in the end, Dewey argues for the development of habits that allow for continued growth, which is the heart of education. We aim to inculcate in the young those habits that will give them the capacities to acquire knowledge on their own and then to evaluate it and to use it to achieve their goals, adjusting to circumstances and learning from their experiences.

The under-examined component here for Dewey is the selection and evaluation of goals. We do a reasonable job of injecting values into our students, but we rarely give them the means by which to select and evaluate goals or to consider their goals in the light of their values. Often they will choose a goal because everyone else is choosing that goal, for example, to make lots of money. But they usually haven't thought through the conflict that arises, for example, between being happy and pursuing a fortune. It may just be, and I will put this forth as an unargued for premise, that, for the long run, the most important thing we can teach our students is how to select and evaluate goals in the light of what they value and what they know. How do we accomplish that goal? Perhaps we can learn something from engineering design processes.

1.2 The Design Process

Design is at the heart of the process of engineering. It consists of a set of steps that lead a team of engineers from the initial postulation of an idea for a product through to its manufacture and marketing. If we break this process up into its major components, design, manufacture, marketing, we find that common to all is the notion of a feedback loop. We begin with some basic assumptions about what is possible, knowledge of some of the constraints under which we will be operating, and a specific goal, i.e., you are to build a *whatsit* that will do XYZ. Next comes the process of laying out ideas, examining the assumptions behind those ideas, proposing means by which those ideas can be put into action, and then returning to the original objective to see how far these deliberations have taken us toward that goal.

Walter Vincenti, a Stanford University aeronautical engineer, in his *What Engineers Know and How They Know It,* notes that design as a process

> . . .typically involves tentative layout (or layouts) of the arrangement and dimensions of the artifice, checking of the candidate device by mathematical analysis or experimental test to see if it does the required job, and modification when (as commonly happens at first) it does not. Such procedure usually requires several iterations before finally dimensioned plans can be released for production. (p. 7)

The key notion here is that the design process is iterative. Once you have laid out the artifact, determined its components, tested it and found that it fails, you go back and redo.

Later Vincenti breaks the process down into finer levels.

1. Project definition – translation of some usually ill-defined military or commercial requirement into a concrete technical problem for level 2;
2. Overall design – layout of arrangement and proportions of the airplane to meet project definition;
3. Major-component design – division of project into wing design, fuselage design, landing-gear design, electrical-system design, etc;
4. Subdivision of areas of component design from level 3 according to engineering discipline required (e.g., aerodynamic wing design, structural wing design, mechanical wing design);
5. Further division of categories in level 4 into highly specific problems (e.g., aerodynamic wing design into problems of platforms, airfoil section, and high-lift devices). (p. 9)

He goes on to note that

> Such successive division resolves the airplane problem into smaller manageable subproblems, each of which can be attacked in semi-isolation. *The complete design process then goes on iteratively, up and down, and horizontally through the hierarchy.* (p. 9, emphasis added)

In short, the design process is a process by which goals are given, the means to achieve those goals are broken down into smaller parts which are then put back together to see how it all shakes down and the process begins all over again.

1.3 Students as Self-Designers

We already know that students in particular are involved in self-design, often much to their parents' consternation. Fashion, music, dancing, language, are all appropriated by the young in ways they hope will allow them to exhibit their individuality. And if you point out to them the irony of the loss of individuality when they adopt the codes of a group, they have something of a legitimate response when they note that it *their* group. Nevertheless, the *art* of self-design is not foreign to the young. Our objective is to turn the *art* of self-design into the *science* of self-design, if you will. To do that we have to introduce the young to the notion of the consequences of their actions. Thus, "what do you think your body is going to look like in 30 years when that tattoo has faded and is just an ugly black smudge on your arm?" Put that way, of course, the question will only be met with overt hostility. So, the challenge is to find away to make students fix goals and to think of the consequences of their actions, and to correct for bad choices or unintended consequences.

The key, I believe, is in the distinction introduced at the beginning of this essay, active versus passive. We should seek to make our students active participants in their own lives. However there are problems with the distinction. First, it seems as if some people don't want to, or simply can't be active players. Further, I am not sure there is a single method for achieving the desired end. Third, the entire idea that everyone should be an active player needs to be examined. Instead of attending to these issues, important as they are, I will first try a different approach.

I have argued elsewhere that human beings are artifacts (see Chapter 13). That is, by the varied processes of mate-selection, be it deliberate, through falling in love, or by arranged marriages, we have made ourselves what we are today. With the possibility of gene therapy on the horizon, we can design our children, fix genetic "defects," etc. No matter that if we start selecting for "beauty" we may be doing our offspring a grave disservice as conceptions of beauty are notoriously fickle. Be that as it may, this is not exactly what I mean when I say, as in the title of this essay, that human beings are technological artifacts. I mean that we are technological artifacts by way of the enhancements *we* select for *ourselves*.

We have all experienced a power blackout. Nothing wakes us up to our almost total dependency on electricity as when it is not there. The interesting thing here is that, for the most part, we did not choose to live in an electrified world – we were born into it. Some, for either political or economic motives are leaving "the grid", choosing a solar power source or some other. But they are not choosing to give up the technologies that electricity powers, whatever its source. However, others are also beginning to ask questions like "Should I buy that SUV?" given the impact on the environment that increased gasoline consumption produces. But, they are not giving up the idea of a personal means of transportation. Nevertheless in choosing not to pollute the environment as much, they are making some sort of statement. And in making that statement they are saying something about themselves, just as they are making a statement about themselves by insisting on having a personal means of transportation. Some people I know do not own a television set. They claim they can get all the information they need from newspapers and they prefer to read. Some people do not own a cell phone, not because they can't afford one, but because they prefer the freedom of being able to walk around and not be summoned by the sound of a phone going off in their pocket, but they have a phone at home and at the office and they answer them when they ring. Part of what we in the United States refer to as a "Soccer Mom" entails owning a mini van or similar vehicle to haul the kids around to soccer games or piano lessons. Everyone has to have a home entertainment center, fancy or minimal, a house without music is uncivilized; but note, it is increasingly rare that the music is self-generated. We no longer play the piano at home for the purpose of making music – instead we put on the stereo. In short, you are the sum of your choices of technological enhancements, be they books, TVs, stereos, SUVs, or cell phones. The pieces of technology you have selected to be part of your life are as much a part of you as the color of your hair.

Several factors are at play here. First, the idea that you are the sum of your choices is not original – it is certainly a major theme in the work of Jean Paul Sartre. What may be new is the idea that in choosing to employ certain artifacts and technologies in your life they make you what you are. Imagine a teenager without his or her cell phone – they would be a totally different person. Think of the changes you undergo when your car is in the shop and not available to you. Sometimes, I am told, the feeling one has when deprived of a favorite piece of technology is like drug or alcohol withdrawal. You are simply not yourself.

If the above claims are correct, it is but a short step to the conclusion that human beings are themselves technological artifacts. If we are the sum of our choices and if our choices entail making certain artifacts part of who we are, then given our account of design earlier, it follows that we make ourselves what we are in a far more concrete sense than Sartre may have thought. Recall that design is an iterative process proceeding from an idea to a finished product. Testing to see if the mini choices we make along the way are going to all fit together to make the thing we are planning work. Making adjustments and corrections if they do not is an integral part of the process. The feedback loop is essential to the process; it is what makes the product work at the end; it is what makes it possible for us to change.

1.4 Active Versus Passive, Redux

Let us, for the time being, put aside the unfortunate fact that there are passive people and that they are perfectly happy being so. Maybe they are not "perfectly happy", but they are at least not willing to try to change. Instead, let us concentrate on the active person, knowing that there are individuals out there who engage the world and try to make their way in it, overcoming obstacles, seeking their specific vision of the good life. Finally, let us understand that the world they engage is increasingly populated by a multitude of technological artifacts, the choices of which to adopt will impact who they are, what future choices they have, and who they will become. If those individuals are our students, how do we prepare them to navigate that world?

First, note that we cannot *teach* them how to deal with the technologies that will be available to them. For one thing they are much more adept at manipulating the new information technologies, in particular, than we are. Second, we cannot begin to imagine what new and transformative technologies the future has in store for them. Third, you cannot anticipate the situations they will encounter and therefore you can't teach them how to respond to those. We can show a student how to use a screwdriver, but we cannot predict the circumstances in which she will have to use one nor what type will be needed.

If we can teach our students anything, it is that they will be faced with choices, and that they do have the ability to make choices, and they must evaluate the results of their choices, how to do that, and, finally what to do with the results of that evaluation. How will we know if we have succeeded? Crosby, Stills, Nash and Young had the answer years ago:

> Teach your children well,
> Their father's hell did slowly go by.
> And feed them on your dreams,
> The one they picks, the one you'll know by
> ("Teach Your Children" on *De ja vu*, 1970)

All that said, is there anything more we can say and do? There is; actually there are many things to do. The first is to recognize that many of the technologies that are second skin to our students, are also transparent to them. They do not appreciate the extent to which they have incorporated these technologies into their very being. Thus, a major challenge is to find a way to bring that to their attention in a non-judgmental way. If they are to learn how to and what to choose when we are not around, we must give them those tools. Telling them something is bad for them isn't, by itself, enough. For one thing, they won't believe us. Giving them the means to find out for themselves that something is bad for them is better. Furthermore, letting them find out that something they thought was going to be good for them and wasn't, *and* also giving them the tools to correct their mistakes so that they learn from them is even better. What I am proposing, in effect, is that we teach students how to use something akin to the design process on themselves. This of course, follows from what I have been arguing for when I made the case for human beings as technological artifacts. If we are the sum of our choices and our chosen technological artifacts, then deciding what kind of person we want to become entails learning how to make choices and how to evaluate them and how to evaluate their consequences, and, most importantly, what to do in the light of those evaluations. Our goal should be to make this process a habit, in both Dewey and Peirce's sense.

What I have in mind is an intellectual process that begins by acknowledging that even our youngest students come to us with a given knowledge base, some values, some fears, some hopes, expectations, etc. Given that background set, students should be provided with situations in which they must made decisions, i.e., they must choose between options. Having made their choices and acted on them, there will be results. The important part comes next. If there is teaching to be done, this is where it takes place. The students must be taught how to evaluate their choices and what to do after that. It is crucial that the outcomes are described in the language of "expected" or "unexpected", *not* in the language of "good" and "bad", for we are not in the business of approving or disapproving of their choices. Rather, we are helping them develop a habit of evaluating the outcomes of their choices in terms directly applicable to their own values and goals. Furthermore, an unexpected result is not necessarily a bad one; while unexpected, it might, nevertheless, lead to alternatives not previously anticipated. If, however, the outcome is not what was expected and, furthermore, not desired, then the student has to be shown how to revisit her previous background assumptions and bases and try to discover what it what that she thought she knew, believed, valued, etc. that needs to changed, updated in the light of the outcome, or straight out rejected.

Essential to this process is for the student to learn that making mistakes is not such a terrible thing and that there are good things that can come from mistakes: we can learn from them in such a way as to possibly avoid making them in the future. It is in the iterative learning process that we set the stage for the next step: choosing among technologies. Learning how to learn from mistakes and successes is the key to the entire process. And key to that is developing the ability to evaluate your starting point, your assumptions and background knowledge. Further, we absolutely must learn how to avoid the fear of failing fallacy. If you are afraid to fail, you will not choose, but not choosing is itself a choice, the choice to be passive. And to be passive is to put yourself in a position where you cannot learn.

If we now take this iterative learning process and apply it to the problem at hand we can begin to see some positive results. First, it is important to remind us of the problem with which we are dealing: how do we prepare our students to make intelligent choices about the myriad technologies they are and will be faced with, insofar as the choices of those technologies help create the person they will become?

The wrong thing to do is to approach the problem by announcing that technology is bad or out of control or some such. They know otherwise. Rather, if you are convinced that a particular technology is undesirable, focus on the consequences of choosing, allowing, adopting that technology on them. The difficult part here is getting them to understand that there are short term and long term consequences. Understanding long term consequences is not easy, even for experienced planners. Putting the issue to them in terms of what kind of a person they will appear to be to the world if this or that consequence occurs will, however, get their attention.

The fact of the matter is that we are our technologies. We would not be designing them if they were not essential to who and what we are. Accepting that as a starting point is not to give up, it is to recognize a fact of our nature. Further, while there are potential negative consequences of any choice, it seems more productive to look at the positive outcomes. In particular, since every choice we make closes down some options and opens up others, we can start our students thinking about their choices of technological enhancements in terms of the options they open up for them. For example, a fancy new computer comes on the market. You can take your hard earned money and buy it, or rest content with your perfectly good old computer and spend the money on some courses that teach you how to extend your current capabilities. Seeing the opportunities that your choices afford you is the key. The active person can design him or herself to be the most flexible and creative person, if they have the habit of iterative reflexive thinking.

Chapter 2
Technology and the Objectivity of Values

2.1 Introduction

Our task concerns the problems of learning to live with technology. Since the characterization of technology is itself an issue of some debate, let me begin with some preliminaries. I treat technology as humanity at work in the world. That is, technology is not a thing in itself; it is the techniques and methods, including machines, tools, social systems. etc., we use to make our way in the world. Given this perspective, let me now, for the purpose of focusing our efforts here, rephrase the objective of learning to live with technology in the following way: we are concerned with the problems created by the methods we use to manipulate and investigate the world. As such, the philosophical problems of technology are problems associated with the reasoning we use to develop and employ these methods and techniques and to assess the consequences, expected and otherwise, of their use. These problems range over a variety of issues. I will be concerned only with the nature and role of values in assessing technologies. And even then, my worries are narrow and restricted to the problem of structuring the debate over how best to assess technologies.

Concentrating on how to assess technologies raises the question of the appropriate framework within which to evaluate technological developments critically. As it turns out, there is more than one framework. Each is used under different circumstances, but they are all needed if we are to proceed rationally. The problem of evaluation is one of the major philosophical issues technology elicits and I am deeply concerned by the unproductive use of ethical frameworks in this domain. I will be arguing that the use of ethical categories to discuss technological developments is for the most part misapplied. This is not to say that under all circumstances ethical evaluations of technology are misplaced. Rather. I seek to differentiate ethical claims from other forms of assessment and to find an appropriate schema for indicating where each approach best functions. I am also not claiming that questions

---

Appeared in *Technology and Life World*, 1989, edited by T. Curry and L. Embrey. Reprinted with kind permission of The Center for Advanced Study in Phenomenology.

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_2,

of values are not legitimate. But it is important to see that not all questions of value are ethical ones, and, further, that the rational assessment of technology requires an appeal to objective values. In what immediately follows I construct an argument based on a discussion among a group of biologists on the issue of whether or not to continue recombinant DNA research. My reconstruction does not intend to capture faithfully the full flavor of the argument; it attempts rather an abstract version of it for purposes of identifying a type of problem that is typical in debates about technology.

In a 1977 debate at the National Academy of Sciences (see Research with Recombinant DNA. National Academy of Sciences. Washington, 1977). Ethan Signer, Professor of Biology at MIT, argued that using the techniques of recombinant DNA to develop more food is so risky it ought to be stopped. His argument went as follows: we are currently not doing all we already can to help starving people (perhaps implying by that we are not using the enormous grain surplus in America to its fullest). So, if we are not doing what we already can, why try something dangerous? Clearly the argument is confused. Signer runs several complaints together. First, there is the worry over the failure to use our grain surplus to alleviate hunger. Then there is the second worry, that the use of recombinant DNA is dangerous. While the two worries are not necessarily related, they do share a common feature: both are based on certain value judgments: first, about the appropriateness of using our grain surpluses in certain ways, and, second, the risk involved in the use of recombinant DNA techniques.

Two points need to be emphasized. First, the claim that using the techniques of recombinant DNA is dangerous is a value judgment. When one hears it from a scientist one tends sometimes to forget that. But it is true nonetheless. Scientists make value judgments, even about the facticity of certain aspects of their own areas of research. This claim about the dangerous nature or recombinant DNA is not as outrageous as Robert Sinshimer's assertion 3 years earlier at the 1974 Assilimar conference. Sinshimer argued that we should abandon any form of recombinant activity because it was against nature. Nature, he alleged, has set a barrier beyond which we should not go. How he knew this was the barrier was not revealed, but that was the claim. Sinshimer has since modified his stance, but it is important to note that both claims, that recombinant DNA is dangerous and that it is against nature, are value judgments made by scientists, judgments that exemplify the traditional is/ought fallacy. I examine this further below.

But before outlining my second point let me add one further note here. It may be remarked that recombinant DNA research has proceeded at an incredibly fast rate and that, given what we know today, it is unfair to criticize biologists who were worried 25 years ago when virtually nothing known about the potential results of such experimentation. In a sense it is true. My concern here is not with the rationality of the specific claims made in a state of relative ignorance but rather with the type of argument they represent: a form of intellectual bullying in which the authority of science is weaved into a value-laden position established independently of the facts of the case under discussion.

Returning to the Signer argument, the second point I want to emphasize is the inappropriateness of the ethical judgment about using grain reserves as a criticism of recombinant DNA technology. The fact there are other ways to accomplish a goal does not in itself mitigate against one particular option. I will also return to this problem later. First, let us look further at values and value judgments.

2.2 Types of Judgments

Many types of value judgments have been made about the kind of threat posed by recombinant DNA technology. These come primarily in the form of predictions about potential undesirable consequences if the results of recombination experiments do not pan out as we anticipate. These judgment/predications prey on our inability to know the full extent of the risk associated with results of this technology when it escapes human control. (See Sinshimer's comments in the 1977 Academy Forum, pp. 74–79.) The problem these judgments present is that they often obscure or confuse the facts of the case, thereby blocking our efforts to assess rationally the situation. This is not to say that value judgments play no role in assessing our technologies. The question before us concerns the kind of role they play.

Part of the problem here concerns our understanding of the nature of values and how they function in evaluative situations. New technologies, and some old, are often seen as presenting a threat to the life style (read "values") of some group or other. In learning how to deal with new technologies we need to learn how to assess the strength of such threats. But, more importantly, ultimately we are faced with the more difficult problem of determining the significance of the threat itself. Let me explain.

Objections to the introduction of new technologies are usually based on the claim that some valued state of affairs will be endangered. The snail darter was allegedly threatened by the construction of a dam. This dam was opposed because, largely ignoring the evolutionary history of species, that is, the fact that not all species survive over time even without man's intervention, some individuals made the value judgment that endangering the existence of a species is wrong. The fight that ensued was not over the efficacy of the dam, but over conflicting value systems with regard to importance of a near extinct species versus the needs of the human community. In other words, the battle was over values. And, unlike battles over facts, there is no way to determine which of these values are most appropriate.

If we had some account of values which showed how they could be as objective as facts are supposed to be, we might be in a position to short circuit many arguments over technologies. That is, if we could show that the preservation of a species is an objectively good thing and not merely an opinion, we should be better off. But even that won't do. We need more. We need both a theory of objective values, that is, a theory that explains how values can be objective like facts are supposed to be and, in addition, a theory that ranks values, i.e., that shows which ones are more important than others. If we had such a scheme, then if there were questions about whether or not we should proceed with some technological development, we could resolve them by appeal to the facts of the matter. The other option is to remove as much of the value dimension as possible from the discussion.

Unfortunately a theory of objective values is utopian. For one thing, it isn't obvious that values are there-in-the-world like events. If they were, surely by now we would have found them or have discovered some clue as to how to go about finding them. Second, there is also some question as to whether or not facts are there in the world. On some accounts, what count as facts arc a function of the theory being used to investigate the world. Hence, facts are descriptions of events or states of affairs in the language of one theory or another. And, on an extreme account of this view, as theories change, so do the facts (see Kuhn 1962). If facts change, then any search for objective values based on an analogy with the objectivity of facts breaks down. While I do not endorse the extreme version of this view, on which all facts change as our theories change, I do think that some sense can be made of the idea that some facts, those that are more theoretically based, change as theories do. (I will not develop this any further because it takes us afield into the murky lands of the realist/anti-realist debate that are not germane to our concerns here.)

Our other option for resolving conflict is to remove as much of the value dimension from the discussion as possible. Unlike our first option, which was utopian, this is just wrong-headed. To make progress we shouldn't ignore values – for that can't be done – rather, we must distinguish between those values about which rational discussion is unlikely and those about which we can proceed in a reasonably confident manner. The distinction I have in mind derives from Richard Rudner's observation that there is an entire range of values that we tend to overlook, but which still count as values (Rudner 1953). These are *epistemic values,* the values that guide our search for knowledge. They include truth, objectivity, measurement, justification, etc. And they are to be distinguished from those values operative in the other area, the search for the good life, i.e., aesthetic values. I call these "aesthetic" values rather than ethical or moral values because I see ethical values as ultimately deriving from broader aesthetic considerations. It should also be noted in passing that it is not being claimed that these two categories, epistemic and aesthetic, exhaust the range of values. They are just the two under consideration here.

2.3 Epistemic Values as Objective Values

The difference between these two domains, the epistemic and the aesthetic, is striking. And while the values by which they are characterized function analogously to Kantian constitutive rules, in that they determine legitimacy of actions, assertions, etc. within those domains, we are often slow to sec that epistemic values are values nonetheless.

There are many explanations for this state of affairs. The fundamental one is that epistemic values are not as vulnerable to the open-Question challenge as aesthetic values. Under scrutiny they may turn out to be as weak as their aesthetic counter-parts, but the initial reaction is not so intense. Epistemic values have a close connection with some sort of objectivity, meaning by that, with the way the world is. Thus, when we speak of a "true" statement, there is a sense that such statements relate to the world in a way that makes it possible to show without doubt that those statements have the property which they proclaim. The statement "snow is white" is true because the world really is the way it says it is. And it is that kind of "showability" that enforces the view that epistemic values like truth are immune to the open-question challenge. For what sense would it be to say, "Well, yes, the snow really is white, just like the sentence which you uttered and said was true claimed, but is that sentence really true?" This being the case, I want to argue that epistemic values are as close as we can get to objective values. That is, claims involving conflicting *epistemic* values can be resolved by a method which has a definite stopping point, namely, the world.

When we turn to particular cases, the situation is obviously more complicated than what has been suggested. But the line of attack here is straightforward. Epistemic values arc epistemic values rather than aesthetic values precisely because of their relation to empirical knowledge. Remember that by "aesthetic" is meant not merely the values of art but those of morality and ethics. Empirical knowledge is about the world, and if epistemic values fail to yield such knowledge, they cease to function in the epistemic domain. One way to see this is to look at the history of the theory of knowledge. The history of epistemology is the history of our efforts to discover those values that will most accurately guide our search for knowledge. And in our efforts to achieve a better characterization of knowledge over time we have had to change values. Thus, starting in the seventeenth century we eventually gave up the search for knowledge as certainty because we discovered that, as a value, certainty was non productive. There are other cases of value rejections and replacement as well. For example, one of Galileo's major contributions to the development modern science was his insistence on the importance, i.e., value, of measurement as a way of distinguishing scientific truths from metaphysics. As he pointed out over and over again in *The Dialogue* and in the *Discourse*, appeal to metaphysical principles must be replaced by rigorous measurement, for this is the only way to be sure of truth. In effect, Galileo was challenging the metaphysical principles used by seventeenth century Aristotelians with the open-question argument. Consider the following from the *Dialogue* where Galileo is attacking Aristotle's proof of the perfection of the world that appeals to the fact that the world has three dimensions:

> To tell the truth, I do not feel Impelled by all these reasons to grant any more than this: that whatever has a beginning, middle, and end may and ought to be called perfect. I feel no compulsion to grant that the number three is a perfect number, nor that it has a faculty of conferring perfection upon its possessors. I do not even understand, let alone believe, that with respect to legs, for example, the number three is more perfect than four or two; neither do 1 conceive the number four to be any imperfection in the elements, nor that they would be more perfect if they were three. (Galileo 1632, 1967, p. ll)

Galileo sought a method of providing reasons not open to the kind of objection he is raising here against Aristotle. And the general search for constitutive values in epistemology has generally proceeded in a similar manner, seeking ways to avoid charges of arbitrariness and metaphysical superficiality. The result has been a set of values that have come to occupy a place of honor, having survived the trials of experience (cf., N. Goodman 1953, Chapter 2). And to the extent that calling something "objective" is honorific, epistemic values are objective.

The point to be secured here is that truth is not just a property of a limited set of statements that meet certain logical conditions. By trial and error we have come to the conclusion that truth *ought* to be a property of those statements if they are to count as knowledge. The same holds for other epistemic values. Thus, on normal accounts, a statement does not have to be justified unless we want it to count toward knowledge; likewise for the role of measurement, confirmation, evidence, etc. The reason we have been able to select these values is that we know how to determine when we have knowledge, i.e., when these values have done their job.

Part of the reason we fail to recognize these concepts as values is that in epistemology we tend to use them in the formation of "conditions of adequacy," thereby disguising their normative features with an overlay of empiricism. But the reason they *are* conditions of adequacy is just because they meet our developed intuitions about what can and ought to count as empirical knowledge. This, however, is not enough to differentiate epistemic from aesthetic values. Aesthetic values also function as conditions of adequacy for exactly the same reasons. They too meet our developed intuitions concerning what can and ought to count as contributing to the good life. Nevertheless, there is a difference. For unlike the situation with aesthetic values, epistemic values do not lead us into an infinite regress looking for the ultimate conditions of adequacy. For when all is said and done, epistemic values can be shown to lead to better or worse conditions of adequacy because what is supposed to count as knowledge can always be tested against the world. The test of knowledge is whether we can use it to manipulate the world. There is no counterpart test for aesthetic values that escapes the open-question challenge.

Despite the efforts of some moral philosophers to consider the empirical consequences of ethical theories, moral categories are, for the most part developed *apriori*. This is not necessarily bad. The articulation of moral problems, categories, and specific values is a response to typically abstract philosophical questions deriving from our normative search for an acceptable account of the Good Life. Those categories stipulate the context for moral and ethical discourse. That context is the world of the normative ideal, of what ought to be the case (cf. Rawls 1971). Thus, ethical theories, unlike the theories of science are not intended to be directly testable by daily experience. Furthermore, it is not clear what to do if, when tested, they are found to be flawed in some way. There seems to be two options: revise the theory or change the world. The problem is that there no way to decide between the two and no *prima facie* reason to prefer one to the other. This general set of problems is exacerbated when we invoke moral considerations in dealing with the problems technologies create.

2.4 Problems with Ethics

Despite the fact that our technologies are the fruit of our labors, there is a major problem with using normative categories to assess them. Approaches of that kind often fail to address a technology in terms of whether or not it will do the job for which it was introduced. Instead, the technology is confronted with a set of ethical standards, formulated under different constraints than were employed in its creation. Thus, rather than asking if the technology works as it is supposed to, we compare the effects of the technology to an ideal picture from some other conception of what ought to be the case. In other words, we confuse two dimensions of our analysis, criticizing what is the case by reference to an account of what ought to be the case and rejecting appeals to what ought to be the case by reference to the facts. The point here, however, is that *it is not always clear that what is claimed ought to be really is what ought to be.*

Take, for example, the worry voiced by Signer over using recombinant DNA techniques to help grow more food. Signer's original argument was directed against the use of recombinant DNA because (1) no matter what science does to help produce more food, the problem of keeping people from starving remains a political one, and (2) we can't assess the risk of DNA, hence we should avoid incurring any risk. Here I would like to update Signer's worry and use his complaints to exemplify a general sort of problem. In particular, I will develop a Signeresque *Scenario* – not to be attributed to him – but compatible with the sort of approach he takes. To differentiate our scenario from Singer's original worry I will refer to it as S*. S*'s first objection is that we ought not to use recombinant DNA techniques because we are not doing what we already could be doing, i.e., using the enormous surplus of American grain to alleviate the problem of starvation. Reconstructing the missing argument, let us assume the reasoning to be something like the following.

(1) In those cases where the US Government feels it ought to alleviate some problem or other, we first ought to use available resources.
(2) The US ought to help.
(3) The US has available resources in the form of unused grain.
(4) Therefore, The US ought to use its grain surplus to relieve.

There are several problems with this argument. The first is that it is not at all clear what decision procedure was used to pick the grain surplus as the appropriate available resource. Second, no thought has been give to the *consequences* of using our grain in Africa. This is notable since there was some consideration given to the consequences of using recombinant DNA techniques; namely, it was said to be dangerous. But there are also dangers associated with sending US grain to Africa, three in fact. (1) Our grain may contain a biological contaminant that can be extremely destructive of native growth, such as the parasite that caused Dutch Elm Disease in the US (2) There is the social and economic danger of creating a dependency on the US (3) There is the social and economic danger of undermining the autonomy and stability of the way of life of the recipients of the aid. All three of these are real possibilities and yet they were not raised by Signer. Rather, and instead, the specter of unloosing an alleged dangerous, man-made creation was invoked.

This argument against the use of this bio-technology is skewed for several reasons. First, it is incomplete. It fails to elicit the dangers associated with the choice of sending grain. Second, it asserts that DNA bio-technology is dangerous without considering the facts. Again, we recognize that in 1977 the facts were unclear. But I still am concerned with the type of argument S* exemplifies. It makes it *sound* like it is a fact that recombinant DNA technology is dangerous. Further, if we are to compare what we know about the probabilities associated with the occurrence of the other *three* dangers elicited so far, we will find out that it is considerably more likely that we will create either a dependence on the US or release a dangerous organism already in existence in the US than we will create some kind of a monster agent using the technology of recombinant DNA. The reasons for this conclusion are several. First, the safeguards required then, as well as now, for the development and release of an artificial agent are not just strict, but considerably more strict than those for guarding against the other two dangers. Second, given the time it takes to develop and test an appropriate grain using recombination, the probability that we will have witnessed at least one of the three alterative scenarios will increase substantially.

I propose that S*s solution to the problem of aiding the starving people of Africa, i.e., use our grain and avoid the dangers of recombinant DNA, typifies the kind of reaction one finds to new technological developments. Those, like Signer, who offer such argument as S* did not obviously spend a lot or time thinking through the available options, examining them in the light of our knowledge about the kinds of problems they present and assessing their comparative merits. S*s response mixed epistemic, aesthetic and even political-economic frameworks, resulting in a confusion of the facts of the situation and placing the discussion on the level of emotion rather than reason. In other words, we have here what is, by analogy with Gilbert Ryle's notion of a category mistake, a framework mistake. If the argument against DNA research is that it is dangerous, then this discussion ought to take place exclusively in the context of epistemic considerations and then *all* possible dangerous outcomes ought to be considered – including the dangers of using US wheat. If it is, rather, a question of moral obligation, then we ought to use an aesthetic framework – what should not be done is to start with one set of claims, i.e., epistemic, and then without warning shift to an aesthetic point. The unintended shift to the aesthetic framework is signaled by suggesting as a solution to the problem that we ought to share what we have. It is then reinforced in the claim that recombinant DNA technology is dangerous. It isn't just that sharing is an ethical notion. Rather, in suggesting that by sharing our grain we will be doing all we can, we short-circuit the discussion of the available options and on suspicious grounds. This confusion of frameworks would remain even if S* had divided things up by suggesting a two stage process, one for the short term and one for the longer term. Say, for example, that what we can do in the short term is send grain, acknowledging that there is a certain danger in doing so and attempting to guar against it, while at the same time attempting to develop a long term solution to the problem.

By having a biologist describe recombinant DNA as dangerous we have created the illusion of a scientist asserting a fact, when he is merely asserting an emotional opinion. This is complicated by the fact that S* has behaved in a fairly unscientific manner by omitting a consideration of the dangers associated with the option S* advanced. This provides the ammunition to defend my view that S*'s entire approach has been conducted in the context of an aesthetic framework, where advocacy does not necessarily require full and considered examination of all available options. Rather, only when honesty is the supreme value will you be forced to give equal weight to all options. In that case there is no basis for choosing among them since honesty applies to the behavior of the investigators and does not reflect the merit of an option. Furthermore, in this case, those advancing S*choose to work within an aesthetic framework rather than an epistemic one. This may have a negative effect because of the failure to deal honestly with the problem. Rather than proceed by adhering to epistemic values, the use of aesthetic values along with epistemic ones has confused and misdirected the argument.

2.5 The Methodology of Technological Decision-Making

On the other hand, all is not lost. By distinguishing between epistemic and aesthetic values we have the basis for a methodology of technological decision-making. For now we can distinguish between two tasks:

(a) discovering all there is know about a given problem, including its history and the options open for its solution, with their relative merits and demerits, and
(b) choosing among our options.

In the process of discovering our options, epistemic values help constitute the framework for investigation. In choosing among our options, aesthetic considerations have initial priority. What should not be allowed is to have aesthetic considerations intrude into what is primarily an epistemic context. Thus, referring to recombinant DNA technology as "dangerous" is inappropriate in the context of attempting to determine whether that technology can provide an option and what its costs will be. Everything we do is dangerous to some extent. What is important to know is not that an option is dangerous, but what its consequences will be – only when we know that can we determine the relative merits of one option over another. What constitutes negative consequences will depend on the aesthetic values of the decision makers. Some consequences are more clearly negative than others, while other situations may produce results for which a choice is very difficult due to comparative weight of the values being threatened.

Does this mean that all technology related decisions are ultimately aesthetic ones, that it is impossible to be objective in our assessments? Not at all. If in the determination of our options and their strengths and weaknesses we allow the conditions for knowledge to guide us and, if we follow those restraints as closely as possible, then the resulting situation will be one in which the nature of the conflict among our values, if there is such a conflict, should be clear. When we are clear about those conflicts we are in a much better situation to decide rationally which course of action to pursue. Let me explain.

As noted earlier, disputes over what to do about technological innovation usually revolve around the perception that there is some threat to a life style or a set of values. When the facts of the situation are as clearly before us as possible and our options are clearly expressed, where there are differences they will rarely be over the facts, and when they are those differences can be settled by getting the facts straight. The differences can only be over the consequences of those facts for a certain privileged state of affairs. The desirability of that state of affairs is a matter of value. Only when those values are in the open can a reasoned decision be made. At what level or in what framework is that decision to be made? I suggest a third framework in addition to the epistemic and the aesthetic, the rational.

There are *three* stages to good technological decision-making. The first is epistemic, the second is aesthetic and the third is rational. Objectivity enters at two of those three stages: the epistemic and the rational. We have already seen how objectivity enters at the epistemic level in the form of epistemic values. I am now suggesting that objectivity also has a role in deciding between two options that differ because of their aesthetic components. In other words, *we can decide between aesthetic value* judgments *on rational grounds.* It is important that these value judgments are about predicted states of affairs, for this claim is not initially about decisions about values, it is about the consequences of choosing an option.

When we speak of making rational decisions, we mean employing a principle of rationality to make judgments, The principle I favor is the Commonsense Principle of Rationality (CPR). CPR employs a model of simplicity which is a virtue perhaps only in these contexts (see Pitt 1989). It only requires that we learn from experience. Thus, when faced with deciding among options that differ by favoring one aesthetic value over another, and where there is no obviously superior option, CPR instructs you to consider what we have already learned, that is, consider other similar situations. Thus, we have learned in other cases that even with the best intentions in hand, we can disrupt an economy by undermining its cycle of production. Simply providing a starving nation with grain will not solve the problems that situation presents. We need to provide them with the means for feeding themselves. In the African case, we are not just faced with famine. The lack of rain has resulted in the consumption of the seed store. This may not seem too important, we can always send them seed. But there's the rub. American seed doesn't grow where it is sometimes needed. The destruction of the seed store means the termination of a 100 year old or more research project in applied selection by the African famers. For decades they have been saving the best seed from each crop and using it to plan the next year's crop, saving from it the best in turn, etc. While American seeds have been selected for optional results in our environment, it doesn't follow that they will flourish in Africa. Finally, from past experience we know that recombinant DNA technology can produce good results, as in the development of artificial insulin and microbes that consume effluents. And as we carefully introduce artificially altered biological agents into the open environment, we have been increasing our understanding of the mechanics of recombination and the effect of natural selection on these agents. Thus, based on CPR if we were limited to one option, we ought to choose to develop a grain using bio-technology that will meet African needs. But CPR also tells us that for a problem of this magnitude we have often needed more than one approach. So the result would be a decision to invoke a variety of solutions, some short term: feed them our grain now; others long term: send them our seed and see if it will work and also get to work on a research program aiming at an artificially altered seed that will survive in the African climate. Now, granted, the solution proposed here has elements not available to Signer in 1977. But again, it is the kind of approach I am suggesting that has merit. Signer simply ruled out the recombinant DNA option because of unknown factors.

I have claimed that good technological decision-making has at least three separate levels and that objectivity enters at two of them. At the first level, the epistemic level, objectivity is a function of the values we have developed that are constitutive of our concept of knowledge. Once we have uncovered all the options, we should make our judgments with regard to which ones fit our aesthetic considerations best. But those value judgments themselves need to be assessed and this is where CPR invokes our experience. We need to learn from our past decisions based on value judgments. Not to learn from our mistakes is to be irrational. It is in this manner that we hold out hope for the evolution of our aesthetic values away from a priori analyses toward the rigors of survival in an increasingly complex world. And this potential for rational assessment of aesthetic judgment is the new and compelling aspect of our technological age.

Acknowledgment My thanks to Richard Burian for his helpful suggestions on earlier drafts and a special note of appreciation to Ronald Druzina, despite my unwillingness to concede.

References

Galileo, G. 1632, 1967. *Dialogue on the Two Chief World Systems*. Trans. S. Drake. Berkeley, CA: University of California Press.

Goodman, N. 1953. *Fact, Fiction and Forecast*. Cambridge, MA: Harvard University Press.

Kuhn, T. 1962. *The Structure of Scientific Revolutions*. Chicago, IL: University of Chicago Press.

National Academy of Sciences. 1977. *Research with Recombinant DNA*. Washington, DC: National Academy of Sciences.

Pitt, J. 1989. "Simplicity and the Aesthetics of Explanation". In Rescher N., ed., *Aesthetic Factors in Natural Science*, Series in Philosophy of Science, Vol. 12, pp. 27–34. Pittsburgh, PA: Center for Philosophy of Science.

Rawls, J. 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.

Rudner, R. 1953. "The Scientist Qua Scientist Makes Value Judgments". *Philosophy of Science*, 20, 1–6.

Chapter 3
Anticipating the Unknown

The Ethics of Nanotechnology

3.1 Introduction

The prospects of a fully exploited knowledge of how to manipulate the nano world can be frightening. The possibilities imagined by both advocates and opponents of research and development in the nano world range from the mundane like pants that can't be stained, to the horrific in military applications. The problem is what should we do now to ensure the benefits of nano technologies and to avoid the horrors. This situation is in principle no different from the one we faced with the advent of techniques for splicing and recombining DNA in the 1970s. How do we go about shaping the future when the genie is out of the bottle? Specifically, how should we think about the ethical and social consequences of developing various nano-technologies?

In this paper I propose that the only form of ethical reasoning capable of giving us some purchase on the set of issues future developments in nano technology present comes from Pragmatism. After offering some suggestions as to why more traditional philosophical theories of ethics are inadequate, I sketch a pragmatic theory of ethics and show how it helps us deal with these issues. Fundamentally, pragmatic ethics forces us to a kind of wholism that is both naïve and yet, nevertheless illuminates the nature and scope of ethical situations. It is, in William James' sense, forward looking, rather than a search for first principles. And, following a lead from Charles Saunders Peirce, the founder of Pragmatism, it asks us to think about developing a conception of The Good Life that, in the long run, can receive universal endorsement.[1] In short, the proposal developed here asserts that concentrating on

Expanded version of "Anticipating the Unknown, Ethics and Nanotechnology" appeared in *Ethics and Nanotechnology*, edited by B. Fabrace, 2008, Springer. Reprinted with permission of Springer Academic Publishers.

[1] The concept of The Good Life, capitalized, is a regulative ideal – as such it is in constant motion – an evolving ideal that develops as we learn more about who we are, what we can do and where we are going as a world. As I will argue below, there is no one conception of The Good Life currently in hand, and there may never be one, but trying to articulate what that might be seems to me to be a exemplary goal.

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_3,

the particular merits or demerits of a particular action in some limited context or other is simply the wrong way to think about ethical issues. For if ethics is concerned with The Good Life, as it is, it should be The Good Life For All. And, as we shall see, taking the "for all" seriously has far reaching consequences. Others are also thinking about a pragmatic theory of ethics. Jozef Keulartz, Michiel Korthals, Maartje Schermer and Tsjalling Swierstra in *Pragmatist Ethics for a Technological Culture* offer a proposal that augments some of what is presented here. In conclusion I will consider some of their recommendations.

My point of entry is fairly straightforward: in considering the consequences of various nano scenarios, those scenarios have to be played out as fully as possible, thereby giving us a sense of the full scope of the impact of these technologies. Not only are there short term and long-term consequences, but there are also local, regional and global consequences. It is only when we understand the full range of the impact of our actions that we can rationally decide among our options. Anything else is merely the exercise of prejudice. But we cannot consider the full range of our actions merely in terms of physical effects. What these changes signify is to be understood against our values and our goals. That is, they press against our concept of The Good Life writ large. I will elaborate on the notion of The Good Life below when I talk about the range of ethical theories open to us.

There are then three components to the view I wish to sketch here, a problem about ethics, a theory of rationality, and the concept of The Good Life. But first some thoughts about pragmatism and ethics.

3.2 Pragmatism

Unsurprisingly, pragmatism is different things to different people. But both Peirce, the founder, and James, his close friend for a while, and his protogé, whatever their ultimate differences, agreed on Peirce's original principle:

> Suffice it to say once more that pragmatism is, in itself, no doctrine of metaphysics, no attempt to determine any truth of things. It is merely a method of ascertaining the meanings of hard words and of abstract concepts. . .All pragmatists will further agree that their method of ascertaining the meanings of words and concepts is no other than that experimental method by which all the successful sciences (in which number nobody in his senses would include metaphysics) have reached the degrees of certainty that are severally proper to them today; this experimental method being itself nothing but a particular application of an older logical rule, "By their fruits ye shall know them." (Peirce 1906, 1955, p. 270)

Pragmatism was originally conceived as a method for determining the meaning of words and concepts. That method required that one consider what James would call "the practical cash value" (James 1907, p. 46) of the concept. I interpret this to mean some thing like, but not equivalent to, Wittgenstein's understanding that the meaning of a word its use. But it is more than its use; the meaning of a word is exhausted by the sum total of the possible inferences one can make using that concept. But that sum is in principle unknown since the meaning changes as the concept is employed in novel circumstances, so that meaning is evolving.

The second thing to note here is that Peirce especially saw this as a method to be applied to refining the meanings of scientific terms. In his 1907 book, Pragmatism, James explicitly rejected Peirce disavowal of metaphysics when he announced that "The pragmatic method is primarily a method of settling metaphysical disputes that otherwise might be interminable." (1907, p. 42)

There is a second respect in which James parts from Peirce and for which reason I would opt to stick with Peirce's account. Since Descartes most definitely, but even before, epistemology has been burdened with a solipsistic demon. Knowledge was and (in the school of justified-true-belief) remains a function of the individual knower. The major Peircean pragmatist move was to endow the community of investigators with the decision power to ascertain whether some claim or other could count as knowledge. The community set the criterion and revised it as various experiments proved earlier formulations inadequate. James returns the decision making power to the individual – his pragmatism is very subjective and individualistic and thus falls prey to the same problems that beset traditional epistemology. It is for this reason, that the community should be the source of ultimate criteria, that the kind of pragmatism I am advocating here is more in the spirit of Peirce than James.

While I am advocating a Peircean pragmatism, it should also be clear that I am extending Peirce's method to questions of ethics.[2] It is not clear he would disapprove, but neither do I clearly have his endorsement.

A third feature of pragmatism that is important for our purposes is the emphasis on the long term. This ties into a fourth point, so I will deal with the two together. The fourth consideration has to do with trying to see the big picture. The big picture is a sense of how we and the world will work together in the long run. What we do now has consequences and we need to try to see what they will be in the long term.

Pragmatism then is the attempt to clarify the meanings of key concepts and words over the long term in order to get a better sense of the whole so that we may act better today.

3.3 Ethics

I made a rather bold claim in my opening paragraph, i.e., that only a pragmatic theory of ethics can handle the potential problems posed by nanotechnology. This claim has already been challenged by a colleague, Tom Staley, who says, naturally, that it is false. So, let me put this in perspective. There are roughly four major kinds of ethical theories in the western tradition of philosophy: utilitarian, deontological, virtue, and pragmatic. I am going to argue that of these four, the first three fail, leaving us with pragmatism. But I will do more than offer a proof by elimination. I also offer a positive argument in defense of pragmatism.

It is tempting to cheat here. The move I have in mind, but cannot capitalize on, is this: only pragmatism is theoretically committed to the idea that we must learn

[2] I am now going to drop "Peircean Pragmatism" and simply speak of "Pragmatism".

and benefit from our mistakes. Under pragmatism it is possible both to learn from our mistakes while still being ethical, hence, since learning from our mistakes is the heart of being rational from a pragmatist point of view, only pragmatism offers genuine solutions to ethical problems. Q.E.D. To develop this argument would appear to many to win by default – for the theory of rationality I am appealing to is not the only one available to us. I happen to think it is the best one and that it captures in the most fundamental way what we seek in a theory of rationality. But I cannot argue for that here, it is too big a job – so, with apologies for the appearance of arrogance, I refer you to *Thinking About Technology* where I develop the view in some detail.

I will not attack utilitarian, deontological and virtues ethics theories separately. For our purposes here I want to suggest that one argument holds equally well against all three: they fail to guide our actions, with an emphasis on "guide" and "actions". Each approach fails for different reasons, but the result is the same. Some reasons why they fail are: for utilitarians, it often happens that we cannot calculate the consequences accurately enough to determine what produces the greatest amount of happiness in either the short or long term, or, for that matter, what constitutes "happiness". For deontologists, G.E. Moore's open question looms as large today as when it was first posed. We may have arrived at a definition of "good" or "right", but is it really good or right? I frankly don't understand virtue ethics, maybe because I don't know what virtue is, but here is my best shot. With help from another colleague, William FitzPatrick, virtue ethics shifts the focus from principles of to the character of individuals. So we ought to seek to do those things that improve our character. Here, the issue is similar to the one facing deontologists. Unless there is something else against which to measure whether this or that virtue actually improves character and such an improved character has independent merit, who can say?

But when all the philosophical arguments are concluded, the fact of the matter is that when people act, they do not consult deep and profound philosophical theories first – they do what they do and only when challenged do they seek justifications, and even in those cases it is rarely a philosophical principle they pull out of their bag of excuses.

The Original Ethical Question was "what is the nature of The Good Life?". Somehow, after 2,500 years of philosophical perversion, that question has come to be wrongly understood as "Should I be a utilitarian or a deontologist, or a virtuist?" But there is more to living The Good Life than adhering to (or, rather, professing) a set of rules that have been shown to be bankrupt for hundreds of years. And one reason the now classical ethical systems have failed to provide the guidance we seek is because their defenders have forgotten that we are concerned with living The Good Life, which doesn't always translate into performing isolated moral actions. Not all actions are moral actions, and not all actions designed to contribute to The Good Life are moral actions. And those who would have us think so would have us replace an awareness of the genuine complexity of living with superficial formulaic solutions. In short, whether or not to go to a nude beach is not a question that can be decided by appeal to the principle of doing what is best for the greatest number, but it is a question about the kind of life you want to lead. And to be frank, whether or not you want an all over body tan has nothing to do with anyone else's well being and don't let moral Fascists tell you otherwise. I may be personally disgusted by seeing you in the all together, but there is nothing moral about my disgust, just the application of a perfectly good set of aesthetic principles!

3.4 Philosophy and Pragmatism

So, if we are interested in living The Good Life, instead of fleeing from it, then we need to locate the domain of philosophical inquiry most appropriate to solving that problem. And there's the rub, for it is not as easy as one might think to locate the appropriate domain wherein lie the right principles to help you decide about and then live The Good Life. Today we tend to characterize philosophy as being composed of five areas, epistemology, metaphysics, ethics, logic and the history of philosophy. Unfortunately this way of capturing philosophy doesn't cover the whole area of enquiry we call Philosophy. If we see Philosophy in the traditional way, we end up with some notable leftovers, such as philosophy of science, political philosophy, aesthetics and more.[3] Now those fascinated with neatness will try to stick the leftovers in one of the established areas. Thus, philosophy of science gets shoved under logic or epistemology, depending on your preferences. Political philosophy gets lumped with ethics, or under it, another mistake; and we are always left holding the questions of aesthetics in our hand looking for some sort of philosophical rug to sweep it under, hoping no one will catch us.

The mistake, of course, is to try to package neatly something as messy and complex as philosophy. There is no obvious reason why philosophy should come readily dividable into those five packages. To see this, we need to understand what the goal of philosophy is. Those who adhere to the five box theory seemed to have forgotten what we are trying to achieve as philosophers. We know about the goal of physics, to understand the components and structure of the physical universe. Biology is concerned with the nature of life and living things. History with the nature of the past. But what about Philosophy? What is its goal?

According to Wilfrid Sellars,

> The aim of philosophy, abstractly formulated, is to understand how things in the broadest possible sense of the term hang together in the broadest possible sense of the term. (Sellars 1963, p. 1)

If Philosophy's goal is to see how it all hangs together, why should we assume that it all comes together under five simple headings or three or ten? We shouldn't because it doesn't. The complexity of the world demands a philosophical method that accommodates that complexity. That doesn't mean we should abandon the five categories or three or ten, they have their roles. It just means, among other things, that if we insist on them only we will probably end with philosophy being irrelevant to the world in which we live since as new topics come up, like nanotechnology, the

[3] See Chapter 5.

opportunity to address them is lost because there is no canonical box into which to put them. On the other hand, if philosophy is to proceed as a dialogue through which we attempt to see how it all hangs together, then we cannot dismiss the approaches of the past out of hand, for the present is a function of the past and to understand the present we need to deal with the past. It is for this reason that many of those with whom we have this dialogue known as philosophy are thinkers from the past. To do philosophy requires understanding them, understanding how they expressed their ideas and how to interpret what they said. Many of the words we use as philosophers obtained their meanings from earlier philosophical systems. So we must begin with the legacy that has been bequeathed us, the world as it is, the language as it is spoken, our history as it is written, and then proceed to analyze, correct, rearrange and put it together in a more coherent form as we learn more and more. The problem being, of course, that as we learn more and as relationships become more complex, the more we will need to keep analyzing, rearranging and resorting, continually trying to make it all make sense. But surely that is preferable to trying to live according to rules put forth in a story written centuries, nay millennia ago, in a world vastly different from our present one with players and social arrangements now foreign to our understanding.

What I am after is this: the philosophical job is on-going, it never ends, because the complexity of the world is as much a function of what human beings do as anything else. And since we have this seemingly infinite capacity to mess up whatever appears to work, there remains the constant need to keep trying to rearrange things so that they fit together – where making things work in the world of ideas is to see how it all hangs together. This also means that no matter how pretty a picture we manage to create at any given time, we must accept the fact that it will shortly be out of focus because the world will not stop for temporary philosophical perfection.

The conclusion that follows from this discussion is that there is no absolutely preferable way of seeing how things hang together, since the bits and pieces we have arranged so nicely to justify doing what we want to do today will soon be augmented by new pieces and players, or impoverished by the loss of old ones. The world will not stand still for us – we live in a sea of constant flux without any permanent dry land on which to stand. And no matter where you go, everyone else will be trying to find some place firm enough on which to stand, frantically jockeying for position only to find the sand underfoot being undermined by yet another ocean current. So we live in a world without absolutes, without constants, one that is constantly in change.

Well, things are even worse than that. Like my hero, David Hume, I feel it is important that we allow our philosophical ruminations, to be carried to their logical end, to reduce us to melancholy, if you will, for only then will we allow a different control to take over.[4] So for now let reason continue as our guide to complete and total dismay.

[4]See Hume (1739, 1888, pp. 263–274).

3.5 The Law of Unintended Consequences (LUC)

Let us assume that history is wrong and that our favorite ethical theories are not defeated, meaning by that that they have not been shown to be internally inconsistent, or based on false premises, or irrelevant. There remains, nevertheless, one fatal consideration that forces us to find them less than satisfactory as guides to living The Good Life. That is the one law of human existence that is incontrovertible, the Law of Unintended Consequences. Our best efforts notwithstanding, the best laid plans of men and mice will come a cropper because of our in-principle inability to predict the future with complete accuracy. What do I mean?

The fact of the matter is this: when we act we set into motion events over which we cease to have complete control. When you consider how many actors there are in the world, and how many actions we each initiate, then anticipating the results of compounding their effects on each other simply cannot be done. It is not just that there are too many variables. With a big enough computer and fine-tuned programs, we could probably make accurate predictions – on one assumption. That assumption is that the world is deterministic. However, even if the physical world is governed by laws of nature, the social world is not governed by deterministic laws, and if you factor in the inconstancies of nature, such as tsunamis and the erratic actions of 6 billion people, is it any wonder that we cannot predict the future?

The Law of Unintended Consequences says this:

> Calculate as accurately as you will, the future will be different than you predict because you cannot factor in all the possible consequences of even one new action. No matter how hard we try to make things work out, something always seems to happen which spoils the result.[5]

Consider the following somewhat strained example, based on fictional events at my home university. You reserved rooms in Blacksburg's best hotel for graduation the day you came to Virginia Tech as a freshman. The week before graduation there is a horrible thunderstorm which blows the reservations program in the hotel computer. Furthermore, the hotel has recently been bought by a private investor and there is no national backup. Further, the original owner destroyed all records after it sold the motel because it needed the storage space for current matters. Net result: no hotel reservations for your parents, grandparents, or your Swedish great-grandmother who has flown in just for this occasion and who at age 102, after hearing of this disaster has a heart attack and expires. Unlikely? No. You did all the right things, you made early reservations; each year you called to make sure you still had them. You alerted your family to the need to make early plane reservations for great grandmother, you studied hard so you would graduate on time, etc. You did everything right and still lost. And, in the immortal words of Kurt Vonnegut, "so it goes".

---

[5] A number of year ago I was sent by a well-intentioned Dean to a conference on Chaos Theory hosted by the US Navel Academy at Annapolis. The ideas I was introduced to there concerned mainly predictions about the physical world. But as I read what I have written here I now realize how deeply I was impressed by those ideas and they apply to the human world multi-fold.

So, not only are your favorite ethical theories worthless, because they are flawed and unworkable, because no matter what you do, as this example is intended to show, there is a certain sense in which you are doomed to fail.

If I may be so bold, at this point I would like to make a suggestion. Since we are, in so many words, doomed to not doing the right things, no matter how we try, then maybe we shouldn't be concerned about doing the right things, where "doing the right thing" means meeting the criteria of some abstract ethical theory. Instead, why not be concerned about the way in which you go about trying to do the right thing. I am talking about form, method, or, as I prefer, style.

3.6 Way of Going and the Good Life

In the world of horse showing, in particular, hunter-jumpers, there is a phrase I would like to appropriate: Way of Going. It refers to that component in the competition in which points are awarded for the way the horse and rider together make their way over the obstacles. For their way of going, their sense of each other, of their being a team working together, listening and responding to one another, they receive points, as well as for not knocking down any jumps.

In a world without absolutes, in a world in which we are constantly changing and responding to new demands made by nature, other people, and institutions, what better model for us than to evaluate ourselves in terms of our way of going? Let us take a look at what is involved here. But so you are not misled, I am not going to start with nice and neat definitions and deduce all the relevant consequences. That is exactly the approach I reject. What I will be doing now is trying to give a feel for the kind of view I am advocating.

First, we need to reevaluate the conclusion of the graduation story. You did nothing wrong. Several events transpired over which you had no control. In a world where that is understood, Grandmother might not have died just then. Instead everyone would have crammed into your one room apartment and made do, being left with a marvelous experience to reflect back on over the years, and you would walk away with wonderful stories to tell your grandchildren about trying to get 127 people showered and dressed and out to graduation in the football stadium by 8:30 am, where then it rained.

Why would this Pollyanna ending result if I am right? The key concept in "Way of Going" is team-work. The horse and rider get more points the more they seem to act as one in accomplishing their task. Living The Good Life means working with your fellow human beings to create the best possible world. How you go about accomplishing this is not easy, nor is it obvious what these loaded concepts mean. Their meaning is to be worked out in the times and places in which you will find yourselves. Furthermore, we need to recognize that sometimes things work and sometimes they don't. Our job is not to decry the failure of ethical principles, but to find something that works at the time it is needed. Does that lead us down the dreaded path of relativism? Not necessarily so.

3.7 Common Sense Pragmatism

This is where pragmatism makes its entrance. The fundamental principle of the form of pragmatism I am advocating is the principle of rationality introduced earlier: learn from experience or Common Sense Principle of Rationality, CPR. It starts by recognizing that when faced with selecting among options, we do not make our choices in a vacuum. We come to any situation equipped with a set of values, goals, and background knowledge. We rely on some subset, or maybe more, of what we bring to decision-making in general. But when we select an option, and act on it, that is not the end of the story. As I have been emphasizing, actions have consequences. Under CPR we are required to follow through and consider the consequences of what we did and then to reevaluate the assumptions, goals, values, and background knowledge we used to make the decision that led to the action that had those consequences. This is exactly what Peirce meant when he spoke of clarifying our concepts over time. As we move forward we need to constantly reevaluate. Once we have done that, we must make whatever appropriate adjustments we can to those assumptions, goals, values, and background knowledge to improve our ability to select the best option the next time. This is not to say that the option we selected had the expected consequences, but only that whatever the consequences, they ought to affect the set of factors we used to make that decision. If things turned out the way we hoped, then we can say that those things we employed in making that selection seem to be solid, for the time being. If things did not turn out the way we expected, then we need to reconsider the importance of that goal, or the place of that value in our preference ranking, etc. In so doing, we are trying to make coherent those various factors that come into play when we make decisions and act. In other words, we are trying to make it all hang together in a way that helps us to do the right thing, which is to fashion and live The Good Life.

3.8 Common Sense Pragmatism as an Ethical Theory

Common Sense Pragmatism, based on CPR, is a method for making choices that lead to actions that have consequences for our conception of The Good Life. Why does this make Common Sense Pragmatism the only ethical theory that we can use in the face of the uncertainties raised by the possibilities of nano technology? As Anna Russell might note here, "Remember nano technology?" And how did a method, Common Sense Pragmatism, become an ethical theory?

Well, we also have a bit more than Common Sense Pragmatism. In our discussion so far we have noted that the world is far more complicated and difficult place than even the most diehard cynic would have it.[6] Not only can we not act assuming that everyone else can be held still. We are all acting and our collective actions have impacts on each of us and on the social and physical environments in which we live,

---

[6] Well, perhaps not Nicholas Rescher. See his 1998 *Complexity*.

and all of this ramifies and ramifies. It is truly amazing that we can make any plans at all that come out half way close to what we had hoped.

We also have the role of the concept of The Good Life to consider. It is clear that there is no one view of The Good Life that has priority or that even claims the allegiance of a majority of the world's population. Under those circumstances, how can we expect the concept of The Good Life to play any kind of a significant role in ethical deliberations? I propose we take a page out of Peirce's book. According to Peirce, scientific enquiry occurs in the context of a community of investigators. What constitutes knowledge is a function of the norms of the community as they evolve over time. Further,

> Different minds may set out with the most antagonistic views, but the progress of investigation carries them by a force outside of themselves to one and the same conclusion. No modification of the point of view taken, no selection of other facts for study, no natural bent of mind even, can enable a man to escape the predestinate opinion. This great hope is embodied in the conception of truth and reality. The opinion which is fated to be ultimately agreed to by all who investigate, is what we mean by truth, and the object represented in this opinion is the real. (Peirce 1878, 1955, p. 38)

I would not propose that there is by necessity one perfect arrangement by which humankind should live. But I am convinced that in seeking for that one perfect arrangement, and measuring our actions in terms of how they contribute to or detract from contributing to that effort is a very good way to determine how to act. For a pragmatist, the meaning of our actions is to be found in their consequences – consider the consequences. When deciding on what action to take, we should consider the consequences, both for ourselves, and for others who will be affected by those actions. But I am not talking about just the few individuals I can imagine being affected by my actions. When I say that we should have The Good Life in mind, I am suggesting that you ask yourself this question: will this action here and now contribute in the long run to my current view of The Good Life? Further, if you are, and I suggest you must be, rational, then it would follow that you should reevaluate your conception of The Good Life in the light of the effects of your actions, and that this should be a constant activity. In this way, your and my conceptions of The Good Life will be a constantly evolving empirically informed ideal.

There are some goods we can agree on from the start: a room over our heads, foods in our bellies, health care, safety, etc. What we see at work in the world are different efforts to achieve these goals – and it may be the case that for different peoples, different arrangements might make sense – but part of the living The Good Life is acknowledging that and factoring it into your conception of The Good Life.

So that is the general idea – be rational according to CPR and seek The Good Life, knowing that it is always a changing ideal but one on which humankind can, with good intentions, most probably, in the long run, come to agreement. Thus if you are committed to The Good Life in the long run, the ethical counterpart to Peirce's commitment to convergence on the truth, is the normative injunction: seek to bring about The Good Life.

3.9 Common Sense Pragmatism, Ethics, and Nanotechnology

Lots of people are engaged in the design and manufacture of nano technologies. It seems to be everywhere. In our local paper, The Roanoke Times, In a March 2004 Sunday special report, there was an article on a start-up firm in Wytheville, Virginia, making bucky-ball shaped nano-structures based on a discovery by a Virginia Tech researcher. This device is intended to be a vehicle for carrying medicine into the body. Is this something we should encourage or actively work to stop?

The reason I believe pragmatism provides the only viable framework for dealing with nano technology is that the perceived ethical problems associated with nano technology today derive primarily from a fear of the unknown. Actually, it is fear not so much of the unknown as a fear based on speculation of the "what-if" sort. Thus, people worry about nano particles that could be sprayed like a gas in military contexts that would eat a person's insides – particles so small no filter could stop them from being inhaled. To be frank, no ethical theory is going to deal with that. When it comes to military technologies, I am afraid that ethics is not the issue – beating the enemy no matter what is.

No, we need an ethical framework for dealing with issues that really are, to some minimal extent, under our control. So, when we worry about the bucky-balls noted above, the problem is determining the consequences of introducing these devices. These devices will be expelled from the human body – how will they affect the water supply, the soil, the air? Consider this parallel: antibiotics have decreasing efficacy because they are also being used to cure food animal diseases. We ingest beef treated with the same antibiotic that we also use to cure a human infection, but because we are being constantly exposed to the antibiotic through the beef we consume, it ceases to have the effect it should when used to combat infection in humans.

There are numerous scenarios out there in which nanotechnologies are seen as the key to eternal youth or immortality – nano thingies will destroy cancers, delay aging, fight infections, etc. Some thoughtful people will ask "Is curing cancer a good thing?" The pragmatist, seeking to clarify the meaning of the question, turns it into a different question. The question we face is not so much "is curing cancer a good thing?" as, rather, "What effect on The Good Life will a population without cancer have?" On the surface eliminating cancer is a good thing – in the short run – for me. But in the long run what is the effect? Likewise for wrinkle free pants – currently being manufactured. What if those pants turn out to be not only wrinkle free, but also non degradable? Landfills get filled. If the nano treatment the pants receive contributes to their remaining in good condition, good and useable for a much longer time, then what is the effect on the garment industry, and on the cotton growers, etc.? Cameras on cell phone seemed like a not so bad idea, even a cool gimmick – but today many gyms won't allow cell-phones in locker rooms because some men are taking pictures of their friends without any clothes on and sending them to their girlfriends or worse.

The problem with nano is that we don't know yet what it can do – hence we need a way of thinking about it that not only has us considering the consequences of this type of device or that – but considering the consequences in terms of their impact on our way of living – i.e., The Good Life. It is not enough to think about the consequences, the consequences must be weighed in the light of something else and I propose that the future wellbeing of all of humanity is the place to begin, recognizing that that future is constantly changing as we learn about more about the present. But the question remains as to how to do that.

3.10 A Different Approach

In *Pragmatist Ethics for a Technological Culture*, Jozef Keulartz, Michiel Korthals, Maartje Schermer and Tsjalling Swierstra. provide an excellent discussion of the need for a new vocabulary for ethics vis-à-vis technology and offer up a proposal for solving the problem of how to consider the global well being of humanity. That said, I have some concerns. I raise them in the spirit of continuing the exploration of the pragmatic ethics they have begun and not merely to be contentious. My two concerns are: (1) what I perceive to be a move towards ethical colonialism, and (2) the authors' conception of pragmatism.

As I understand the authors' argument, it goes like this: (1) life in a technological world is dominated by change – therefore there can be no universal principles (I agree); (2) new technologies make for new possibilities for action (right again, especially since possibilities for action raise the specter of new moral problems); (3) the best way to decide what to do with new possibilities and new moral problems is to have as many people as possible in the discussion (idealistic and probably not really a good idea because not everyone has something contributory to say on every issue); (4) new possibilities and new solutions require a new way of speaking – therefore we need creativity, and the best way to get that is to have everyone at the table (not clear).

At a conference held at Virginia Tech on the meta-ethics of moral value in March 2003, Sara Williams Holtzman gave a fascinating talk about the ethics of mountain top removal mining in which she wanted to accord mountains moral status. In the discussion that followed, I proposed that she was engaged in ethical imperialism. After reading *Pragmatic Ethics for a Technological Culture*, alas, I feel compelled to accuse Keulartz, Korthals, Schermer and Swierstra, of being unknowing allies of Holtzman and equally committed to ethical Imperialism, or more to the point, ethical colonialism.[7]

What do I mean by ethical colonialism? It is the attempt to endow everything in the world as an actor with moral value. It is to deny that there are other types of values, such as epistemic values, which have their own integrity and can operate in an ethically neutral framework (see Chapter 3). Consider this claim from their opening

[7] Andrew Garnar suggested the shift from "imperialism" to "colonialism". For he is correct in that the idea I am after is subjugation of new lands and the imposition of a new set of values for indigenous ones.

essay: “The example of the pill makes it clear that technological artifacts possess a written-in or built-in normativity. They embody particular options and restrictions, and thus reinforce or alter existing role divisions and power relationships” (p. 9).

There is a certain seductive allure to this Latourian idea that material things have a kind of agency. By virtue of their existence in our field of action, we are forced to accommodate artifacts and that, it is mistakenly claimed, gives them agency. I, however, find it somewhat problematic to attribute agency to a tree simply because I have to walk around it.

The idea that technological artifacts possess normativity follows the same line of thought as Latourian material agency. That objects are to be used in a certain way, or come to be used in a certain way, however, does not mean they possess normativity. Nevertheless, the authors in their Introduction note that, “Technological artifacts carry a script or scenario within them; they require particular role patterns and role divisions and lay down a specific ‘geography of responsibilities’” (p. xvii). Now, why is that the case?

The example they use to cash this idea out is the birth control pill – they point to the fact that it has given women control over their bodies and over when they will bear children – altering the power relation between men and women. They also point out the fact that the pill has facilitated the separation of reproduction from sexual activity, and its use has altered our perception of family planning. All of that is to the good – no objections from me here – but they neglect to mention that the pill was finally developed by a Catholic researcher, John Rock, for the express purpose of regularizing menstrual cycles of women who were having difficulty becoming pregnant because of irregular cycles. The point was to have more children, not fewer. If the pill embodies normativity, it is the norms of a tradition in a male dominated society. Now surely that is not the desired outcome of the example.

Let me try an alternative account, one that is not in conflict with the desired end of a pragmatic ethic, but one that places normativity in people, not in things. I do not object to the authors’ claim that technological artifacts have normative significance. But the normative significance is a direct function of how people choose to view them and use them. It is the use to which artifacts are put that exhibits the normativity of the users, not the things. And this is a very pragmatic point of view. As I argued in *Thinking About Technology* (2000), in discussions of technologies, the emphasis should be on the decision-makers, not on the objects. Whatever normativity there is with respect to the designing, manufacturing and use of artifacts, is to be found in the values, both epistemic and non-epistemic, involved in making the decisions to do this rather than that, to use this material rather than that, and to use this tool rather than that. Decision-making is a value-dominated activity. With respect to technological artifacts, whenever they are employed, it is because a decision has been made to do so. Making decisions is an inherently value-laden activity since it always involves making a choice. Understanding the normative dimension of technological artifacts requires an analysis of the factors that played into the relevant decisions. If you want to understand the normative dimension of the birth control pill, ask a woman why she chooses to use it. Ask the manufacturer why those materials rather than others – or why this shape for the dispenser rather than another. If I may be so bold, there is, therefore not one normative dimension to technological artifacts but many – and that in part is why there is so much discussion over technologies that promise the most. The more a technology promises, the more choices have to be made.

The insight that our authors are in risk of losing is that new technological artifacts open up possibilities for human action – which is what I think they mean by the "geography of responsibilities." The responsibilities, however, are not in the objects, they are ours. For example, it is our collective responsibility to come up with a protocol or two regarding human cloning. It makes no sense to say the responsibility lies. . .where? in the process? But what is the process other than what people do? Somewhat in jest I propose to my students that now that human cloning is possible, not only will someone do it, but it also means that men are now facing extinction. Talk about changing the power relations! – if the only thing men are needed for is reproduction, then we are no longer needed. Given recent events, peace might have a better chance if men did not control the decision making process in our government. Women control the economy – they are smarter than we are and they live longer – we are done for.

To return to the issue at hand – the normativity of technological artifacts – it seems to me that by placing the normativity in the people making the decisions rather than in the technologies, we make the possibility of the authors' end-state, creative democracy, more attainable. Why do I think this?

As I see it, what makes change problematic is that change, especially technological change, threatens a person's perception of the good life and in so doing challenges his or her values. No one likes to have his or her values challenged. Our values are at the core of who we are – not all our values, but the basic ones, like, for example, protection for our children. If a proposed technological change, such as locating a nuclear power plant next to my house, is perceived by me as a potential threat to my children, then you certainly can expect me to object. And because there is little rational argument or rational deciding on our values, the possibility of rational discourse when there are clashes of values is very low unless something intervenes. Therefore, just bringing people to the table is not enough – it is not enough to get them to listen to one another and it is not enough to generate the kind of creativity needed for making decisions about the new possibilities and the nature of the new moral problems these possibilities bring. Several things are at issue here. The first concerns getting people to engage. The second concerns the meaning of "creativity".

Concerning getting people to listen to someone with a different point of view is a rock bottom problem; it permeates all peoples and societies. This is the same problem as understanding the possibilities a new technology offers – it is the problem of overcoming our fear of the unknown. Often people refuse to engage in discussion with someone who holds radically different views from theirs because they are afraid of the challenge the new ideas may pose to their own views, views for which they often know they have no defense. That is, they are afraid of what they don't know or what they fear they may lose. Likewise, it is not so much that people object to technological innovation because of what they know about it – it is rather what they don't know what makes it so difficult to have a reasoned discussion. So I agree, something like creativity is called for – but I think we can be a bit more precise. If I am right and the problem is fear of the unknown, then what we need to do is to eliminate the fear by making the unknown more familiar, or, rather, to make the technology appear more familiar, so that what it can do is not so threatening. I think the way to do this, to reduce the fear of the unknown, is to use metaphor. But, to see why this may work, we need to get a better handle on what we mean by “creativity”.

Elsewhere (Chapter 12) I offer the following account of creativity: To be creative is to produce variation given the constraints of the materials and other parameters within which you engage in the deliberate design and manufacture of the means to manipulate the environment to meet humanity’s changing needs and goals. The problem in seeking creative solutions to ethical problems posed by technological innovation is in not knowing where to start. My suggestion is to start with the way we talk about our technologies and to probe what we think are their possible ramifications. If we can come to a common language through which to discuss the problems, we may have a chance to actually find solutions. But, unlike the positivists, I am not proposing that we develop a formal language from the start. Rather, I think we should circle the problem, trying out different ways of talking about it until we find one that satisfies all parties. The way to begin this process is through metaphor.

Metaphor, by its very nature, gives meaning to the new by way of associating it with something already understood. Irrespective of what Al Gore had to do with it, calling the world-wide web the information super highway was very helpful to many people in coming to grips with the potential for this new technology. It also helped to open up some of the ethical issues, like privacy. Calling it an “information” highway raises the red flag that should come up whenever issues of information are discussed. Further, because it is not just a case of potential eavesdropping, the manner in which this phrase serves as a metaphor becomes clearer. By using metaphor, however, we are not doing what our authors do not want us to do, which is to live in the past. The fact that we can rely on what is already understood, does not entail that we stop there. Metaphor extends the use of language, it changes the meaning of words, words that had a familiar meaning now mean even more. The material world is not the only thing that changes constantly – so does language – just try understanding a 16 year old today. But because it looks like the old language, we are often not aware of the extent to which language changes – unless you are French and keep a careful watch over the purity of the mother tongue. But how do you say website in French?

The way to find the right metaphor or metaphors is by applying the pragmatic method. The pragmatist’s first maxim is “consider the consequences”. If we keep in mind the consequences of using this metaphor, we can work our way toward a metaphor that captures what concerns us in terms of the consequences of allowing this new technology. For ethics, this does not reduce to mere consequentialism. The consequences the pragmatist considers are not just the effects of his or her actions; it is the effects of those actions on his or her knowledge base and on his or her values and goals. Translated, this says that in the evaluation of choosing A over B or C over D, there is more to consider than merely the physical consequences, there are also consequences for your vision of the good life, for the rock bottom values that constitute what you hold most dear, what you hope all people value. If you lie, what does that do to your self-image? If you use a gun to kill deer, what does that say about your conception of a civilized human being? If you use 5,000 pound bombs and withering machine gun fire to kill enemy soldiers and civilians alike for no clear reason, what does that say about the character of a country's leaders? What a pragmatist does is to consider the consequences and then using a feed-back loop to return to his or her assumptions, values, knowledge, beliefs and readjust in the wake of what has transpired. Yes, it has a bit of relativism associated with it – but the second maxim of pragmatism helps to derail the slide to total relativism.

The second maxim is: the ultimate arbiter is the community. However you adjust your values, there are still the values of the community that override and with which the values of an individual must co-exist, and that must be considered as well. But what happens when the community fails, when the election process is subverted, when the leadership does not listen to its people? I don't know, but that is a problem for all, not exclusively for pragmatism.

So in closing, let me summarize:

- Technological artifacts do not contain values nor normativity.
- Democracy will not solve all our problems because you have to get people to listen to others – not just talk at them.
- Using metaphor to demystify the new may help in getting people to actually communicate so they can talk reasonably about new technologies and old.
- Nevertheless, the two basic pragmatic maxims can serve as a basis for an evolving ethical system.
- Consider the consequences.
- The community is the ultimate arbiter.

References

Hume, D. 1739, 1888. In Selby-Bigge L.A. ed., *Treatise on Human Nature*. Oxford: Oxford University Press.

James, W. 1907, 1955. Lectures on Pragmatism. Cleveland, OH: The World Publishing Company.

Keulartz, J., Korthals, M., Schermer, M. and Swierstra, T. 2003. *Pragmatist Ethics for a Technological Culture*. Dordrecht: Kluwer Academic Publishers.

Peirce, C.S. 1878, 1955. "How to Make Our Ideas Clear". In Buchler J. ed., *Philosophical Writings of Peirce*. New York, NY: Dover Publications.

Peirce, C.S. 1906, 1955. "Pragmatism in Retrospect: A Last Formulation". In Buchler J. ed., *Philosophical Writings of Peirce*. New York, NY: Dover Publications.

Pitt, J. 2000. *Thinking About Technology*. New York, NY: Seven Bridges Press.

Pitt, J. 2011 (forthcoming). "Successful Design in Engineering and Architecture: A Plea for Standards". In Braun H.-J. ed., *Creativity in Engineering, Mathematics and the Arts*.

Sellars, W. 1963. "Philosophy and the Scientific Image of Man". In *Science, Perception, and Reality*. London: Routledge and Kegan Paul.

Chapter 4
Don't Talk to Me

I'm Listening to MY Music

The iPod is one of the pernicious developments of recent technological innovation. It, even more than electronic gaming, has fostered what seems to be the ideal environment for the social solipsist. Electronic gaming, once the scourge of mothers and fathers trying to communicate with their kids, has now evolved into a social phenomenon, where groups participate. But iPod owners revel in the splendid isolation provided by a hand-size player and a pair of earphones.

One of the more joyful sounds is that of people engaged in conversation. Talking to one another is the most exhilarating form of human activity. The lilts, the accents, the rhythms, the modulation of sound is more complex and more rewarding than even, for example, a Mahler symphony. The sound of the human voice is a joy and it brings joy. The iPod, however, has managed to do what even Big Brother could not: silence that voice. Worse, it has turned iPod users into antisocial beings, those who avoid human interaction. The spontaneity of the social has disappeared and the silence of the anthropoid now rules.

Lest you think I exaggerate, take a look around you. The subways are silent. In a recent issue of *Wired*, a reader wrote to ask if was ok to tap another subway-rider on the shoulder whose iPod volume was turned up so high he (the iPoder) couldn't hear his cell phone ringing.[1] It used to be the case that the most vibrant place in the Philosophy Department at Virginia Tech was the end of the hall where the graduate student offices are. If you got bored with what you were doing and wanted to liven up a bit, you would wander down there and you could always find a good, loud, animated philosophical argument in progress. Today, there are few sounds down there, only drones sitting at their desks plugged in and staring vacantly. If you interrupt one of those reveries, you are met with surly stares and impatience.

Walking down the sidewalk on a college campus you used to hear students greeting one another, yelling to friends, arguing, making plans for the evening. Now: silence, walking slumped over, staring at the ground, plugged in. IPod users avoid

Reprinted by permission of Open Court Publishing Company, a division of Carus Publishing Company, from IPOD AND PHILOSOPHY edited by Dylan Whitkower (Popular Culture and Philosophy series Volume 34), (c) 2008 by Carus Publishing Company.

[1] "Earth to Rocker: Reality calling" in *Wired*, March 2008, p. 50.

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_4,

eye-contact. They don't want to engage in conversation. They want to listen to their music. Their entire body language signals avoidance of human interaction.

Consider the following situation. We all love to go to the beach. There is something primordial about the pull of the ocean. Then there are the sounds of the waves, the birds, children laughing and parents yelling at them to be careful. But the beaches have been invaded by the brainsnatched; the iPoders. They walk the beach heads down, listening, not to the sounds around them, but to their music. Why go to the beach if you are going to avoid everything that is there? An early morning stroll amid the sounds of gulls and pelicans calling, with the waves lapping at your feet is one of the more relaxing things to do. What makes it relaxing are the sounds, the ambience of the waterfront, the feel of the sand between your toes. But can you have that experience when you are plugged in? I doubt if you even feel the sand.

So what is so wrong about all this? It might be argued that it is I who has missed the boat. What the iPod does, it will be argued, is provide people the opportunity to disengage from the roar of contemporary living, to collect their thoughts and even to meditate. There is, it is said, simply too much noise in the world today, and the iPod provides a means for limiting the impact of that noise on our fractured and stressed being. It helps create a haven wherever we are and for the most part in whatever we are doing. What is wrong with that?

Well, nothing, as such. That is, there is nothing *wrong* with it. The iPod itself is a piece of technology. As such, it is neither right nor wrong, good nor bad. In another place, I argued that it is the use to which individual technologies are put that gives us the context in which to say a technology is good or bad (Pitt 2000). Here I would like to extend that somewhat and argue that it is not just the use to which a technology is put that allows us to make normative judgments regarding it, but the consequences of its use as well. It is the consequences of using the iPod as an escape mechanism that are so bad.

In *Brave New World*, Aldous Huxley offered us *soma* to achieve the desired state of bliss and state control. Who would have thought that the contemporary version of an imagined science fiction drug would arrive in form of a music player? But it has. The iPod is our form of soma. And the reason it is so successful is that each one is individually programmed by its owner. We, or should I say, Apple, has found the way to achieve the perfect isolated state of bliss by having iPoders program their own version of musical heaven. No need to worry if this or that version will fit all. Furthermore, the iPod has overreached even Linus's blanket as the ultimate comfort giver. We now have iPod stations and special speakers so that when you unplug you can keep the music going. But what happens when two people who share an abode each have an iPod? Do they share the station, or does each have his or her own and do we then have iPod station wars? I doubt it, since the whole point of an iPod is to avoid the social and conflict is social – my guess is that one or the other will simply plug back in and avoid conflict.

So now we have come to the set of iPod consequences that are most pernicious. As alluded to above, the iPod is an antisocial tool. By that I mean that the consequences of using iPods to create your own haven, into which you can retreat and ignore the world around you, are dangerous. The ubiquitous use of the iPod may be one of the final nails in the coffin of social skills.

So what is the big deal, what is so important about social skills? Social skills are the means one has to interact with others in a productive fashion. The ability to interact with others in a productive fashion is the key to technological innovation and to a successful democratic government.

Technological innovation in our high-tech world is a product of lots of brainstorming, group projects, feedback loops and team building. Even Steve Jobs had a co-inventor working with him. Technological innovation is not the product today of single individuals working alone in backrooms. It is the result of ideas being tossed around by people willing to try something new, something they may not have invented all by themselves, and work with it and others to improve, modify, and convince others to build and market the gadget. Good ideas do not speak for themselves. They need advocates, and they need advocates at all levels of their development, from the first glimmer of a thought to the polished thing in front of you. The ability to articulate your ideas and to be an advocate for them also calls for different types of skills. Articulating your ideas requires that you have the ability to express clearly your thoughts using readily understood means such as clear language, obvious diagrams and useful metaphors.

Being an advocate for those ideas requires all of the above and more. The "more" is the ability to successfully interact with others – to know how to read body language, the ability to present yourself as open and approachable. To develop those skills requires experience and lots of work. These skills will not come to you when are you slouched over a desk, lost in the world of your iPod.

What worries me about the lack of argument in the graduate student offices is that it is during those informal discussions that philosophy students develop their skills at argumentation – skills that require more than rigorous logic; skills that require command of rhetorical strategies and knowledge of how to make eye contact and good use of body language. These students may know the details of Kant's transcendental arguments, but if they can't defend their interpretation in person, on the hoof, then their futures as successful philosophers will be severely limited. Likewise, their futures as teachers and, more importantly, as productive members of society at large.

Good teachers must be able to interact with their students in ways that draw students into the discussion and help them expand their own skills. You might object that to be a philosopher you don't have to be a teacher. Ok, then – except for the very rare individual like Martin Luther King or Ghandi, the real philosophers of the twentieth century – philosophers are professors. How are you as a professional philosopher with a Ph.D. going to make a living? There are no more patrons today.

Now, this argument does not apply only to philosophers. There are so few jobs available today that don't require social skills that anything that impedes their development ought to be carefully scrutinized.

Above, it was also claimed that social skills are necessary for one to be a productive *member* of society. I would like to take this one step further. Social skills are necessary for there to be a productive *society*. Today we are experiencing an

America in which social skills apparently are not deemed important. How can it be that we tolerate the rantings of radio shock jocks and TV talk-shows in which interlocutors scream at each other and show no sign of courtesy and respect? A society in which intolerance for the differing views of others and a fundamental lack of respect for the other is considered entertainment is a society in which the value of social skills has been lost. The ability to interact with others in a respectful and productive fashion is essential for a democracy to function. If we no longer think democracy is a good thing, then why not scream at someone you disagree with instead of entering into a civil discussion? Why not stick your earplugs in and tune out the world when taking to people becomes such a burden? Why not tell everyone else "Don't talk to me, I am listening to MY music!"?

Part II
Methodological Issues

Chapter 5
Against the Perennial

The Changing Face of Philosophy

As philosophers are wont to do, my philosophical career has been centered on the search for a philosophically adequate, i.e., universal and perennial, account of, in my case, scientific change. It has slowly dawned on me that efforts to find a single theory of how science changes seek the wrong grail. Not only has the search for a perennial, systematic, and universal explanation of scientific change been mistaken, but any such effort with respect to any of the major concepts we employ to discuss and analyze the sciences and our technologies is not only doomed, but also rightfully fuels the idea that philosophy is irrelevant. The idea of the perennial lays out the philosophical landscape as if it were static. However, a little reflection on the history of philosophy immediately tells us this is not so. The landscape changes, as it should, because philosophy is a living phenomenon responding to the issues and challenges of the day.

One of the ways the philosophical landscape changes is that new areas of philosophical concern develop. One of the more recent of these areas is the philosophy of technology. For most of the last century, it has been a field in search of a place to fit. Part of the failure to see the philosophy of technology as a legitimate part of the philosophical landscape is due to the current commitment by contemporary philosophers to the perennial, i.e., to philosophy as a set of fundamental questions that do not change over time. In what follows, I challenge that view.

5.1 Continents

To set the stage, consider the metaphor of a "landscape". The traditional view of philosophy can be understood in the language of continents. We traditionally break Western Philosophy into five major divisions – albeit arbitrarily. Roughly seen, they are Value Theory, Epistemology, Metaphysics, Logic and History of Philosophy. Further, they are usually portrayed as fixed in location and general outline. However,

---

Appeared in *Techné: Research in Philosophy and Technology*. Vol. 7 No.2. Reprinted by permission of the editor.

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_5,

if we see Western Philosophy as a 2,500 year conversation, it seems inevitable that any such set of distinctions when exposed to that amount of working and reworking will eventually collapse. I will return to that point at the end.

Nevertheless, if these are the continents, the question becomes on which one(s) do we find the philosophy of technology? Surely that depends on how the philosophy of technology is characterized. And given the range of issues that arise in the philosophy of technology taken broadly, it is not surprising to find outposts on each continent.

Thus:

- *Value Theory* – includes ethics, action theory, social and political philosophy, aesthetics, and moral philosophy, and it alerts us to normative questions concerning the origin of technological projects and the consequences of technological innovation and development, among others.
- *Epistemology* – including some areas of philosophy of science, philosophy of language, and logic, addresses question of the foundation of technological assessment, for example.
- *Metaphysics* – including some aspects of ethics, philosophy of science, aesthetics, philosophy of language, and logic, turns us to the status of technological objects and processes.
- *History of Philosophy* – all of the above plus as much gossip as you can squeeze in – is used to plot the impact of technological innovation on the human condition, etc.

Philosophy of technology is everywhere, hence, nowhere. On this view, the landscape metaphor fails to illuminate.

5.2 Tectonics

The failure of the landscape metaphor doesn't mean we have to abandon talk of continents. We can change our static model of continents to the dynamic view that continents sit on floating plates that are constantly moving and colliding with one another. The result of this change of model is that while there is much activity underwater, the major visible changes occur on the continental surfaces. This, I think brings us closer to a reasonable way to look at issues in the philosophy of technology and it may force us to conceive of philosophy in an altogether different light. For example, I propose the not unreasonable thesis that social, cultural, and economic forces play a role in the changing character of philosophy and that these can, at first pass, be understood in the context of the metaphor of tectonic plates as the plates themselves. Philosophical change is a response or reaction to social change and not a motivating force in and of itself. The changes and the meaning of a new philosophical landscape must be understood against the background of social and cultural change.

There has been a spate of new work from people like Peter Galison (1998), Davis Baird (2004), with antecedents in the work of Gaston Bachelard (1940, 1978) and Ian Hacking (1983), that have in common some epistemological concerns, especially the ways in which human knowledge is "embedded" in technological artifacts. Let me refer to this group of scholars and others worrying the same and similar issues as The New Technists. I say these are new issues – it may be that Heidegger (1977) had something to say about this – but, not in the way the New Technists have – for they are concentrating on actual artifacts and the manner in which they reflect human knowledge. The heart of the matter is that by concentrating on the way in which artifacts are the product of human activity, and yet stand alone in some sense, new dimensions of epistemological concern have been opened up.

For example – if, as the New Technists argue, artifacts embody knowledge, then we need to find an account of knowledge that makes sense of this idea – traditional philosophical discussions of knowledge such as justified true belief (JTB) won't work here. This issue of the definition of "knowledge", and similar sites, remind me of eruptions and similar geographical changes that occur when tectonic plates collide.

These kinds of issues can't be located in the traditional philosophical static continental landscape – they can better be understood, as emerging from the clash of tectonic plates, and the movements of the "plates" are powered by technological innovation and change. In short, the arrow moves from technological innovation to social change to philosophical change.[1]

Because of its long history, we tend to think that philosophical concerns have basically remained the same over time. This is due in part to our general insistence that philosophical questions are perennial, without really looking at their historical dimensions seriously – good historians of philosophy do this, but the rest of us seemed to have missed the message (see Ariew (1999)). We import current concerns and analytic techniques like formal logic into the work of earlier philosophers, making it seem as if our concerns were always present to the minds of now dead philosophers. In so doing we readily ignore the historical social contexts in which philosophy is done. When we pay attention to historical contexts we also see that few, if any, philosophical questions are perennial except in the most trivial sense.[2] It is only when we accept the historically contextualized nature of philosophy itself that we can truly understand its history, and only then can we understand the emergence of new areas of philosophical concern such as the philosophy of technology.

---

[1]Be alert to the fact that any such singular conception of order will surely be falsified by the historical record. What is proposed here is at best a schema to assist in reorienting our thinking about these relationships.

[2]Thus, when we ask "What is the Good Life", the appearance is of the same question Socrates asked. But, by virtue of the fact that the answer that would have satisfied Socrates will not satisfy us shows that, in a fundamental way, the same question has not been asked. Likewise, the answer that would satisfy us, if we only knew it, most certainly would not satisfy Socrates.

Today the issues that newly raise their heads concern how to understand or to characterize a world increasing dominated by man-made artifacts. The traditional philosophical questions of the good life or the nature of reality or knowledge have been reconfigured in the light of strange, even perverse questions concerning the things we have made. These are new problems. As noted above, one version worries the problem of how these artifacts embody human knowledge. Thus, to borrow from a talk by Davis Baird, we can read time off the face of a watch without knowing how the watch works, but we can only do that because the watch, as artifact, embodies the knowledge of its makers. This is not the same problem as differentiating between theoretical knowledge and craft knowledge, or between knowing how and knowing that because, it is claimed, the object needs, requires, nay, *embodies* the theoretical knowledge and the knowing that – but just how is the question.

A second manifestation of this new set of problems concerns the way in which artifacts embody *meaning* in the same way or similarly at least as paintings and novels do.

These two questions seem to highlight a shift in our philosophical focus. In the first case, it is not a question of what "knowledge" means, but how it is manifested in the product of our labor. Likewise, the question of the meaning of artifacts does not seem to signal a call for a semantic analysis of objects, but rather asks for a performative account of what the objects say about themselves and maybe even us.

The standard analyses of "knowledge" aren't going to work when we ask how artifacts embody human knowledge. Let's look briefly at a couple of tries: (1) in the standard account of knowledge as Justified True Belief (JTB) – our belief that something is the case counts as knowledge just in case that belief is justified and it is true. When we turn to artifacts the epistemic status of belief is not the issue – nor is truth nor is justification. Those questions don't make sense when we try to unpack the sense in which *an object* embodies knowledge, for objects do not have beliefs. (2) If we take a different account, say a pragmatic theory of knowledge, the usability of an object or the consequences of its use tell us nothing about how human knowledge is embedded in the object itself. For one thing, the use to which the object is put may not reflect its intended use, e.g., the birth control pill.

Similarly when we turn to meaning. If we ask the question how objects bear meaning, we are not asking for the significance of the object in our lives. We are asking for something different, for, for example, the significance of the object in its own right.

So, on the one hand, it appears that if we take seriously the questions raised by the increasing dominance of technologies in our world we will have to ask and solve different kinds of philosophical questions. But, on the other hand, this seems odd.

Both the possibility of different kinds of questions and the oddness, I propose, are due to a continued reliance on the old model of fixed continents. That is, if these really are new questions, then to answer them we have to abandon the view that philosophical questions are perennial. If these really are questions in the philosophy of technology, then we seem to have to do one of two things: either (1) add a new field called the philosophy of technology and try to figure out how it fits in with the others, or (2) subjugate the rest of philosophy to these newer concerns. The first of these options seems to me to characterize the history of the philosophy of technology in the twentieth century. Efforts to legitimize the philosophy of technology have proceeded by trying to set it up as a new branch of philosophy. The second approach keeps trying to raise its flag, but it is having a hard time finding a mountain top on which to stand – see the scene in John Bormann's 1981 film, *Excalibur*, when Arthur, flush with victory creates the fellowship of the Round Table. You need a mountain top from which to declare a new order.

On the other hand, what if we take the plate tectonics metaphor more seriously? The collision of plates and continents does not leave one the victor and the other vanquished – nor does it result in different kinds of things, we still have mountains, rivers, valleys, just in new arrangements. Rivers get diverted, new mountain ranges rise and where there were oceans, deserts come to be. What does this say about these new issues in the philosophy of technology? It strikes me that among some of the things that can be offered is a possible explanation of why these new issues stand so starkly against the philosophical landscape and still do not ring true from a phenomenological perspective.

5.3 Meaning

Coming to the issue from an analytic stance, consider the question "Can artifacts mean anything in and of themselves?" I think not. The rejection of the idea that artifacts have meaning or the ability to communicate their meaning is based on the same reasons we would conclude novels do not mean anything. This is not to say that the works, sentences, paragraphs, etc. in a novel cannot be said to express this or that. It is, rather to deny that a novel, or a hammer, in and of itself has some independent "meaning"

From an analytic point of view, to say a word means this or that is to imply there is a semantics involved – i.e., a set of rules for determining how a word or an object can be said to stand for this or that. In short, for an object to "bear meaning" is to place that object in the logical space of symbols, which is to say we know how to reason about it – how to make inferences with regard to it. This is my version of the Peircean/Sellarsian/Wittgenstein "use" theory of meaning, wherein the meaning of a word is its use or the role it plays in the complex of inference patterns of our language. We take symbol x to "stand for" y – but there is nothing inherent in x that it be a "stand for" sort of thing. That is, when we accept that x is a stand in for sort of thing, we accept it as a symbol in our system of symbols – on the other hand to say of an object that it has meaning *simplicitur* is to say something different, or is it?

Two viable options present themselves: (1) We can say objects have no meaning outside our symbol systems, or (2) objects do have or bear meaning – it is just that the symbol system by which they can be said to have meaning isn't ours.

If we agree that objects only have meaning in the context of one of our symbol systems, then we *can* talk about the meaning of a novel, because we can say that it stands for an expression of this or that sentiment or insight into the human condition, and we know how to reason about those sorts of things. The meaning of "meaning" in this context is metaphorical.

On the other hand, if we say of an object that "it has meaning or bears meaning in and of itself", and by that we mean "independent of our symbol systems", then we are on the dubious road to the metaphysical reification of meaning. This is an unjustified and unjustifiable move. It gives us meanings in the world – something that cannot be determined outside of some interpretive scheme. There are many such schemes. Assigning meaning to anything requires the use of some such scheme. But even though some scheme or other must be used to assign meaning to anything, it is not necessarily the "right" scheme. To determine that we do in fact have the right scheme by which to assign meaning to a thing implies that we can argue about the object's *real* meaning independent of seeing it as a symbol while using a symbol system to accomplish that end. That makes no sense. That is, we cannot conclude that objects embody meaning simpliciter using a symbol system in which meaning is contextual in some form or other. To say that we can is to wallow in post-modernist fantasies.

Another option, one equally unpleasant, is to suggest that the real meaning of x is determined through a symbol system alright, but one to which we do not have access – something like an appeal to God – at that point reasoned discussion has left the room.

From a Pearcean standpoint, words, objects, novels are meaningful only to the extent that they are symbols and subject to a semantics, i.e., a theory of meaning. This means that objects cease to be meaningful when we no longer know how to reason about them or how to use the symbols for them. There is a warehouse in Maryland that belongs to the Smithsonian Institution. It houses a collection of what we believe are clearly human made artifacts whose use we cannot figure out – here it is, what it does no one knows – those artifacts have lost their meaning – or more precisely, we have lost the ability to speak meaningfully about them.

But does the same case hold for knowledge? That is, can we use the same tactic we used on the notion of meaning to explain away the notion that artifacts embody human knowledge? This one is tougher. The difference is that we can watch a watch-maker – the old fashioned kind – make a watch. We can watch him select the gears (perhaps even make them) and fit them together with appropriate winding mech-anisms – take it apart and redo it if the first effort doesn't fit his fancy. We can observe his efforts to fit it in a case and test it for its time keeping accuracy. In this case we can clearly see the knowledge being put into the watch. Or can we? Aren't we employing a similar interpretive scheme – inferring intentionality where they may be none? For the "skilled" watchmaker substitute a chimpanzee. I am sure you can picture a chimp sitting in front of several boxes of components, picking up one piece and cocking his head while turning it around, then putting it down and picking up another piece and eventually assembling a watch. Or imagine a set of boxes of watch components on a shelve and a cat jumping up on the shelve and then knocking over one box after another as it tries to negotiate its way, spilling the contents on the table below, with the result being a watch! In this case you are unaware of how the watch was put together. This last case may appear far-fetched. Regardless, doesn't the same point hold? However the watch was constructed or put together, given the watch, we *infer* that there is something called knowledge in the thing – it isn't that the thing *has* knowledge in it somehow – rather, we infer that to be the case.

Let us return to the Maryland warehouse. Here we have loads of artifacts – "clearly" human made objects, of whose function we have no notion. As noted above, they have no meaning because we infer the meaning of an object and if we don't know what they are supposed to do, we can't import meaning to them because we can't make any further justified inferences about them.

What about knowledge – could knowledge be embedded in them? I think not. For to answer in the positive requires a host of assumptions to which we are not entitled. That it looks like a human made artifact does not mean it is one. Further, assume for a minute that it is a human made artifact – from the fact that a human made it, it doesn't follow that the human knew what it was doing when it did so. And if it didn't know what it was doing, didn't intentionally do action A to achieve result B, then in what sense does the machine have knowledge embedded in it? And, finally, consider the Rube Goldberg machines – marvelously entertaining devices that in our teleological world *do* nothing. They are the equivalent of a tinkerer putting parts together with no end in mind and no plan – just putting parts together – in what sense does that machine have human knowledge embedded it in?

Where does that leave us? There seems to be at least one sense in which we can say that an artifact embodies human knowledge – which was the case of our watching the watchmaker build a watch. But, then again we are only entitled to infer he knew what he was doing – and even that inference is restricted, for if he did what he did because he was merely trained to put part A in slot Z and he has no understanding of what he is doing, then the sense of knowledge begins to erode. On this approach objects do not embody knowledge. At best we infer they do and sometimes we may be right.

Finally, let us return to our philosophical landscapes, plate tectonics, and a possible explanation of what this is all about in some sense of "all about". The philosophical landscape has changed. The increasingly technological – i.e., human-made, features of our world are forcing us to ask some different questions. Some of these questions in turn arise from two things: (1) not recognizing the changing landscape and (2) from remaining immersed in the old way of thinking about things. If we accept the view that philosophical questions ought to be understood as historically contextualized, then the old way of thinking about *things* has some historical basis to which we should return, if only briefly. Here is a first stab at identifying that basis.[3]

Descartes (1637) created our current problem by making a mess of things when he introduced the mind/body distinction. He separated what he wanted us to think of as our essential being – our minds – from the peripherals like our bodies. The

[3]Clearly what follows is at best a potted history of an idea. However, I am convinced that with some work it can be filled out and elaborated. What I fear is that it will look a like Heidegger's view, but perhaps without the gloomy conclusion.

mind/body distinction gave Marx some philosophical license to describe the result of the relation between the man on the factory line and the objects with which he interacted in terms of alienation. That is, unlike the classical craftsperson, whose labor could be seen as transformed into an object he could use to procure goods, the assembly line person never really gets the opportunity to identify with the product of his or her labor being involved with only one small part of the construction of an object. This point becomes instantiated into a permanent bifurcation between human beings and their world. Descartes set up the general formula and Marx gave it specific content.

However, just as there is an historical context for understanding our current state, there is an historical antidote. Hume (1739, 1888), in part reacting to Descartes, tried to set the philosophical discussion back on track but, for a variety of political reasons, he never got the hearing he deserved until now. In Books Two and Three of the *Treatise*, Hume lays out a thorough-going naturalism which places us back in the world and as part of it. We are not things that look at the world through odd lenses and think about it – we are fully engaged parts of the world, we cannot be alienated from our labor for it is what we are, things that we do. What we do may or may not be intentional. The meaning of what we do is what we say it is, and over time that will change, depending on our perspective, and when we (the human race) are gone there will be no more meaning. Likewise for knowledge. To seek to place meaning and knowledge in things is to adhere to the Cartesian mistake of taking us out of the world. To see our objects as apart from us is to continue to buy into Marx's view of the separateness of things from their makers and it represents a philosophically bankrupt way of trying to correct the new problem of seeing how objects embody human knowledge.

Ironically the change in the landscape, which comes from a variety of factors, some of them the results of modern scientific research, is the reassertion of *homo faber* as the model for thinking about philosophical issues and the displacement of the Enlightenment ideal of reason. It also gives us the basis for rethinking our approach to the questions we raised earlier about knowledge and meaning.

The objections raised earlier were based on a contemporary analytic framework of sorts. But if philosophical landscapes change, maybe we are speaking from the bottom of the ocean and don't realize it. That is, if, contrary to the Cartesian/Marxist perspective, we are in fact *in* the world and part of it, as opposed to being commentators on what we see from some abstract point above it all, then maybe we need to rethink the status of human made objects. If it can be said of humans that they have knowledge, and the objects that humans make are, as we are, both in and of the world, then we might be able to talk about objects having knowledge in the same way that humans do. Or, to put it differently, maybe we need to rethink the category of human made object, i.e., artifact. At this point in time it is traditional to proceed with an analysis of things as this or that and then try to figure out how what we know got poured into this or that. Instead, consider what would happen if we accept as a primitive category in the scheme we use to think about the world, the assumption that artifacts, *by virtue of being human made objects* embody, exemplify, human knowledge? If we start there, with the assumption that in fact this is the case, then the question of “how” disappears and other questions can come to the fore; questions such as “what is the process by which we create our objects so as to embody our knowledge?” And through an analysis of that product we may discover that the object was constructed by rote training. Does that mean that the objects don’t embody knowledge? No, it simply shows us that while this person doesn’t have the knowledge that the object embodies, the fact that what he constructs reliably works as it is advertised to means that someone had that knowledge and packed it in such a way that it could be passed on by training. To reconceptualize the issue by rejecting the Cartesian dualism and embracing a form of naturalism changes the philosophical landscape, as long as we actually reject the old and embrace the new. The new view takes human action and its product as primitive and concentrates on the processes by which goals, desires and values are transformed by those actions and products. Abstract, perennial questions and universal answers are no longer, in this new configuration, the *desiderata*. Understanding the means by which we impact the world and how it impacts us is the new goal. Philosophy has changed, thanks to philosophers of technology, and that is good. We should stop trying to fit this very good new wine into old chipped and cracked bottles. It is time we gathered at the mountain top, looked around, and appreciated the new landscape.

References

Ariew, R. 1999. *Descartes and the Last Scholastics*. Ithaca, NY: Cornell University Press.

Bachelard, G. 1940, 1978. *The Philosophy of No*. Paris: Pellicanolibri.

Baird, D. 2004. *Thing Knowledge*. Berkeley, CA: University of California Press.

Descartes, R. 1637. *Meditations on First Philosophy*. Trans. M. Crewe. Indianapolis, IN: Hackett Publishing.

Galison, P. 1998. *Image and Logic*. Chicago, IL: University of Chicago Press.

Hacking, I. 1983, *Representing and Intervening; Introductory Topics in the Philosophy of Science*. Cambridge: Cambridge University Press.

Heidegger, M. 1977. *The Question Concerning Technology and Other Essays*. Translated with an introduction by W. Lovitt. New York: Garland.

Hume, D. 1739, 1888. *A Treatise on Human Nature*. Oxford: Oxford University Press.

Pitt, J. 2003. “Against the Perennial”. *Techne: The Society for Philosophy and Technology Quarterly Journal*, 7, 2.

Chapter 6
Philosophical Methodology, Technologies, and the Transformation of Knowledge

6.1 Introduction

There are many methodological approaches employed by scholars working within the general area of the philosophy of technology. This methodological richness is rightly the cause of a certain degree of pride felt by many members of The Society for Philosophy and Technology. For we clearly work hard at encouraging and maintaining a form of methodological pluralism that has evaded the mainstream American philosophical community. Philosophical provincialism also seems to threaten productive interchange between divergent points of view on the Continent as well, witness the almost total breakdown of communication between the so-called postmodernists and analytical philosophers. Nevertheless, despite our sense of having preserved diversity of viewpoint within the Society for Philosophy and Technology, there is a different perception of our work held by the larger philosophical community. To put it bluntly, work in the philosophy of technology is deemed largely marginal.

And yet, it is equally clear, given the pervasive character of technology in modern life and with the apparently increasing rate of technological innovation and dispersion, its effect on our ways of living and on our values, that technology is a central, if not the central, feature of the human world. As such, it demands philosophical examination in all its various aspects and manifestations. So if technology is so central, then why is the philosophy of technology so marginal?

Of the many possible reasons for this state of affairs, I want to concentrate on the one I believe to be at the heart of the problem. After examining the cause of the problem, I will suggest an alternative approach to exploring questions in the philosophy of technology, which approach ought to make our work not only more acceptable to the philosophical community, but also to society at large. What I am about to put before you is framed by the fear that because we have taken the tack we have up to this point, that we will continue to be ignored and to experience further

Appeared in Technology and Ecology, Proceedings of the VII International Conference of the Society for Philosophy and Technology, edited by Larry A. Hickman and Elizabeth F. Porter, 1993. Reprinted by permission of the editors.

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_6,

disintegration as an organization. I also want to emphasize the need to find a way to maintain and encourage the diversity of viewpoints that makes this organization so exciting. So, another precautionary note I am not trying to impose a single point of view on us. Rather, I am trying to find a way to accommodate and nourish our diversity while at the same time providing us with an entry into the larger philosophical dialogue. I am convinced that the work of the members of this society should and could be of immense value to our philosophical colleagues.

6.2 What's Wrong with the Philosophy of Technology?

The problem, as I see it, is that most of the work in the philosophy of technology is perceived, rightly or wrongly, mostly wrongly, as being expressed within one highly charged ideological framework or another. Furthermore, that framework is perceived to be highly antagonistic to technology. It is not just that many of us are critics of specific features of various technological intrusions in our lives. This by itself would not be enough to cast suspicion on our work. Not all social critics are viewed as working at the margins of the philosophical world; consider much of the work currently appearing under the heading of social and political philosophy. Social criticism per se is not the problem. It is rather the context within which the criticism occurs that is the villain.

I labeled this context "ideological." By that I mean two things: (1) what sets the context is a discreet kind of conceptual scheme; (2) to call it a conceptual scheme is to identify it as a structure for thinking; as such it can be misused by individuals who, for whatever reason or cause, assume the guiding principles of that scheme are inviolate. In short, an ideology is a pathological conceptual scheme. By that I do not mean that the scheme is pathological, but rather, that the people who employ it do so in a manner that can be interpreted as such. It is, of course, people who are pathological, not conceptual schemes. What does this mean for us? It means that one of the reasons some work in the philosophy of technology is ignored or marginalized is because those outside the circle view that work as being conducted within an isolated and insulated context, a context whose users refuse to admit any challenge to its assumptions. And that means that it is not seen as being includable within the wider philosophical discussion.

6.3 The Aim of Philosophy

Now, this diagnosis of what I have called "our problem" assumes a certain view of philosophy. I would like to lay out that view, my view, and discuss its implications for our work. My view of philosophy derives in large part from a rather remarkable characterization of philosophy by the twentieth century American pragmatist, Wilfrid Sellars. According to Sellars,

> The aim of philosophy is to understand how things in the broadest sense of the term hang together in the broadest sense of the term. Under things in the broadest possible sense', I include such radically different items as not only cabbages and kings', but numbers and duties, possibilities and finger snaps, aesthetic experience and death. To achieve success in philosophy would be, to use a contemporary turn of phrase, 'to know one's way around' with respect to all these things, not in that unreflective way in which the centipede of the story knew its way around before it faced the question, how do I walk?' but in that reflective way which means that no intellectual holds are barred. (Sellars 1963, p. 3)

In short, the aim of philosophy is to make it all fit together and to know how it fits together. It is to know what all the pieces are and to know how to move them around and to know how they are connected. That much is Sellars. From now on it is me and my interpolation of that view. This conception of philosophy assumes no privileged point of departure. It does not assume, for example, that all discussions must be carried out in terms of the ethical, moral, or value implications of the topic under discussion. It does not assume that philosophical issues are questions of power seeking. It does not read every question of technology as one of political correctness. It does not assume that in the process of finding out how it all hangs together there is already one set and agreed upon value system that determines the relevance and importance of all other considerations. It means, likewise, that not all philosophical issues are to be approached as matters of language. Further, it means that there must be some demonstrable relevance of philosophical ruminating to living, which is, minimally, the "finding one's way around" part.

The philosopher then must be constantly reevaluating what she knows as she finds new things and new problems to consider. Under this conception of philosophy, a philosopher cannot come to a problem with a predetermined solution or approach unless he or she is willing to allow that those assumptions are only starting points and can, nay, must, be examined for their appropriateness as part of the process of inquiry and synthesis. This view of philosophy also means that the philosophical process of trying to make it all hang together in the broadest possible sense of "hang together" is not and cannot be complete. For as we discover new things about us and about the world in which we live, we must be constantly reassessing and reformulating our account of how it all fits together. If such constant reformulation and reevaluation is the steady state, then any assumption of privilege must be mistaken. And just to make the point as clean, or as obnoxious, as it can be, the opposite view, i.e., the assumption of an a priori privileged perspective in the doing of philosophy, I will call fascist philosophy. It applies to all schools of philosophy from analytic to existentialist, from empiricist and rationalist to politically correct. You simply cannot do philosophy with the Sellarsian aim in mind if you use the crutch of a school or a single method.

And now, I suppose I would do well if I were a bit reflexive. Is the claim of no privilege itself not a privileged claim? I think not. For it does not assert in a non-refutable way that there can never be a privileged perspective. Rather it says that there is no a priori privileged point of view that one can justifiably bring to any and every philosophical problem that is immune from challenge. One must start somewhere and somehow. What I am saying is that both starting point and methodology are constantly up for grabs, depending on how we are doing. I am also suggesting that I do not think that it will ever be possible to establish once and for all time a privileged point of view because we are constantly in a state of coming to know new things. That is an empirical claim that may be false; however, I am willing to live in that state of indeterminacy for the moment. But note that if we do think that we have achieved that exalted state of privileged point of view, then we will also have come to the end of human creativity, for only when there is nothing new under the sun, to quote the preacher, can such a stance make sense.

Let us now enumerate some of the parts to be related in this constant effort to figure out how it all hangs together. In the context of discussing that technological marvel, the space program, and taking our cue from Sellars, we are looking to see how such disparate things as scientific instruments, scientists, space shuttles, laboratory experiments, our concepts of knowledge, science, and standard experimental conditions, hang together and cohere not only with themselves but with challenges to the value of the space program, cries for feeding the needy and housing the homeless, and various assertions about the way in which large scale technological projects such as the space program politically disenfranchise us. It is in this context that I wish to examine the topic of changing knowledge.

6.4 Changing Knowledge

Knowledge changes. Furthermore, it changes in two ways. First, in terms of its content, i.e., in terms of the specific things we know. Second, in terms of what we mean by the very notion of "knowledge." Not only do we know more than our forebears, but we can and do expect our children to experience even greater wonders than we have. The fact that what we know changes is not in dispute. However, it remains a question, if not something of a mystery, as to how the content of our knowledge changes, and what affects such changes in the content of our knowledge have on our concept of knowledge.

According to the current popular story, science is portrayed as a major player in the process by which knowledge changes. Beginning at least with the Copernican Revolution, science, it is alleged, has explored the world around us, revealing nature's secrets in increasing detail and at an accelerating pace. I think this picture is wrong.

Science is not what forces us to change and correct what we know. Science is not responsible for our new vision of an expanding universe. Science cannot be credited with revealing, in ever-increasing detail, the structure of, for example, the human genome. At least, science cannot do all these things by itself. Rather than credit science with increasing our knowledge, I want to argue that it is the technological infrastructure of science, rather than science itself, which is responsible for these monumental changes. In other words, the popular and well entrenched view that science is responsible for how knowledge changes is a myth, a popular, well-entrenched myth, but a myth nonetheless.

Changes in the technological infrastructure are what makes it possible for us to cast off false views and replace them with what we hope is a continuously improving understanding of our world. The picture of science as the major mover responsible for the transformation of knowledge is inaccurate because it leaves out the role of technology, both in the generation of knowledge and in the development of science. In the popular story, when and if technology is begrudgingly included in the story of human progress, it is always as an afterthought, as, at best, a nice benefit of scientific research. That is not only false, it is historically myopic and potentially dangerous. One last point here. When I refer to the technological infrastructure of science I am not just indicating scientific instruments used in experiments. No, I mean such complicated social structures as the control room at Houston complete with enormous television screens, computers, various computer programs and telephone hookups, the building itself, the communications network, the space shuttle, etc., the astronauts, all the personnel required to make the whole thing work. In other words, I mean the infrastructure that makes the doing of science in space possible.

What I intend to do now is look first at the major change in our conception of scientific knowledge which Galileo helped establish. I will then look at how, once technologies were accepted as part of the knowledge generating process, they have become increasingly crucial to it – so crucial that they have in fact transformed our conception of knowledge anew, under our very noses, without our being aware of it.

One way to see how knowledge changes is to consider changes in the criteria by which something is said to qualify as knowledge. When the criteria change, the things for which it is the criteria can be said to change. In previous work I have sometimes called these criteria "values." I am changing terminology here to avoid unnecessary confusions. But for those who know my previous work, I am not changing much by way of the general position. The criteria associated with knowledge I will call epistemic criteria. What I have in mind here are such notions as truth, conceptual economy, usefulness, justification, and simplicity, among others. These are the notions according to which we determine if some claim or other is going to count as knowledge. Thus, while the statement, "There are 231 mountains on the moon," may or may not be true, in order for someone to say that they know and not merely believe that there are 231 mountains on the moon, that statement must be true. Truth can function in two different ways: first as a property of a statement, i.e., this statement is true; and second, as a criterion by which we determine if a statement qualifies as knowledge and if its user can be said to know something. Because we are using the notion of truth to determine the acceptability of the claim being made, truth functions here in its second sense, as a criterion; and that is what I wish to look at in greater detail. Since epistemic criteria are the key to what we mean by knowledge, if the criteria change, then what counts as knowledge changes. That means that what we thought we knew 200 years ago, today we may decide no longer counts as knowledge. Furthermore, this need not be due to the discovery of new facts that cast the old knowledge into doubt. It can be a simple case of our changing our mind as to what counts as knowledge.

To understand what this means and what it implies for our conception of knowledge, let us take a look at an historical example of change in epistemic criteria.

The example concerns Galileo's efforts to incorporate mathematics and observation within what were then the standard criteria for scientific knowledge. After looking at Galileo's arguments, I will turn briefly to our contemporary situation. If, as Galileo's case shows us, the development of new instruments opens up scientific research by way of developments in the technological infrastructure, then we need to ask what happens to our understanding of what counts as knowledge when, as is now the case, the conditions for observations change again, such as in the new environments of space. Most importantly, we need to consider the effect on our conception of knowledge made by the entire technological infrastructure that allows us access to space and how that infrastructure, employing as it does new techniques and machines and data processing devices, impinges on our understanding of the components of knowledge. In short, in what sense is the data gathered in space that is digitalized and transmitted by radio to a receptor on earth, retransformed into codes, and finally, through further computer programs reconstituted into an image to count as an observation? But first, Galileo.

6.5 Galileo and Mathematics

Galileo did two things that had a major impact on the then prevalent Aristotelian conception of knowledge. He insisted on the role of mathematics in scientific knowledge and he changed our understanding of what counts as an observation through the use of his telescope. Let us look at these in order, first the role of mathematics.

Shortly after he perfected the telescope in 1609, Galileo secured a job with the Duke of Tuscany as his chief mathematician and philosopher. Galileo's title at Florence is important because it tells us something about the organization of science in the seventeenth century and that has important ramifications for understanding what constituted knowledge at that time.

In the sixteenth and seventeenth centuries Italian universities were dominated by the Catholic Church and the conceptual framework of Aristotle that had been acquired primarily through the writings of Thomas Aquinas. The predominant way of thinking about the world, that is, the Aristotelian/Thomist framework, also included a taxonomy or structural organization of the sciences in terms of what was supposed to be each science's proper domain of inquiry. There were many deep and maybe even perverse reasons for this structure, but we cannot risk going into them or we will find ourselves in the middle of an Umberto Eco novel. For our purposes, it is enough to remember that the sciences were divided into major sciences and subfields, not unlike today.

Furthermore, also like today, their precise order cannot be said to be set at any time during this period. There were many variations, depending on a variety of factors. But Galileo had intimate knowledge of and seemed to be concerned with the doctrine most favored at the Collegio Romano, the home institute of the Jesuits, located in Rome. On this account there are five total sciences, understanding by the notion of a science a source of knowledge. These five were a science of God, a science of intelligences, a science of being in common, a science of natural bodies, and a science of quantity, i.e., mathematics. It is most interesting for our purposes that mathematics was not to be applied to corporeal substance, that is, matter. The subject matter of mathematics was "nude quantity," i.e., matter considered only in terms of necessary connections and not through relations of cause and effect. That means that the proper subject matter of mathematics was abstract relations among quantities. Mathematics could not be applied to physical matter (for an account of the structure of the sciences in Galileo's time see Wallace 1984). In other words, physics as we know it today, i.e., mathematical physics, was not possible.

One way to view Galileo's methodological research program is to see him, and a few select others like Kepler and Clavius, as engaged in the preliminary conceptual battles that made it possible to make sense of and to accept the views of the new mathematical physicists like Descartes and Newton. One of Galileo's primary considerations was to incorporate mathematics into our very conception of how to describe and reason about the world. He makes this case in a number of places, but most notably in his famous *Dialogue on the Two Chief World Systems*. In that work, Galileo urges us to use mathematics wherever we can, especially in talking about the physical world. His criticism of Aristotle is not that his arguments are bad ones, but that they could be so much better if framed mathematically. Consider what he has to say at the very beginning of *The Dialogue*. It is the first argument discussed, namely the argument of the followers of Aristotle, called Peripatetics, that the earth cannot be a planet and move as do the other planets. The first step in the Peripatetic argument is to show that the earth is complete and perfect. (We will not consider the rest of the argument; my objective here is merely to demonstrate Galileo's point about the use of mathematics.) First Galileo presents Aristotle's argument:

> [The earth] is not a mere line, nor a bare surface, but a body having length, breadth, and depth. Since there are only these three dimensions, the world, having these, has them all, and having the Whole, is perfect (Galileo 1632, 1967, p. 7).

Galileo's response is not to attack the proof and show that it is wrong, but rather to help it along. To this end he notes,

> To be sure, I much wish that Aristotle had proved to me by rigorous deductions that simple length constitutes the dimension which we call a line, which by the addition of breadth becomes a surface; that by further adding altitude to this there results a body, and that after these three dimensions there is no passing further so that by these three alone, completeness, or so to speak, wholeness is concluded (Galileo 1632, 1967, p. 8).

He then goes on to draw a little diagram using the basics of geometry to prove the very point.

Galileo's strategy throughout the *Dialogue* is to argue for replacing and/or augmenting the convoluted semantic arguments of Aristotle's followers with mathematical proofs. Each time an Aristotelian proof is offered, he counters with a mathematical one making the same point, only in more intuitive and obvious fashion. His apparent objective is not to disprove Aristotle, but where possible to show how his ideas can be improved by employing mathematics, and, where that is not possible, to argue for the correctness of the mathematics. He concludes,

> It is best to have recourse to a philosophical distinction and to say that the human understanding can be taken in two modes, the intensive or the extensive. Extensively, that is, with regard to the multitude of intelligibles, which are infinite, the human understanding is as nothing even if it understands a thousand propositions; for a thousand in relation to infinity is zero. But taking man's understanding intensively, in so far as this term denotes understanding some proposition perfectly, I say that the human intellect does understand some of them perfectly, and thus in these it has as much absolute certainty as Nature itself has. Of such are the mathematical sciences alone; that is, geometry and arithmetic, in which the Divine intellect indeed knows infinitely more propositions, since it knows all. But with regard to those few which the human intellect does understand, I believe that its knowledge equals the Divine in objective certainty (Galileo 1632, 1967, p. 103).

Since Galileo is here talking about mathematics, it might be thought that he is concerned only to claim that we can know some of the truths of mathematics as well as God can. And if one were to concentrate on this short paragraph alone that would be so. But if we look further we will see that more is going on. First, let me emphasize one thing from what I just quoted. Galileo says that it is the mathematical sciences alone that provide certainty in knowledge. This means that none of the other sciences do. So, to the extent that you can have knowledge at all, it must use mathematics. Second, in the rest of Day, he always follows each mathematical proof with an empirical example. In many ways this resembles the old logical positivist's idea of interpreting a formal abstract language using only the observation terms of normal language to give you the language of science. Galileo does not say this, but the regular way in which he follows every proof by an empirical example suggests that he is urging his reader to draw the parallel between the points of the proof and the physical situation in the example. That is what gives you knowledge of the world.

The point of examining this dimension of Galileo's work is to provide an example of how, by emphasizing a new or different kind of epistemic condition, you can change the very conception of knowledge. Galileo urges us to consider an alternative epistemic criterion, alternative to the Aristotelian. For the Peripatetics, knowledge is based on the writings of Aristotle. If you are an Aristotelian, to show that a particular claim is a knowledge claim requires fitting it into the categorical scheme of the great philosopher, which precluded mathematics from providing knowledge of the world. On the other hand, in his new account of knowledge, Galileo insists on the value of providing mathematical proofs and empirical counterparts to those proofs. It is not enough, he is saying, to merely cite the words of some approved authority. This marks a major turning point in Western science and in our conception of knowledge. For, until the language of mathematics is acknowledged as a legitimate means of expressing knowledge, mathematical physics cannot develop into the powerful tool it has become today.

6.6 Galileo and the Telescope

Let us now turn to the impact of Galileo's telescopic observations on knowledge. Using his telescope to make careful observations of the heavens, Galileo saw things no one had seen before such as the phases of Venus and the moons of Jupiter. These observations had a devastating effect on the Aristotelian conception of the structure of the solar system. Let us call this the Aristotelian theory. This theory was complicated and yet very elegant. Beginning with the claim that there are only four elements, air, earth, fire, and water, and the principle that each element has its appropriate place, earth being the heaviest, its natural place was at the center. The universe had no top or bottom; therefore if all the earth matter in the universe seeks its natural place, which is down, it will come to the center of the void. So we get the view with the earth motionless at the center of the universe and the heavens in rotation around it. If you jazz this up with metaphysical assumptions such as the only motion appropriate to the heavens is perfect circular motion, and some later theology which argues that the earth being at the center is as it should be according to the Bible, there can be no other centers of action in the universe. Imagine their surprise when Galileo reports his observations of the moons of Jupiter, showing that there is yet another planet with moons rotating around it, i.e., another center.

There are other observations we could discuss and other effects, but the point here does not require that we do a detailed analysis. Basically, what we have is the impact of observations on the Aristotelian theory made possible by technological innovation. The impact is significant enough to force reconsideration of the theory, and ultimately it is responsible for its downfall. Galileo forced us to reconsider the adequacy of a theory which was used to explain why the universe is the way it is through non-empirical abstract metaphysical reasoning. In addition, by pushing for the superiority of framing knowledge claims in the language of mathematics and having them backed up with empirical observations, he makes it possible for Newton to turn to the work of Kepler for a set of purely mathematical relationships in order to put together a new theory to replace Aristotle's. We need only mention in passing the fact that in creating that new theory, Newton also had to invent a new form of mathematics, the calculus, to see how far we came in a short 100 years in our understanding of the criteria for knowledge.

Before we leave Galileo we should also note one further impact of the telescope. Not only did using the telescope produce observations that challenged Aristotle's theory. The use of the telescope challenged the very notion of what constituted a legitimate observation. Now we had the ability to go beyond the unaided human eye, which restricted what we could see, to view the universe. There were challenges to Galileo's use of the telescope. Some called the observations of the Medicean planets artifacts of the instrument. There were also claims that one could not use an imperfect instrument to view the perfect heavens. But what was now under dispute had not been under dispute previously: what does it mean to make an observation?

6.7 Space and Beyond

Let us jump now to the present and the impact of space-based experiments on our conception of knowledge and changes in the criteria associated with our contemporary account of knowledge. The March–April 1990 issue of American Scientist contained a piece entitled, "Effects of the Space Environment on Space Science," by Joselyn and Whipple. In that article Joselyn and Whipple carefully review the variety of factors which affect space-based experiments, causing instruments to produce what they call "ephemeral data." As they relate, the environment of space affects our experiments in ways hard to correct for. The solar wind, solar flares, electromagnetic radiation, all produce particles and forces we have to be aware of and account for. Likewise, the very materials we use to construct both our spacecraft and the instruments they carry actually interfere in the information to be generated. Some of these factors can be anticipated, others cannot.

What does this mean for our concept of knowledge? To begin with, one basic point seems settled. Given the random influence of the multitude of factors that interfere with our understanding of the significance of the data from space provided by our instruments, any conception of knowledge embodying any sense of certainty must be abandoned. Second, the very fact of space-based experiments forces us to confront in unmistakable terms the technological infrastructure of science, and the extent to which we depend on that technological infrastructure. We often speak in casual fashion of the link between science and technology. In so doing we generally take an ideological stance with regard to assessing their relative significance. Science, it is argued, is an example of pure intellect and *ipso facto* must be superior to any form of technology. But when the very possibility of the science is shown to rest on a massive technological investment such as the space program, the question of superiority should be, at best, blurred.

However, I want to argue, we should not rest content with just blurring our understanding of the relation between science and technology. We are in a position to dissolve old distinctions and reconfigure the entire relation. The fact of the matter is that space-based experiments would not be possible without the technology behind the space program. Thus, the fact of doing science in space forces us to face the fact that this science requires this technological infrastructure. The example of experiments in space may seem to force the issue in an artificial way. But, on reflection, it is easy to see that any mature science absolutely demands an extensive technological infrastructure. Where would microbiology be without the ultracentrifuge and a host of sophisticated machines and counters? Is it possible to do astronomy today without computers and computer programs, cameras, and mountain top observatories with a variety of telescopes, mounts, buildings, and supplies? Particle physics is almost too easy a target. But when all is said and done, the fact remains that contemporary mature science requires much more than a theory about a domain. Focusing on space-based experiments brings that sharply into view.

The relationship between science and its technological infrastructure that we come to see when we concentrate on these experiments also allows us to see the problems that emerge when the technological infrastructure mediates the science.

The claims of scientific theories are seen now through the lens of the machines and devices of the infrastructure. This raises the question of the extent to which the technology transforms and influences the formulations of the theory. More importantly, and we are finally at the point where we can concentrate on the truly important issue, when we consider the significance of the technological infrastructure for space-based science, we can isolate some of the presuppositions we have not recently paid attention to regarding experimental practice.

If, as Joselyn and Whipple suggest, not only do the materials involved in the space stations and orbiters affect the experiments, but the environment of space itself also makes a difference, then we need to examine the kind of difference. With respect to the materials used to build space stations, etc., that perhaps may be merely a matter of fine-tuning. The real problem comes from the randomness of the influences of the environment of space. If we cannot anticipate with any degree of regularity the environmental influences, then what happens to the bedrock concepts behind the reliability of experimentation? One such notion leaps out: standard conditions. If the environmental influences of space are sufficiently random that we not only cannot build in safeguards against them, but also cannot calibrate our instruments to account for them, then what do our space-based experiments tell us? Another way to ask this question is: to what extent does an empiricist theory of knowledge presuppose standard conditions? The answer must be totally. “Standard conditions” is a fundamental epistemic criterion. If we reject or even reformulate the concept, then we have changed the meaning of knowledge. Furthermore, by emphasizing the degree to which what we know is dependent on our technology, we are left in the following paradoxical situation: the better we get at devising and constructing the means for learning new things about our world and universe, the less we know. Let us look briefly at these two points in some detail.

Despite his attack on Aristotle’s conceptual framework, Galileo still agreed with Aristotle’s definition of knowledge as certainty. If we know something, we are certain about it, not merely psychologically, but logically certain. That is what Galileo was talking about when he said that intensive knowledge gives us an understanding of necessity. This view lasted until David Hume destroyed it in his 1739 *Treatise of Human Nature*. Since Hume’s devastating attack on certainty, philosophers having been trying to come up with an account of knowledge which acknowledged that whatever we say we know must be bracketed by a certain probability. How to do that and still give an account of knowledge that is intuitively plausible is the big epistemological problem. The reason for the problem is the unabashed attachment we have to the idea that knowledge must be related to truth. However, if, as Kant did, we can recognize the fact that we cannot ever get to the truth about the world, as our data about the effect of the space environment shows us, then we may be able to make some headway on knowledge.

Kant showed that we could not ever know the way the world really is because our way of thinking about the world can never be tested against the world in a naked, unbiased way. We always interpret what we see and experience; it has to be that way. Our knowledge is, therefore, necessarily contaminated, just as our data from space are. The interesting feature of the Joselyn/Whipple article is their observation that, on the one hand, if we guard too heavily against the features that contaminate the data, we will not get anything worth using and, on the other hand, if we do not guard against these interferences, we can not use the data anyway. It would seem, therefore, that data cannot contribute to knowledge, a most unintuitive result. But this *is* the dilemma of knowing. We must trust our data, knowing that they are not trustworthy. Furthermore, we do not know how far to go in not trusting the data, since we cannot compare them against the world to know if we have made the right adjustments. Our continued use of instruments to provide data from space does not make this a new problem; it merely reveals the depth of the problem.

Perhaps one more look will help us make the case more convincing. Optical astronomy has come a long way since Galileo's little eight power hand-held telescope. We do not need to turn to the Hubble to see that. Not only have telescopes grown in size, but the necessary support systems have become more complicated. The truly large telescopes require massive housings, highly sophisticated background technologies to produce the machines and lenses, electricity to run the equipment and, once cameras are introduced, all the apparatus needed for quality night time photography and the optical theories to support interpretations of the products, computers to calculate position, manage the photography and coordinate the systems. But there is more yet; consider the contrast between Galileo's original hand-held telescope which we can take into the countryside, and a typical mountain top astronomical installation, with roads, electrical generators, sewage systems, housing, buildings to house the various types of telescopes and the computers and the other equipment. But even then there is more, for we need to consider the auxiliary support systems which the main system needs to carry out whatever theoretical investigations are in order. There is, for example, the entire support system that developed and produced the computers and the cameras and the programs and the space technology to launch telescopes, satellites, interstellar probes, etc. There is the optics of the camera, the new types of film... shall we stop here?

Astronomy is the science of the heavens. Its function is to describe the constitution of the universe in terms of the relative positions of its parts. To accomplish this goal astronomers need to be able to see the heavens. And so we have the elaborate technological infrastructure of the optical telescope. But to assume that the components of the universe are limited to those that can be seen by the human eye is absurdly homocentric. So if you add to the optical infrastructure and the radio telescopes and the theories upon which they are based, spectral telescopes, the use of high speed computers to not only control the telescopes, but to generate and interpret to at least the first and second order information they generate, the computers and the computer programs necessary for all that, the launching of space-based telescopes and the technological systems behind that, the infrastructure behind the computers, etc., the list goes on. If you add all that in, the technological infrastructure of astronomy appears to swamp the goal of the science. But there is more, for at each stage, the development of the instruments is constrained by the fit with other instruments and the theories with which they interact and sets of instruments and their backup systems. The result of employing these systems forces restructuring theories all the way down the line.

Just reflect on the original disaster with the Hubble and you will see the extent to which the systems of the technological infrastructure interact and affect one another. It is not just that new observations force revisions in the description of the heavens. The questions include how do you integrate spectral telescopy with optical? Do the theories behind the instruments cohere? One of the hot issues in cosmology today is the problem posed by the fact that the visual picture of the universe provided by astronomy does not cohere with the predicted mass of the universe. So now everyone is looking for dark matter. How do anomalous results from one instrument, e.g., excessive red shift, affect the other theories?

We look with awe at the picture that the new space probe, appropriately called Galileo, sent back to Earth on its way to Jupiter. If we think about the technological infrastructure behind the pictures, we get some sense of what is involved. The pictures are not simply sent from the space vehicle, traveling at high speeds in its own trajectory, to earth, also traveling at high speed and on its own trajectory, the "pictures" are transmitted as electronic code. That means they have to be disassembled, sent, reassembled, etc. The machinery, the programming and the capacity for mistakes is enormous. If you add the testing of scientific theories to the problem, and the interaction between the theories and the technological infrastructure, as well as among themselves, there can never again be a simple history of the ideas of science, nor should there be.

If the science is astronomy, or even cosmology, once we understand what it takes to do cosmology today we must turn to the technological infrastructure to understand its results. It is no longer possible to say, "Science tells us...," and it is certainly misleading to say, "Science and technology tell us...," for no one has taken the time to spell out what that means. When we do spell it out we will find what we really wanted to say was, "The technological infrastructure within which scientific theories are being developed and transformed makes it possible for us to describe and explain the universe in the following way." This contextualization of our science is extremely important. The kinds of things we come to know about the universe, or to put it more dramatically, the universe modern science reveals to us, is a function of this complex interaction between theory and technological infrastructure.

The second point noted above was that the better we get at building instruments and devising ways to use them, the less we know. We know less because we do not know how to filter the data. This too is not new. We never knew, a priori, how to interpret whatever data we got. In sum, we really do not know what we know. And we do not know what we know because the more data we get, the less we know what to do with it. Finally, thanks to the technological successes that makes space science possible, things are going to get worse, not better. Technology not only drives science, its continued development forces us to radically reconsider our conception of knowledge. For knowledge cannot be data nor can it simply rely on data. Technology may drive science, but it also may not contribute to knowledge.

This leaves us with the following final problem. Since the word "science" comes from the Latin "scientia," meaning knowledge, and if space science puts us in the position of knowing less and less, in what sense is it science? In short, we may be confused about the meaning of knowledge because we have identified it too closely with science.

6.8 Standard Conditions

The source of the problem here is the extent to which the efforts to make space science and its fancy new experiments yield new knowledge reveal the inadequacies of one of the major criteria constitutive of our current concept of knowledge. That criterion is standard conditions. The problem of ephemeral data is due to our inability to know which variables to account for, i.e., what are the standard conditions for space-based experiments? The problem is fundamental because those conditions will change depending on the experiment, since the knowledge the experiment is supposed to yield relates to earthly phenomena to be found in earth's environment, not to celestial phenomena in a celestial environment.

The very criteria for knowledge are under attack here. Our current sense of knowledge rests heavily on the notion of experience. We must be able to back up our claims by appeal to empirical data that count as evidence. These conditions for evidence are otherwise formulated as standard conditions. If they cannot be specified, then our evidence is in doubt and our knowledge shaky. This then is how current technological entrenched space science is transforming knowledge. It is forcing us to reconsider the notion of standard conditions as an epistemic criterion.

6.9 Conclusions

Finally, in closing, we are led to question the very value of such a large scale venture as the space program, especially if it leaves us with a totally bankrupt conception of knowledge. And while we are at it, we might as well note that the expense of the enterprise is itself suspect given large scale social needs we have at home. On the surface, three things seem obvious: (1) If the space program can not generate knowledge, then has it not lost its primary justification? (2) If we have been captured by our initial investments in it, which capital outlay now forces us to continue investing in a fiercely financially debilitating program that we can not give up because of the adverse economic effects it would have, then is not someone like Langdon Winner right about the autonomy of technology? (3) Should we not give up the space program and spend all our money on the poor and the homeless? Well, and I am sure this will come as no surprise, my answers are no, no, and no. Here is why.

First, the fact that our current criteria for knowledge are inadequate does not mean that the space program has no justification. The fact that we are having a difficult time figuring out what we have actually come to know from these new ventures does not mean the ventures are flawed; it means our account of knowledge is. But this is no big deal. Our account of knowledge is constantly being revised, as I have tried to show, in the light of new conceptual and technological developments and factual information. The fact that on the Aristotelian scheme mathematics could not yield knowledge did not mean that knowledge could never be had. We had to change our criteria and our conception of knowledge. The fact that on the Aristotelian/Ptolemaic account of the universe there could be no other place than earth where objects revolved around a planet did not mean that the moons of Jupiter did not exist. By way of analogy, the discovery that they did meant that we needed a new way to explain and accommodate that set of phenomena.

As I have suggested, given the epistemological problems presented by the space program, we need to give up simple-minded empiricism as a foundational criterion of knowledge and reevaluate. It looks like the criteria for knowledge are going to have to be extended to include social structures and institutional criteria as well as some commitment to large scale coherentism. The fact that the universe has forced us to acknowledge that we do not have the intellectual equipment to understand it yet does not mean we should tuck our satellites and space stations and interplanetary probes between our legs, so to speak, and run for home. Our hubris has once again been exposed; so much the worse for hubris. There is clearly philosophical work to be done, i.e., epistemological conceptual work.

Now let us turn to the second point above. Given the scale of the investment, it is clear that we cannot abandon the space program. There are two reasons for this claim: (a) we have already invested so much, and (b) it would be too disruptive to the economy. Therefore, Winner is right, we are slaves of a run-away technological project that has stripped us of our political will and disenfranchised us. Well, I am afraid not. We can stop the program or features of it at any time. The original design of the International Space Station was challenged and downscaled. Just as with other large scale technological ventures such as the super colliding super conductor, the fact that we have invested so much already is no longer an acceptable justification for continuing to invest in that project. This is only rational. If we fail to learn from our mistakes we are irrational. This is Pitt's Commonsense Principle of Rationality, CPR. If we lacked both the will and the power to learn from our mistakes then perhaps Winner would be right. But we do learn and we do act. Winner's claims have been outstripped by real world events. The shuttle program has come to an end, the plan to return to the moon has been dropped. The space program has not become autonomous, because we can interrupt it and transform it no matter what the consequences are as long as we have the will. Thank you Bill Clinton for proving a philosophical point although the cost is fairly high.

Finally, given the problems back here on earth, should we not give it up and spend the money on more important problems? Again, I am forced to disagree. Let me use yet another analogy to underline my reasoning. The late middle ages saw the construction of the great cathedrals of Europe. The building of these magnificent edifices employed hundreds of workers, forced architectural designs and building skills to unimagined heights (bad pun), and brought our level of artistic and physical accomplishment to unprecedented levels. At the same time, this was a period in which human knowledge was having to be reconstructed, the cost of these buildings impoverished societies just crawling out of a period of excessive poverty, banditry, and disease. This was the period of the black plague, of petty fiefdoms waging petty wars with devastating results on the innocent. And yet, to what I assume was supposed to be the greater glory of god, these magnificent buildings were raised amid poverty and disease and rampant human bickering. The cynic could say that this is yet one more example of the corruption of organized religion. It certainly seems so. But to limit the issue there would be to shortchange ourselves. For there is more to the tale.

The ability of the human spirit to rise above its miserable, debilitating, squalid environment, inspired by the search for something transcendent, however misguided, is noble and fundamentally humanizing and enabling. Even if the motivation of those who initiated these projects and those who followed by envy was less than admirable, what they produced was admirable. And so, I would argue, likewise for the space program. This is not to deny that there is misery at home. This is not to deny that more needs to be done to alleviate that misery. But do not deny as well the need to search for more than we can see and deal with here. The space program represents the modern version of the building of the great cathedrals. It is a venture full of all that marks it as a great human undertaking. It has its dark side, its darkly human side full of politics and greed and avarice and the follies of power. But it also gives us what cannot be produced without this kind of mobilization of human resources. It challenges our abilities to accomplish feats of incredible physical achievement like walking on the moon and establishing a permanent human home in space. It does more, however. For in showing us how inadequate our concept of knowledge and value are in the face of change, it forces us to rethink who we are and what we are and where we are and who, what, and where we will be. To meet those challenges we need to be intellectually flexible and ready and capable of being challenged. This requires that we acknowledge that what we now use as criteria for knowing may not be adequate in the face of new data made possible by new technologies. It means that what we now value may not be valuable in the face of new discoveries. It requires a philosophical approach that gives each its due, but none undue privilege.

So, my final question. How do we know if we are in fact employing such a philosophical methodology or if we have fallen back into corrupt ways and have become again pathological? I think I have an answer to that. Is your characterization of a given situation able to handle a variety of different types of philosophical questions asked of it? If you have a problem, is the source of the problem also open to being queried for different kinds of problems? I tried to show how this was to done by showing how the space program presents a challenge to our conception of knowledge and also allows for us to ask questions about its moral standing and its political character. In short, have you characterized a situation in a way such as to close it off from the interests and concerns of other philosophers? If you have, then they have a right to accuse you of being insular and isolated. But if you can present your problem in such a way as to allow other problems to be raised without changing the problematic, then you have succeeded in doing philosophy of technology in a way Sellars would approve.

References

Galileo, G. 1632, 1967. *Dialogue on Two Chief World Systems*. Trans. S. Drake. Berkeley: University of California Press.

Sellars, W. 1963. "Philosophy and the Scientific Image of Man". In *Science, Perception and Reality*. London: Routledge and Kegan Paul.

Wallace, W. 1984. *Galileo and His Sources*. Princeton, NJ: Princeton University Press.

Chapter 7
Working the Natural/Artificial Distinction

The distinction between the natural and the artificial lies at the core of the philosophy of technology. Much of what we say in this field turns on one's position with respect to that distinction. For example, if one endorses the distinction then one is in a position to argue against certain technological innovations such as gene splicing on the grounds that they are unnatural. In other words, the distinction provides certain criteria, depending on how we define the key terms. In what follows I will look at one form of what is usually deemed artificial – human technologies – and examine the extent to which it is appropriate to call a technology such as a space station or a ball point pen "artificial". Focusing on "technology" forces the issue of defining "technology" – and I do offer a definition. The point of attack is to push the question "What work does this distinction really do, if we cannot satisfactorily define the terms in question?" In the end I will show that by attempting to maintain the contrived distinction between the artificial and the natural we commit ourselves to living in the past, for it is a distinction which no longer applies to the modern world.

One of the more unforgettable images in Stanley Kubrick's film 2001 occurs near the beginning when the ape/human tosses the long bone with which he just bludgeoned a competitor into the air and it morphs into the almost completed space station orbiting the earth to the background music of a Strauss waltz. The clash of the primitive with the ultra modern highlights the progress humankind has made through its technologies. But there is another clash implicit in the contrast between the primitive and the modern that speaks as much to the artificial/natural distinction as the contrast between the bone as tool and the space station. That is the contrast between the precursor human and the modern human. It is clear that the space station was a constructed thing, designed, if you will, the result of conscious decisions made over a period of time. What about us? I propose that we too are the result of conscious decisions made over millennia.

Appeared in *Philosophical Dreams*. Oneonta Studies in Philosophy, edited by Douglas Schrader. SUNY/Oneonta Press, 2003. Reprinted by permission of the editor.

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_7,

If we reflect on the selection criteria for marriages over the ages, you may see one aspect of this issue. We know that, by and large, marrying or mating as a function of romantic love is largely a modern phenomenon. This is not to say that people didn't fall in love before today – one need only think of the love poems of Sappho, Catullus, or even Homer's tales, to know that love has been around for some time. However, that is not to say that love was always or even most often the basis for the creation of lasting unions. For most of recorded time, arranged marriages were, and even in some parts of the world remain, the norm. The criteria for the arrangement varied from securing political alliances to acquiring wealth, even to idealizing physical looks. The precise criterion is not the point. The point is that the current configuration of human society is the result of decisions made on the basis of these criteria. There is a certain sense then in which Humanity has been its own best technology, making and remaking itself. I will return to this idea below.

But first the promised definition. In *Thinking About Technology* (New York: Seven Bridges Press, 2000) I argued for a number of points. Two of them are relevant to this discussion. The first is the idea that when attempting to understand a technology, how it became what it is, etc., one must pay attention to the decisions that were made by the people behind and in front of the project. The second point is a definition of technology, a definition that is consonant with this view: technology is humanity at work. I use this broad account for a number of reasons. (1) It captures the fact that our intuitive concept of technology is not, in fact, limited to tools or mechanical machines. With a little reflection we acknowledge that social systems and institutions are also technologies, for example, consider governments, legal systems, political parties, funding agencies, and universities. (2) The idea that it is humanity at work that constitutes technology focuses us on the fact that in our work we transform what is available to us into something else, something someone can use. Finally, (3) because the scope of this definition is so large, it has been suggested that it is useless. I do not think it is useless, but I do think that it forces us to stop talking about TECHNOLOGY, as if it were some one thing, which it is not, and instead makes us direct our attention to particular technologies and to the decisions that led to their creation and their further use.

With this thumbnail sketch of some complicated notions, we can return to the idea above that Humanity is its own best technology. It is no secret that Humanity has been transforming its world in order to achieve a constantly changing and hence, ever elusive, vision of the good life. As we changed the world, we effectively changed our opportunities for future action. We changed, that is, the possibilities for our own transformation. This is not to say that the results were always good, or even as expected. If there is a law of technological growth, it is the Law of Unintended Consequences (LUC). LUC states that things don't always go the way you hoped – which in turn is an acknowledgment of the roles of contingency in our lives. Consider some of what we know. One of our oldest technologies is agriculture. Since its invention, humans have been experimenting with grains in order to obtain a better yield, greater resistance to pests, longer shelve life, greater nutritional value, etc. As we have found out about what it takes to do this and as we have accomplished these goals, we have also seen results in terms of increased size and robustness of our children. But as LUC plays out, we also find that with increased size there are also liabilities, such as increased back problems. The contingency here is the genetics of the human form and how it responds to changes in nutritional value. The point maybe lost here for this appears to be a silly example. But let us look more deeply.

In an important sense, the obvious candidate for an example to make this point stares us in the face: The Human Genome Project. The sequencing of the human Genome is not some new outrageous technology, it is merely the next step in a process humans have been employing all along, the process of creating better humans. The fact that we do not know what a better human should look like is not the issue, the drive to create one appears to be a genetic imperative. We are our own creation, not so much by design, but by particular decisions based on sometimes marginalized criteria. That children are their parents' creation will be all the more apparent when parents are able to select what color eyes their children will have, or even their sex. At the moment we are what we are by virtue of whatever contingencies compelled our parents to mate. In the future we will be what we will be by virtue of the compromises future parents will make when staring at a gene chart as they design their child.

So, what does this have to do with the distinction between the natural and the artificial? I am sure you can see it coming. . .if what humans are and do is constantly changing, where are we to draw the line between what is natural and what is not? The problem stems from the fact that we do not have a clear sense of the meaning of either "natural" or "artificial". What, for example, does it mean for something to be natural? Consider the following: we know our planet is constantly being bombarded by meteorites of various sizes, a process we consider to be natural, knowing what we do about cosmology. Suppose we get hit by a really big one that causes all sorts of disruption, effectively wiping out civilization as we know it. Is it a natural phenomenon? Perhaps. However, if we discover that this particular meteorite was selected and aimed by a powerful alien race and then deliberately propelled at us, we would probably say that it is not a natural occurrence. On the other hand, if the race that brought about this catastrophe was by its nature dedicated to conquest and destruction, then there is a sense in which it is a natural act. It is not a natural act by way of the laws of physics, whatever they may be, but it is a natural act by way of the laws governing the development of that race. If that seems a bit farfetched, let's come closer to home.

Indeed, let's cut to the chase. If Humanity is a constantly self-creating natural entity, then isn't anything Humanity creates equally natural? To pose the question this way immediately prompts a dilemma. On the one hand, if Humanity is constantly deliberately changing itself then why is something else that Humanity changes not natural? On the other hand, if being created by deliberate action rather than the laws of nature is the mark of the artificial, then why is Humanity itself not artificial? This is a nasty dilemma, either everything is natural, or Humanity and everything we create is artificial. Let us look at each horn separately.

What is wrong with everything being natural? Well, for one thing, it doesn't fit our intuitions. This is not to say that our intuitions should have the final say on everything philosophical. However, our intuitions are a good place to start. It may turn out that they need refining, but if so, that needs to be shown. Here is an example of why it just won't do to say that everything is natural. Let us say that we have become a full-fledged space-faring race and one of our ships is approaching a previously uncharted planetary system. Using our gross detection instruments, we discover that of the five planets, two have satellites. One of these moons is composed of various materials, iron, minerals, etc. and it is of irregular form. Rotating around the second planet is the other satellite – it looks very much like the space station from 2001. It too is composed of various materials, but it is also of regular form and it is sending out what appears to be some sort of signal. There is clearly a distinction between these two types of satellites – but what is it? Should we, for example, say that the moons are natural and the space-station-looking-thing is artificial? What assistance would we get by claiming that the space-station-looking-thing is not a natural creation but artificial?

To begin with, we need to address a possible ambiguity. When we say it is artificial, do we mean it is not a real space station or not a real satellite? For surely by definition it is a satellite. And whatever would be gained by claiming that it isn't a real space station, it just looks like one? So that is not the sense of "artificial" we mean. No, by "artificial" in this context we mean that its existence appears to be evidence for the presence in some form or other, at some time or other, of entities capable of building such a place. Now, not everything built by an entity is something we would call artificial – or is it? Is a bird's nest natural or artificial? What about a dam built by a beaver? If the building of a nest is a natural act for birds, is not the building of some form of shelter a natural act for humans, and hence, a lean-to is a natural thing, so is a house made of sod, or of wood, or of brick, or an apartment house made of steel, concrete and glass? This is a slippery slope. Moreover, taking this direction doesn't seem promising since it seems to assume what we seek to clarify.

So, to try a different tact, let us ask what work the distinction between the natural and the artificial is supposed to do. What is the point of drawing such a distinction? For one thing, it assists us in glorifying Humanity. It does this by allowing us to claim that only humans can produce that which is not natural. In this way we continue to endorse the great chain of being, which places Humanity at the top of creation, which does nothing more than endorse a theology, and that hardly settles anything.

There is a second motivation for pursuing the distinction between the natural and the artificial that depends on how you define "artificial". Earlier it was claimed that our discovery of a space-station-looking-thing in orbit around a planet in our newly discovered planetary system could be considered evidence for the existence or for the presence in some form or other at some time or other of entities capable of building such a place. Intuition says further that they had to do more than build it, they had to first, or concurrently with building it, design it. That is, an object which gives the impression of being deliberately built also suggests that it was first designed, for the deliberate building of a thing proceeds according to a plan. This is by way of contrast with a bird building a nest or a beaver building a dam, which seems to be the result of natural selection in some form or other. Given these thoughts, what if we define an artificial thing as one that is designed? Does that help? No. For one thing, it is possible to imagine building something without a prior design. In this case you would simply start banging things together and then see what happens. You may get something useful, and you may not. We have all seen cases of random fiddling which results in something that eventually gets put to work- for something to be useful it doesn't have to have been designed; thus, a rock is used to prop open a door.

If you start at the other end, the distinction still won't work. For instance, here we have this space-station-looking-thing. It is big enough and intricate enough to suggest that it was constructed by intelligent beings. However, the existence of something, however intricate, does not entail something designed. This is basically Hume's response to the argument by design. As Hume tells us, some argue for the existence of God by pointing to the intricate design of the universe and point to how everything seems to fit together just like a marvelously crafted pocket watch. Just as a watch must have a watchmaker, so too the universe must have a designer. Hume responds by noting that the appearance of interconnectedness does not entail design by purpose. Hume is right.

But we shouldn't give up on the connection between the artificial and design so easily. One hallmark of a technology is that it is designed. This is particularly the case if we look at large engineering undertakings. It is also the case when we consider something as commonplace as a ballpoint pen. The fact that the production of artifacts that do not occur in nature all by themselves is the result of a process of design leading to manufacture tells us we should look again at the concept of the artificial. This is sufficiently elementary thinking about this that we should probably not even consider the situation, already in place to some degree in many manufacturing activities, of robots building objects. For if the objects built by humans are artificial in some sense, and if humans build robots and robots build automobiles, then is a robot built automobile a post-artificial object? Is this the world "post-modern" refers to?

One way to reopen the discussion of the artificial is to return to our dilemma. The first horn of the dilemma claims that since humans have transformed themselves and continue to do so with, to some, frightening future possibilities, then everything associated with humans is artificial. So what? Well, to begin with, it might appear that this would affect our conception of what it is to be human. This might not be a bad thing at all. It is no secret that we are confused about who and what we are. There are at least two distinct schools of thought in what has been referred to as the nature/nurture debate, of which our question here is merely one version. If we take the nature side, we are fundamentally biological creatures, i.e., we are what our genes make us, we do what we do because our genetic makeup requires that we respond to stimuli in specific ways. On the nurture side, we do what we do because we have been taught to respond to situations with appropriate behaviors. A third position lies somewhere in between. We are what we are partially because of our genetic makeup and partially because of how we were raised and what kinds of experiences we have had.

Irrespective of how the nature/nurture debate plays out, it is irrelevant to the issue of whether or not, by virtue of our ability to continue transforming ourselves, we are in some sense artificial creatures. If, on the one hand, we transform ourselves by selecting options to become this or that or by way of making crucial choices for our children, because we are genetically programmed to proceed in this fashion, it still doesn't change the fact that we do transform ourselves. On the other side of the coin, if we undertake these changes and transformations because of some rational process, we still effect the changes, and we still are not the kind of creature we were when we started, whatever the cause. In short, it doesn't solve the artificiality question.

Another try: let's look at what seems to be the problematic core of the issue – the fact that we as a species have made choices which have lead to our living in ways that are different from the way other sentient creatures live. Other species are social, albeit by virtue of our classifying them so. Other species build habitats, nests, use caves, etc. Other species communicate, make decisions, create paths. Yes, but not on the scale we do. So is artificiality merely a matter of scale? What do we say about some of the African termite mounds that loom so large. Are they artificial? Well, they certainly are created. But surely we don't want to say that anything created is artificial. Bird nests are created – but intuition here says they are not artificial. And we find ourselves back at the slippery slope we were at earlier.

Once again we are forced to ask what work the distinction between the natural and the artificial does for us. So far we have seen that it does not allow us to identify an object as natural or made simply by virtue of its appearance. The distinction does not find support in the appeal to design, not yet anyway. It does not inform, nor is it informed by, the nature/nurture distinction. It doesn't appear to help us understand technology any better, yet. And it doesn't shed any light on the problem of understanding who we are, yet.

And yet, given that it is a distinction with a long history, surely there was some work it was supposed to do? If we look to the origins of the distinction we find some help. The Greek sophists first introduced the distinction between the natural and the man made to differentiate that which we had some control over from that which we did not. (for discussion see G.B. Kerford, The Sophistic Movement (Cambridge, 1987): Ch. 10 "The Nomos-Phusis Controversy", pp. 111–130) In this way, when they turned to politics, they had the basis for articulating a progressive view of government, arguing that human laws, unlike the laws of nature, are changeable. Hence, if we don't like this form of government it can be changed.[1]

Although not all ideas introduced by the Greeks have survived in their original conception, this one has had surprising durability. Today we reason as follows: Nature and the natural have their own rules and they are beyond our control, e.g., rivers flood after heavy rains. Further, whenever we try to change nature we have to contend with LUC. If we build dikes to constrain rivers and then settle on the old flood plains and rivers still flood, we lose our investments. There are, however, sets

[1]Thanks to my former colleague Mark Gifford for help on this.

of decisions, objects, policies over which we do have control and they are the very things we have created. On this version, if it is artificial it is controllable. While this account of the distinction seems promising, it leads right to the issue of the autonomy of technology.

The autonomy of technology question is this: does technology have a life of its own? Have we reached the point where the momentum behind various technological enterprises is so strong that they are beyond control? If so, that undermines the idea that the artificial is that which we can control. This is the worry behind Luddite claims that technology is out of control and that it is taking over our lives. If this is true, then not only does the natural/artificial distinction cease to do any work, but since technology is not under our control, and that which is not under our control is natural, it would appear that technology has become part of the natural, which is clearly counter-intuitive.

Or is it? In a perverse way it has brought us back to the view that everything is natural. Not only that newly discovered moon, but the space-station-looking-thing as well, and ballpoint pens and glass, steel and concrete apartment buildings. Now despite the fact that my sympathies lie in this direction, I cannot accept this conclusion because it is based on a false assumption. That assumption is that technology is autonomous. However, we can attack the assumption of autonomous technology from two different angles.

First, we can attack the idea that technology is autonomous by appealing to our definition of technology. If technology is humanity at work, and if this definition forces us, as I think it does, to look at specific instances of humanity at work, we will see that technology cannot be autonomous. For one thing, no one group works in isolation from another. There may be the rare hermit who works alone, but ultimately for the most part, whatever work we do must cash out in a social domain of one sort or another. So, for the sense of autonomy that would have us think of total independence, it fails. If we look at the sense of autonomy that speaks to the unstoppable momentum of technology we have a different kind of problem. Since our definition focuses on humanity at work, and short of total annihilation, it seems impossible, undesirable, and just plain dumb to stop humanity from doing whatever work it does (unless it is self-destructive) then, yes, technology is autonomous. But if we look again at the claim that because of the breadth of our definition we must look to specifics, then no, technology is not autonomous. It fails to be autonomous because there is no single human undertaking that cannot be stopped. In this sense, nothing we have undertaken fails to be under our control. Dams we have built can be taken down. Laws we have created can be retracted or changed. Wars can be stopped.

But, let us note that some may not want to buy into this definition of technology as humanity at work. Does that mean we are left with the counter-intuitive conclusion that since autonomous technology is uncontrollable, it must be natural? No, for we can still undertake an examination of the concept of autonomous technology free of any definitional attachments and show that the concept is bankrupt. We begin by asking for an example of autonomous technology, by this meaning: show me a case of a technology so isolated from human constraints that it has a life of its own. One example favored by Langdon Winner is the electrical power industry in the United States. Winner gives us a highly politically charged account of that industry, indicting it with claims of monopolistic and undemocratic disenfranchisement of ordinary people. He further notes the ubiquity of electricity and how it has so completed infiltrated our lives that when the power is disrupted, we virtually stand on the edge of extinction.

This is ominous, but a second look brings hope. Yes, the evidence of the ubiquity of the power industry is everywhere. Frankly, this is a fact I like. When Winner visited my university several years ago, he urged that we turn off our electrical furnaces and heat our homes with wood we gather ourselves. Speaking as someone who can in fact heat his home using a wood stove, it was clear to me that Winner had never done so. To heat your own home exclusively with wood, you need to invest in a chain saw, a truck, a wood lot or have access to a forest, some splitting mauls, or a powered wood splitter, and lots and lots of time. Not only do you have to locate the wood, chop down the trees, cut them up, dispose of debris, haul the wood back to the house, stack it and let it season, you also have to maintain the fire. And because it is not a good idea to try to burn green wood, you need to plan ahead by at least 6 months. Now, frankly, I enjoy a good fire and the house just feels warmer when the wood stove is going. But, I also must admit that I enjoy being able to turn on the electric furnace and sit down and read a good book, or listen to the stereo or even write some philosophy. So the bottom line here is that even if the electric power industry is in some sense autonomous, it is not clear that it is a bad thing. But, is it autonomous? The answer is clearly “no”. For as indicated above, it is perfectly clear that it is possible to do as Winner wishes and to turn off the electric furnace. With properly planning, one could also install a solar array and generate one’s own electricity and become totally self-sufficient electrically speaking.

But, you might object, could all of civilization disengage in that way and still function? But is that the right question? The claim is that for technology to be autonomous, it must be uncontrollable. Disengaging is one form of control. Furthermore, there is evidence that the electric power industry itself can be controlled. First, in my home state of Virginia the industry is subject to government control. Rates are regulated as are the location of power lines. Proposals by power companies to place power lines can be and have been successfully challenged. Furthermore, changes, such as deregulation, can be forced on the industry, as they are currently occurring in the US. So, big as the power industry is, it is not beyond control. And we haven’t even addressed the issue of whether or not it makes sense to speak of the power industry as one unified body. It is, in fact, a collection of competing companies of varying sizes with their own concerns and agendas.

That leads us back to the distinction between the natural and the artificial and the issue of control. We were examining the Greek idea that artificiality implies controllability. In that context we have been exploring the possibility that some things created by human being can escape from our control in such a way as to evade control, giving rise to the idea that technology could be autonomous. There should be now at least good reason to believe that this does not make good sense. So that returns us to the question of whether or not the distinction between the artificial and the nature does any work. Again, the conclusion appears to be negative. Let us look at the notion of control again.

I have been arguing that even large-scale technologies can be controlled, challenged, changed, and, if you look at the nuclear electrical generating story in the US – even rejected. This suggests that perhaps the natural/artificial distinction can in fact do some work. But, alas, this is not so. For once again we find LUC looming before us. Let us revel in our control over technology for a while. Consider the following scenario. The major power company supplying our region, wished to run a very large, 765 kV, power line across some lovely farmland. Well, despite the frailty of the company's arguments about how they need the line, the state government accepted their argument. The good news is that it also accepted the arguments of those opposed to the power line and agreed to an alternate placement of the new line, thus avoiding the farms. But the new site is in one of the most beautiful parts of the state. It attracts visitors from all over who come to admire the natural (!) beauty of the region. Now that view will be permanently spoiled by this massive, ugly, artifact. A view that once brought peace of mind and pleasure will be gone forever. Win some, lose some. So we have some control, but not total control. On the other hand, who has total control over anything?

For it is not just technology that poses the problem of the Law of Unintended Consequences coupled with the concept of control. It is safe to say that much of the activity in the West associated with the growth of industrialized countries has involved the control of nature. We can control nature, to the same extent as we can control anything else. This will not sound reasonable at first, but think about the following parallel. There are some things in nature that cannot be controlled, yet. Take, for example, the eruption of a volcano, or a hurricane.

On the other hand, we can divert rivers, or even dam them. This may have the effect, as it has in the Pacific Northwest of America, of threatening the existence of the salmon, shutting off its route to ancestral spawning grounds, an unintended consequence. We can also seed hurricanes and we can, if not control them yet, at least forecast them. We have eliminated certain diseases – which is controlling a part of nature with a vengeance. We can cut down forests and pave over fields, create cities where there were plains. We can do a lot to nature. But we don't have complete control, yet. There is nothing in principle that says we won't be able to control the forces of nature. The more we learn about how hurricanes are formed, the closer we come to devising a method for defusing them, likewise for volcanoes. And the upshot of this is to turn nature into the artificial, since it can be controlled, not all of it, yet, not all the time, yet, but soon. But if this is not an outright contradiction, it is surely a category mistake. Hence the distinction between the artificial and the natural fails again.

As human beings learn more about how to control events around them, which requires more factual knowledge, and the development of the means to use that knowledge, it appears the natural/artificial distinction has less and less work to do. It may have had a role to play in our infancy, when it appeared we were subject to the whims of nature. That is less the case now, and as our powers grown it will soon be totally irrelevant.

But that is just the point of characterizing human beings as self-transforming creatures and of defining technology as humanity at work. As we learn to control nature and as we get better at predicting the consequences of our actions, we will increase the potential for further transformation of humans. Many years ago, when successful gene-splicing first became a reality, *Time Magazine* ran an issue full of science fiction speculation about how we would soon be genetically engineering people to meet specific needs. One such speculation stands out. Let us briefly project into Time's future. We are now a space-faring race. Space travel will be expensive, it takes a lot of resources to build a vessel that can safely transport human beings through the hazards of the void. Room inside such vehicles will be at a premium, hence, we need to maximize our use of that area. In that light, consider the possibility that certain anatomical features of human beings may not be necessary in space. In zero-gravity, our legs may just be in the way, as they often are in airplanes and cramped busses. Maybe the ideal crewmember on a space ship would be legless. In many science fiction stories being written today, the idea is being floated that a human brain could be hard-wired into the administrative center of a space ship and serve as its pilot.

Add to the science fiction the development today of "artificial" organs and prosthetic devices. The bionic man of the television program of decades ago is much closer to reality today that we ever would have thought. Add to that the knowledge we are gaining from the Human Genome Project and project its influence on the future of the race. All of this suggests that we should finally give up that worn out old distinction between the natural and the artificial.

Sophocles, in his response to the Sophists' formulation of the distinction had it right from the beginning in his Antigone:

> Wonders are many on earth, and the greatest of these Is Man, who rides the ocean and takes his way Through the deeps, through wind-swept valleys of perilous seas That surge and sway. He is master of ageless Earth, to his own will bending The immortal mother of gods by the sweat oh his brow, As year succeeds to year, with toil unending Of mule and plough. He is lord of all things living; birds of the air, Beasts of the field, all creatures of sea and land He taketh, cunning to capture and ensnare With sleight of hand. . . . .(and later) There is nothing beyond his power. His subtlety Meeteth all chance, all danger conquereth, For every ill he hath found its remedy, save only death.

But lest we think this is all praise, Sophocles concludes with a warning:

> But he that, too rashly, daring, walks in sin in solitary pride to his life's end. At door of mine shall never enter in to call me friend. (Antigone, in The Theban Plays translated by E.F. Watling, Baltimore: Penguin Books, 1947, pp. 135–136.)

Sophocles may have been ahead of his time, but he also seems to have been on the right track. Today, as human beings take control of everything in sight, sometimes not so wisely, but certainly with enthusiasm and vigor, it is no longer a question of what is natural and what is not – but rather what is possible and what will be possible.

This then leads us to yet another distinction. For when we ask what is possible and what will be possible, we raise the specter of what is desirable and what will be desirable, reminding us of Sophocles' concluding warning. In other words, once we know what we can do, we need to address the question of what we ought to do.

References

Antigone. 1947. *The Theban Plays*, pp. 135–136. Trans. E.F. Watling. Baltimore, MD: Penguin Books.

Kerford, G.B. 1987. "The Nomos-Phusis Controversy (Ch. 10)". In *The Sophistic Movement*, pp. 111–130. Cambridge: Cambridge University Press.

Chapter 8
Discovery, Telescopes, and Progress

8.1 Introduction

In this chapter I attempt a number of things. Together they constitute some steps toward the development of a new research program. It is a proposal for a new way of conceptualizing the relation between science and technology. It is also an attempt to find a way to escape some old philosophical dichotomies; dichotomies which have kept philosophers of science and philosophers of technology apart. The underlying theme is this: following Derek Price, it seems clear that progress in science is a direct function of increasing sophistication not merely in instrumentation, but in the technological infrastructure which underlies and makes mature science possible.

Price claimed that, “historically, the arrow of causality is largely from the technology to the science” (Price 1963), but it is only part of the story. By emphasizing the causal priority of technology in scientific progress, Price was attempting to overcome a popular characterization of the relation between science and technology in which technology is placed in a second class position, the offshoot of science or sometimes its “handmaiden.” Price was on the right track, pointing out that despite the fact that historians and philosophers of science have a tendency to talk about progress in science in terms of the history of ideas, a significant role is played by technology, a role largely ignored by these same philosophers and historians of science.

The typical history of ideas story of science proceeds by relating that, for example, Newton’s mechanics replaced Aristotle’s and then relativity replaced Newtonian mechanics, The story is usually told in Kuhnian fashion, without any mention of the means by which anomalies were discovered.[1] It is merely announced that following

---

Appeared in *New Directions in the Philosophy of Technology*. Edited by Joseph C. Pitt. Philosophy and Technology Series, Vol. 11. Dordrecht: Kluwer, 1995. Reprinted by permission of Springer Academic Publishers.

[1] A classic example of this can be found in a recent (1987) issue of Mosaic, an official NSF publication: Every so often. in the long course of scientific progress. a new set of ideas appears, illuminating and redefining what has gone before like flare bursting over a darkened landscape. It happened when Galileo realized that physical laws needed to be written with numbers and

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_8,

a certain experiment, it was decided that so-and-so's theory was false and was replaced by another. Thus, a typical bad history would tell you that Michelson and Morley's experiment was developed to test for aether drift, as predicted by Newton's theory. Once it was discovered that drift did not occur, Newton had to be abandoned. Enter Einstein, and all is saved. Very few histories reveal that Newton did not talk about aether drift; the notion evolved over a 100 years in the course of his successors' efforts to adjust his theory in light of their experience with it. Likewise, very few accounts tell you about the details of the Michelson-Morley experiment.[2] The point here is that on the history of ideas account the history of progress in science is made to read like merely the replacement of one bad theory by another once the bad theory is discovered to be faulty. What is ignored in all of this is the technological infrastructure within which the falsification and/or confirmation of theories takes place, to the extent that theories are falsified and/or confirmed.

More to the point, few scholars talk about the epistemology of experimentation. or the nature of the link between experiments and the theories they are supposed to test,[3] or the impact of experiment design and the availability of materials, techniques, and instruments. This aspect of the story of the progress of science/technology is important particularly at times of dramatic changes, such as are marked by the replacements of one major theory by another, because it is precisely at this juncture that what counts as evidence and how it comes to count as evidence is often at issue.[4]

In this chapter, I turn explicitly to the role of the technological infrastructure of science in the growth of knowledge in general. I start by exploring some features of the manner in which Galileo's development and use of the telescope helped create an initial technological infrastructure for astronomy and then move to a sketchy reconsideration of that notion as it occurs in modern guise. In so doing I hope to make plain what is meant by a technological infrastructure of science. Instead of attempting to argue one side or another of the old science/technology debate, some of the issues are recast so as to demonstrate the epistemological importance of a technological infrastructure construed as interrelated sets of artifacts and structures. Furthermore, I want to suggest, if not argue here, that just as it makes no sense to talk broadly of technology, it makes no sense to speak of the history and development or importance of a single artifact, suggesting, as this does, that once invented artifacts

---

invented the scientific method, when Darwin found an entirely different way to consider the evolution of living things, when Freud placed consciousness and emotion in a new context, when Einstein found a radical way to look at space and time, and when Wegener launched an earth science based on continental drift.

[2]Or that both experimenters were Americans and that the experiment was carried out in Cleveland, Ohio at what was then the Case Institute of Technology. After all, with names like "Michelson" and "Morley" they just had to be British and the experiment must have taken place at the Cavendish: didn't they all?

[3]This situation is changing. See Franklin (1986), Ackermann (1985), Hacking (1983), Galison (1987) and Cartwright (1989).

[4]I have discussed some or these issues elsewhere. most recently III Illy (1991).

remain stable over time. The thesis is direct: *the development of new information in a mature science is, by and large, a function of its technological infrastructure*. In short, scientific discovery today almost completely depends on the technological context without which modern science would be impossible. I will not raise the question of the merits of this situation until the end of my discussion, although I will provide a hint: in this age of increasingly theoretical science, the technology behind the science may be our only contact with reality, and even so it is at best a tenuous one. But now let us turn to the question of technology and discovery.

8.2 Discovery

There has been little discussion of discovery by philosophers. To the extent that the issue has been raised it comes in three contexts. First, there are the problems the concept of discovery creates or sets for cognitive science. I will not discuss that at all. Second, discovery is a problem for realism – where the debate hovers over the distinction between discovery and invention. Consider the question, for example, whether it is possible to discover something that does not materially exist, that is, an idea or a theory. The problem of figuring out what this means rapidly becomes tangled, despite the fact that things seem fairly easy at first. One cannot invent, for example, the Americas – they are already there, so we discover what is there. But scientific theories are invented, not discovered, and yet they are supposed to be about what is there. To say we discover a theory makes it sound like the theory has been lying around waiting for us. But that is too Platonic for my tastes, especially since we keep "discovering" the wrong theories–i.e., false ones. On the other hand, we need to avoid making it sound like we invent theories out of thin air-surely scientific theorizing has some relation to what is there. Thus, there is a certain tension surrounding discussions of discovery in accounts of the development of scientific theories, which tension is generally resolved by invoking a temporal ploy-we begin by inventing ways of speaking about situations which have avoided our efforts to understand them until we have some sort of acceptable proof to the effect that what we invented to explain the situation really is there. At that point we say we have discovered these new phenomena, like gravity or quarks. But this is not really a solution; it is more like a wiggle.

The third situation, in which discovery has been a topic for philosophers of science, is as unhappy as the second. It is to be found in the context of Reichenbach's distinction between the context of discovery and the context of justification – a distinction employed so well by Popper in the *The Logic of Scientific Discovery*. Popper made things very difficult with his classic dismissal of discovery as an issue for philosophers by characterizing it as a fit topic only for psychology. His own view is frustratingly obscured through the mistranslation of the original German title, *Logik der Forschung* as *The Logic of Scientific Discovery*, when Popper rejects the very concept of a logic of discovery in the first 5 pages. Surely we would all have been served better if the title of Popper's book had been more accurately translated as *The Logic of Scientific Research* – for it was the structure of *that* process, “Forschung” in German meaning “research” or “investigation,” with which Popper was really concerned. But, the follies of mistranslation to one side, it is nevertheless true that, for the most part, philosophers of science in the middle years of the 20th century bought into the Reichenbach/Popper view that discovery is not susceptible to logical analysis and, hence, is not an appropriate topic for discussion. It was only some years later, following the publication of Kuhn’s *Structure of Scientific Revolutions*, when the locus of philosophical attention shifted to the historical process of science and away from concerns over its rational reconstruction, that discovery once again became an acceptable topic. Only now it posed problems of the second sort noted above, i.e., how does a scientific realist deal with the discovery/invention of theoretical entities?

This issue is currently, to a certain extent, a hot topic. It has taken on a slightly different shape, which is not unexpected given that many old problems never really die, they often reappear cloaked in a different vocabulary and context, wearing new clothes as it were. Today the invention/discovery battle is taking place between philosophers of science who are Scientific Realists and sociologists of science belonging to what is euphemistically known as the Strong Programme. Scientific Realists believe some version or other of the claim that the theoretical entities mentioned by our best scientific theories actually do exist. Thus, for Scientific Realists we eventually do discover the real world. There are varieties of realism but they do not concern us now. The practitioners of the Strong Programme, on the other hand, could be said to be inventionists, although they prefer the term “social constructivist.” On their view, what most of us call the real world, indicating by that that feature of reality that is independent of us, is nothing more than the result of negotiation among scientists with special axes to grind. Reality is invented or constructed. Now I am not a Scientific Realist of the average sort, but this does not leave social inventionism as our only other alternative. I advocate a new form of realism, Sicilian Realism – a view I will return to below. While I am a firm believer in the social aspects of the creation of scientific knowledge, I cannot accept the view that reality is mere invention, so I guess that makes me a Social Sicilian Realist.

8.3 Definitions

If we are going to avoid old conundrums of the sort rehearsed above, and if we are to continue on to make some sense of the three notions I would like to address, discovery, technological infrastructure, and scientific progress, we should back off the old tracks and start somewhere fresh. First, let us consider some basic notions in the form of working definitions:

- DISCOVERY: the cognitive apprehension of that which has not been so apprehended or apprehended in that manner before.
- TECHNOLOGY: humanity at work.

– THE TECHNOLOGY OF DISCOVERY: humanity at work cognitively apprehending that which has not been so apprehended or apprehended in that manner before.

These definitions present a few problems deriving from the realist/constructivist debate. For example, do we, in cognitively apprehending electrons using an electron microscope for the first time, invent or discover electrons? The way to avoid getting stuck back in the very situation we are trying to avoid is to take our definitions seriously. The definition offered above makes no ontological claims, only an epistemological one. One must "cognitively apprehend something new or in a new way." It doesn't follow that such an act entails that what is cognitively apprehended must exist. Thus, this account of discovery hopefully avoids the old problems of the realist and the constructivist, at least in the manner in which the groups mentioned above were plagued by them.[5]

Turning back to the definitions, I wanted to lay them out so as to help clarify some of the issues that are before us. But to understand the role of the technological infrastructure of science, we seem to have both too much and too little in these definitions. Attending to "cognitively apprehending people at work in a new way" is not going to help us explore the sense in which sets of artifacts generate new scientific discoveries.[6] We need something else, we need to know the manner in which further scientific work depends on new developments in the artifacts, i.e., an account of the invention and modification of the relevant artifacts in these circumstances. That is considerably more complicated. Second. we need to define "technological infrastructure."

– A TECHNOLOCICAL INFRASTRUCTURE: a set of mutually supporting artifacts and structures that enable human activity and provide the means for its development.

The notion of *mutually supporting* sets of artifacts is difficult to nail down in the abstract. What ultimately is perhaps most important is not the notion that science

---

[5]Only if your definition of "knowledge" entails existence would you be back in the old ditch in a hurry. Definitions of knowledge that entail the existence of the things that are known usually invoke a truth condition such as in "knowledge = justified true belief." Luckily, there exist accounts of knowledge that avoid the problems truth conditions present. For example, on my account, which I will not belabor here, I distinguish between what is proposed by individuals as candidates for knowledge and the endorsement of those claims by the appropriate social community, An individual may think he or she has found the truth about a particular matter, but thinking or wishing so doesn't make it so. Only when the claim has been endorsed by a particular community does it count as knowledge. The criteria the community invokes may have nothing to do with truth – it may, for example, remain satisfied with coherence or with practical efficiency. But, and this is what counts here, if the community determines knowledge, then inevitably truth will go by the board (Pitt 1983). This is the germ that the social constructivist and most relativists exploit.

[6]But it may bear on the resolution of discipline specific problems by importing techniques and individuals from other disciplines.

works within a framework of interrelated sets of artifacts, but the realization, nay, discovery, that the technological infrastructure has itself grown and developed over time in conjunction with those features of the activity we call science. Thus, I am not claiming that science, whenever and however it is or was practiced, has this kind of technological infrastructure. However, the development of a technological infrastructure is essential if science is going to continue to provide us with new discoveries about how the universe works. In short, after slow and modest beginnings, a developed science requires this kind of technological framework. The sorts of investigations and explanations it is called upon to produce require more than mere unaided human thinking alone can produce. I will return to consider the consequences of this claim later. For now, this is enough speculation; let's start to build the case.

8.4 Galileo and the Telescope

As noted in previous chapters, modern science begins with the scientific revolution of the seventeenth century in which Galileo played a major role. He was an advocate for two of the major technological innovations that made the revolution possible, for the revolution involved more than rejecting the geocentric view of the world in favor of the heliocentric. Crucial to the scientific revolution was the development of mathematical physics, or to put it in another way, the introduction of mathematics as an essential tool of science. Galileo was a prime contributor in building the case for the use of mathematics in physics (Pitt 1991). He was also, if not the inventor, at least the most successful early user of the telescope, the first to use it to explore regions of the universe previously inaccessible to us, to use the results of those explorations to challenge established theories, and as support for new ways of thinking.

I have told the story of the way in which Galileo was captured by the telescope before (Pitt 1987), but let me rehearse it again briefly. There is a temptation to view Galileo's future relationship with the telescope as almost inevitable, but I will resist it. Almost 15 years prior to learning about that device he had already demonstrated his knack for working innovatively with instruments. In 1592 Galileo was appointed to the chair of mathematics at Padua. Shortly thereafter he began giving courses in private on material outside the normal curriculum in order to supplement his income. Somewhere around 1595, following such a course on military architecture, he developed the first version of what came to be known as his military compass, a device for measuring distances and altitudes and for which he composed his first published work, appearing in 1597. By 1599 he had perfected the device into an all purpose calculator, which according to Stillman Drake "was capable of solving any practical mathematical problem that was likely to arise-swiftly, simply, without requiring previous mathematical education, and sufficiently accurate for ordinary practical purposes" (Drake 1978, p. 9). But, in what is beginning to emerge as a familiar pattern for Galileo, the basic idea was "borrowed" from a friend of his, the Marquis Guidobaldo del Monte. I say this is a familiar pattern because around the same time, 1595, he appropriated the basic scheme for what was to become his ill-fated theory of the tides from yet another friend, Fra Paolo Sarpi, and in 1609, upon hearing that a Dutch lens maker had made a device which could bring distant images near he sat down to reconstruct it, checking first with his good friend Sarpi to see if it was feasible and then hastened to beat out an itinerant peddler who was on his way to Venice with a Dutch version. Galileo had already realized that such an instrument would be of value to the Venetian navy to warn them against pirates and he also thought he could achieve some financial advantage by building one and giving it to the Doge of Venice. who was nominally his employer at the time. He succeeded in manufacturing an appropriate device and managed to get it into the hands of the Doge first. But his ploy was only partially successful. The Doge was impressed and ordered Galileo's salary as a professor of mathematics at Padua doubled to 1,000 florins. But the small print in the contract said that Galileo would also not receive another raise for life.

Miffed at being finessed, Galileo had occasion to show his new instrument to Count Cosimo d'Medici when he was home visiting Florence later in the year. Together they tried it out on the moon and discussed the possibility that the dark spots were shadows caused by mountains. But Galileo's eight power telescope was not strong enough to resolve the issue. And now, for our purposes, the crucial events begin to unfold. When Galileo returned to Padua he built a 20 power telescope, confirmed his suspicions about the shadows on the moon, wrote the Count, negotiated a deal, and moved home to Florence to take up his new post as the resident mathematician and philosopher to the Medici. Now free of the restrictions of having to teach a curriculum dictated by the Church, or having to teach at all, and urged on by Cosimo's own interests, Galileo continued to make telescopic observations of the moon and then of Jupiter, discovering its moons. He also became Italy's premier manufacturer and supplier of telescopes, adding to his instrument business. Word of Galileo's discoveries spread, placing him under pressure to publish his findings before he was scooped – priority of discovery being as important then as now. When he finally published *The Starry Messenger* in 1610, the fate of modern astronomy was sealed.

Let me take a minute to defend this rather dramatic claim. Galileo was not the first to use an instrument to investigate the heavens. The astrolabe, a device for determining the positions of the planets and the stars, already had a long and rich history. The quadrant was also a device used to determine positions in the heavens. But unlike the astrolabe or the Quadrant, the telescope produced fundamentally new kinds of information. The telescope did not, as did the astrolabe, merely assist in the refinement of measurements according to an established theory. It produced fundamentally new information about the structure and population of the heavens. It forced a transformation in cosmology. The instrument, in effect, required a major overhaul of theory. What was being demanded of theory then forced a reworking and refinement of the instrument, which in turn pushed the matter even harder toward theory revision. A basic new feature had been added to the activity of science-the interplay between instruments and theory.

Later, but not too much later, the single instrument was to become a complex of instruments. Galileo originally intended his telescope to be handheld for maritime use. But for astronomical purposes it needed a base, then a fixed position from which the observations could be regularized. Tables of sightings could now be corrected and the need for further refinements in the tables would force refinements in the telescope itself.

For example, a major problem in astronomy was determining the size of the planets. For this, Galileo's telescope with its concave lens was not the optimal instrument. In the 1630s it began to be replaced by what van Helden (1989, p. 113) calls "the astronomical telescope" which had a convex ocular and produced greatly improved clarity in its images. It also had a broader field of vision that permitted the introduction of a micrometer into the instrument itself, thereby improving the precision of measurements. This was the kind of instrument Huygens used to measure the diameters of the planets. Slowly, Galileo's simple device was becoming a set of things, each part of which could be refined and in so doing would ramify its effects on the others; perhaps not all the others all the time, but a kind of feedback effect was evident. Furthermore, the availability of increasingly precise measurements of particular features of the observable universe also forced changes in the manner in which the relative distances of the planets was calculated. So now we have the instruments and their refinements forcing changes not only in cosmology but in the auxiliary methods which augment it. In this manner the discovery of the size and structure of the solar system and then the universe was undertaken.

The story could be told without mentioning the instruments. For example, we could say,

> Galileo showed there was more than one center around which planets revolved, forcing a revision of the geocentric theory of the universe. His methods were developed in such a way as to allow for the determination of the distances between planets and the relative sizes of the planets. Modern astronomy continues his program of empirical investigation of the universe.

That says what we have been saying, but the picture it provides of science is, to say the least, impoverished. The mechanism behind the changing ideas is lost, without which mechanism we truly have no explanation.

Perhaps the need for a refined explanation is the proper motivation for including the technological infrastructure of science in our history of culture. And so, by returning to explanation and its associated difficulties, perhaps I have failed to completely extricate us from the old problems. But there is some progress evidenced here. For if we want an explanation for the development of science, we need to offer more than a recitation of the sequence of ideas produced by scientists. We need an account of how those ideas were developed and why they were abandoned and/or refined. We are thus dealing with an issue in historiography. An explanation of scientific progress and discovery requires appeal to some mechanism. That is why the history of ideas approach is inadequate. I am proposing that the mechanism that makes the discoveries of science possible and scientific change mandatory is the technological infrastructure within which science operates. In short you can no longer do philosophy of science, history of science or even sociology of science without the philosophy and history of technology.

If the science is astronomy, or even cosmology, then once we understand what it takes to do cosmology today, we must turn to the technological infrastructure to understand its results. This contextualization of our science is extremely important. The universe modern science reveals to us, is a function of this complex interaction between theory and technological infrastructure. Furthermore, it would seem that with a different technological infrastructure "science" would yield a different universe, or would it? As we attempt to answer this question we also find ourselves back in the realism debate. Let us then turn briefly to realism one more time.

8.5 Sicilian Realism and Technological Infrastructures

If we take as our starting point the fundamental claim of scientific realism -namely, that theoretical entities are real, however glossed, and couple it with the historical awareness that theories change and are replaced, we have a problem: which theoretical entities from which theories are really real? A Sicilian Realist will say that they all are. Sicilian Realism is realism with a vengeance; the universe is a very complicated place, to echo Marjorie Grene and Richard Burian. What we manage to do with one theory/technological infrastructure is to cut the universe at one of its many joints. Optical telescopes tell us planets and stars are real. Radio telescopes tell us there is more out there. Sicilian Realism admits all of this. What Sicilian Realism does not admit as at all necessary is the kind of reduction which normal realism assumes. Thus, atoms, electrons and quarks are all equally real – without one having to be reduced and explained by another. Seeing the universe in terms of atoms is a function of cutting it only one way, and there are others.

What we have to face is the fact that while there is no one necessary way to investigate nature, the mechanisms – read "technological infrastructures"-we develop to assist us set a complicated process in motion in which imagination and creativity is sparked and fed by the interplay between idea and artifact. Artifacts stimulate us to seek uses for them: how to couple them with other artifacts; how to interpret the results. Given different sets of artifacts – by definition different stimulations – we get different results. But we start small and go large in quick order. Compare Galileo's simple telescope with the complex that we need for a modern mountain-top observatory.

What are the consequences of accepting this characterization of the role of technology? Is it not the case, as I am sure some determinist will be sure to suggest, that that means not only that society is run by technology, but now science is too! No, that is not the proper conclusion to draw. It is not a question of which disembodied and reified nonentity, science or technology, controls anything. What a careful look at history will show is that as instruments are made more complex by individuals with specific objectives in mind (objectives sometimes, but not always, generated by theories), a complex of interrelated activities develops through which, by choosing certain ways to augment the technological infrastructure, certain options are opened or shut for theoretical testing and exploration. People still make the choices, and they may choose badly, taking us down a dead end. Or they may opt for a system that does not have the backup to support it. This is what happened to the nineteenth-century astronomer, William Herschel. He built a forty-inch telescope that was certainly a technological marvel. Only there were severe problems. The mounting for it proved unstable. The mirror was made of metal and lost its reflective capacity. It fell into disuse.

My point is that if you want to explain the changing claims and face of science, you have to go beneath the ideas to the technological infrastructure and then you have to look at unraveling the interactions between its parts and the mass of theories with which it is involved. It is that complex that makes it possible to apprehend new things or to apprehend things previously known but in a new way. The discovery of the structures of nature is a function of this complicated mutually interacting set of artifacts, ideas, systems and, of course, men and women. Telling that story puts us in a position to understand finally the nature of the scientific enterprise and how it generates new information. It should also alert us to the sensitivity of the system. With so much depending on so much, there are many opportunities for things to go wrong. Likewise, because of the complex of interdependent relationships, the determination of the accuracy or even of the import of the new information this system gene rates is not an easy task. Scientific knowledge becomes more tenuous and more dearly bought as the technological infrastructure grows, but it is increasingly impossible without it. Sometimes all we know is that this or that system works; we may not know what it is telling us.

References

Ackermann, R. 1985. *Data, Instruments, and Theory*. Princeton, NJ: Princeton University Press.

Cartwright, N. 1989. *Nature's Capacities and Their Measurement*. Oxford: Clarendon Press.

Drake, S. Trans. 1978. *Galileo Galilei, Operations of the Geometric and Military Compass*. Washington, DC: Smithsonian Institution Press.

Fisher, A. 1991. "A New Synthesis Comes with Age". *Mosaic*, 22(1), 3.

Franklin, A. 1986. *The Neglect of Experiment*. Cambridge: Cambridge University Press.

Galison, P. 1987. *How Experiments End*. Chicago, IL: University of Chicago Press.

Hacking, I. 1983. *Representing and Intervening*. Cambridge: Cambridge University Press.

Kuhn, T. 1962. *The Structure of Scientific Revolutions*. Chicago, IL: University of Chicago Press.

Pitt, J. 1983. "The Epistemological Engine". *Philosophia*, 32, 77–95.

Pitt, J. 1987. "The Autonomy of Technology". In Durbin P.T. ed., *Technology and Responsibility*, pp. 99–114. Dordrecht: Reidel.

Pitt, J. 1991. *Galileo, Human Knowledge, and the Book of Nature: Method Replaces Metaphysics*. Dordrecht: Kluwer.

Popper, K. 1963. *The Logic of Scientific Discovery*. London: Hutchinson.

Price, D. 1963. Big Science, Little Science. N.Y.: Columbia University Press.

Van Helden, A. 1989. "The Telescope and Cosmic Dimensions". In Taton R., Wilson C. ed., *Planetary Astronomy from the Renaissance to the Rise of Astrophysics, Part A*. Cambridge: Cambridge University Press.

Chapter 9
Explaining Change in Science

Philosophical theories of scientific change abound and, for the most part, they have one thing in common: they are theories of rational justification for changing scientific theories. That is, they are not about science per se, where science is construed as a social process whose main activity is the generation and testing of ideas about the composition and structure of the material universe. The kinds of theories of scientific change I have in mind are exemplified by the work of Popper (1959), Lakatos (1970), Kuhn (1962), and Laudan (1977). These are philosophical theories whose focus is a philosophical theory of scientific rationality that attempts to provide a justification for abandoning one theory in favor of another. They often proceed by examining the logic of the language of support for scientific theories. Science ought to change, on these views, when current theories are shown to be defective because of failed predictions, or inadequate evidence, or decreasing problem solving ability. Built into these accounts is the assumption that rational scientists ought to accept theories that meet these conditions. In short, these are all theories in the positivist tradition of the philosophy of science, where the center of attention is the logic of philosophical concepts about science. And, to no one's surprise, these treatments of the topic of change in science are sterile and unconvincing.

There are also non-philosophical accounts of scientific change – one hesitates to call them theories – that do pay attention to the social processes of science. Good history of science, both internalist and externalist, social and institutional, provides us with much valuable insight into the workings of the sciences. Then there are the sociological treatments of scientific activity. These, in general, are not so helpful, for they ignore the subject matter of scientific theories and the role it plays in the activity of scientists, concentrating only on the scientists, imposing on them a variety of unsubstantiated psychological motivations for their actions.

Here I am not going to worry about history and sociology, although what I am arguing in favor of has need of both done well. Here I am concerned to develop a *philosophical* account of change in and of scientific theories that really is about science. The heart of this project is to see mature science as an historically

Appeared in *Techné: The Society for Philosophy and Technology Quarterly Journal,* Vol. 3, No. 3, 1998. Reprinted by permission of the Editor.

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_9,

contextualized social process embedded in a technological infrastructure. A technological infrastructure is a complex set of mutually supporting individuals, artifacts, networks, and structures, physical and social, which enable human activity and which foster inquiry and action. Thus, for any particular technological infrastructure of science, the science is but one component of the technological infrastructure. The other components are, strictly speaking, not science, i.e., do not directly deal with the investigation and understanding of nature. However, without them that particular scientific activity would not be possible at that time and in that place. This implies that the activity we call science needs a social environment, which it does, and that science does not proceed in a vacuum by itself, which it does not, and that the engine of mature science is technological, not logical or psychological. Scientific change results in a change in the scientific explanation of the structure and functioning of nature. It, in turn, is the result of changes in the technological infrastructure within which the explanations are generated. For example, new information from the Hubble space-based telescope is providing the impetus for the development of new cosmological theories. Likewise, the creation of the technology of gene-splicing paved the way for new theories of genetic development. And, I will argue later, the more sophisticated and mature the science, the more embedded and indebted to its technological infrastructure it will be.

Now, to speak of the obvious, the account given above is loaded with contentious notions. To accomplish the goal of a philosophically sophisticated and historically accurate account of scientific change, I am proposing some new vocabulary and some different ways of conceptualizing familiar issues. Therefore, to begin with I will spend some time unpacking some of the more superficially obnoxious claims. After doing a little philosophical work here with some examples, I will explore some of the unsettling consequences of this explanation of scientific change.

Let me begin by providing some rationale for introducing new terminology and for offering new definitions of familiar notions. Elsewhere I have argued about the evils of reifying technology, science, government, etc. (Pitt 2000) Reification, mistaking an abstract or general noun for a thing in the world, is responsible for a category mistake with real world consequences. It allows for the misapplication of normative assessments, resulting in claims like “technology is threatening our way of life.” Nothing could be more preposterous. Technology is doing no such thing. It is the application by people of specific technologies in certain ways that sometimes creates problems. In short, there is some truth to the bumper sticker that reads “guns don’t kill, people do.” I have simply translated that insight into a general reluctance to talk about “technology” *simpliciter.* I have also gone further than merely displaying a reluctance, I have offered and defended a definition of technology, angering some, which redirects our attention to people, and reduces the emphasis on artifacts; thus, technology is humanity at work.

Now these considerations clearly have ramifications for my main notion regarding scientific change, which is a technological infrastructure. On this account, a technological infrastructure is that assembly of different forms of work relations among people that makes the doing of science possible. To put it in this way automatically includes the people, artifacts, institutions and networks that constitute the environment within which work occurs. Described in this fashion, it also entails that appeals to any specific development in a science must be historically contextualized, because science must involve the working relationships which make that particular form of social activity what it is at that time. Thus, there can be no general rule or universal explanation for changes in a science, beyond the recognition that what happened was a function of a multiplicity of factors working at that time. In short, it depends on the institutions within which scientific activity occurred and the sources of support for that activity – for example, today it might be the National Science Foundation; in seventeenth-century Florence it was either the university or the court of the Medici – the people, the politics, social influences and fads, etc. The institutions themselves are the contingent product of a variety of historical and social forces.

At this point I need to interject a caveat to forestall shouts of glee by postmodernists. By recognizing the historical contingency of science, it does not follow that science is, therefore, only one activity among others, none of which can claim some sort of epistemic virtue that allows it to be identified as the premier knowledge producing activity. The evidence, which is all around us, is that, in fact, and let me stress the fact of the matter, in fact, scientific activity, of all our activities, is the best at producing the knowledge which allows us to understand and manipulate the natural world. The historical contingency of any particular scientific success or failure does not undermine the fact that nothing has provided us with the scope and depth of knowledge science has. And by science, I mean the set of activities associated with the totality of specific investigations conducted by recognized and accepted practitioners into the structure and make-up of the universe. Now to return to the topic at hand.

If we are going to talk about scientific change, we need to talk about specific scientific changes and the contexts in which they occurred. But, it might be asked, how do we identify the context? The answer is that if it *is* a specific context, then it will be an historical item, locatable in space and time. The technological infrastructure will then be that set of working relationships without which that specific scientific development could not have happened. (Identifying the context is one thing, understanding it is another. This is where the history and the sociology come in.)

At this point two objections come up: (1) to assume that one can identify factors contributing to certain scientific developments, in the counterfactual context that were one of these factors not present, the developments in question would not have happened, suggests a commitment to a dubious sense of social causation; (2) to claim that if a technological infrastructure is that without which the scientific development could not have happened, then is not the door opened to including everything? Let us consider these in order.

First, I am not proposing an account of social causation. Rather, I am offering a justification for selecting the relevant factors for producing an accurate description of a technological infrastructure. Thus, in the historical context under discussion, given the kinds of mechanisms, tools, tool-makers, groups, patronage systems, etc., that actually existed, is it possible to give an adequate explanation of how what happened happened without including factor x or y? In so arguing, it may be the case that several different causal factors are appealed to, but no one single account of causation is being assumed. Thus, the grant from NSF that funded the laboratory in which the crucial experiment took place is causal, but not in the same way that flipping the switch on the microscope is.

The second objection asks whether we are not opening the door to including everything, since it seems that with a little ingenuity, anything can be shown to be relevant to something. To take a trivial example, if we want to explain the change from a geocentric theory of the structure of the universe to a heliocentric theory, then surely this will require that we not only detail the standard and familiar events (Copernicus and the calendar), players (Kepler and Galileo), institutions (the Medici court and the Catholic church), but that we also consider such factors as the educational and familial backgrounds of those who supported the change and those who did not, and the political and economic factors that infused their thinking, the geography of the lands they own, the number of servants they maintain, *ad infinitum*. Where do we stop? The garden of Eden?

Obviously this is not a desirable result. Further, since what actually happened in the past occurred in the seamless flow of time, fixing a context will always be arbitrary to some extent. However, the solution to the problem is one that appears naturally when we are setting it up this way. The point to stress is that the relevant factors to be included as constituting any specific technological infrastructure of science are the ones which make a difference as to whether or not the event in question would have happened. When we are speaking of science, two related criteria for selecting relevant factors come to mind: (1) making a difference means making a difference in the epistemic content of the change in question, and (2) explanatory coherence. Let us now look at each of these in turn.

Making an epistemic difference. Remember we are talking about a theory of change in the process of science. So, if scientist X is led by reason of personal ambition to establish his own laboratory rather than continue to work in Renowned Scientist G's laboratory, and X fails to get funding, and no publishable findings are produced, then it is unlikely that this is a factor to be included in the relevant factors explaining the success of Renowned Scientist G's laboratory in discovering a new mechanism. Someone might try to argue that had disgruntled scientist X continued in G's laboratory, given his disruptive personality, the eventual success of the lab would never have occurred. Now that *is* a counterproductive counterfactual, and does not contribute to our understanding of why G's lab produced the results it did. Hypothesizing as to what might have happened does not affect what did happen.

That was a negative example of sorts. Let us look at a positive example. In a complete explanation of the impact of the Hubble space-based telescope on cosmology, it will be important to include an account of the resources available to the US shuttle program which made it possible for the needed adjustments to be made to the telescope after it was launched and it was discovered that the main mirror was defective. That is, an adequate account of the new changes that are taking places in cosmology due to the observations of the Hubble would not have taken place were it not possible to fix the mirror. And yes, it is important to relate the fact that the Hubble as launched was defective; otherwise, we relapse into the let-us-only-tell-about-successes mode of history of science, which results in an inadequate explanation of why cosmological theories changed. It is inadequate because it ignores factors relevant to having those changes take place. In particular, it explains the acceptance by astronomers of the findings of Hubble observations and their willingness to allow those findings to force changes in their theories. For if the mirror had not been repaired, then the value of the resulting observations would be diminished. That it was repaired, using already agreed upon techniques, is very important. It made it possible for the Hubble telescope to be calibrated. As Alan Franklin argues,

> Calibration, the use of a surrogate signal to standardize an instrument, is an important strategy for the establishment of the validity of experimental results. If an apparatus reproduces known phenomena, then we legitimately strengthen our belief that the apparatus is working properly and that the experimental results produced with that apparatus are reliable. (Franklin 1997, p. 31)

If the Hubble could not be calibrated, then no scientific results would be forthcoming. Important for our purposes is recognizing that the calibration of instruments is crucial to using the instrument to generate new information, but it is not itself doing science. The science can only take place after the instrument is calibrated. But clearly calibration of instruments constitutes just what we have been talking about as part of a technological infrastructure, just as the instruments are part of it.

Now what we want from a philosophical theory of scientific change is an account that explains why this happened rather than something else. Consider the following: for many years I was puzzled by the fact that while everyone acknowledges Galileo's contribution to the Scientific Revolution and the importance of his last book, *Discourses on Two New Sciences,* nevertheless, Galileo's own form of scientific methodology seemed to have died with him. There is no Galilean school of physics; there are no clear Galileans as there are Newtonians. Why is this so? It took me 20 years, but I think I found the answer (Pitt 1992). As it turns out, Galileo's use of geometry is the key to understanding his science. To this end, it is also important to realize that his commitment to geometry was so strong that he urged others not to take up the study of algebra, the new mathematics then being introduced. The reason there are no Galileans is that Galileo, for all his greatness, picked the wrong form of mathematics with which to work. The cumbersome proofs of geometry were quickly being replaced by faster and easier-to-use algebraic methods. Galilean science died because geometry was replaced by algebra and then by the calculus. (It is a bit more complicated, but that is the heart of it.)

But why did Galileo stick to geometry? That requires explanation. An easy and ready account is that he was getting old, and he was virtually blind when he finished the *Discorsi,* which he had been working on virtually all his adult life. It would have been rather difficult to change mathematical methods at this late stage. This would seem to be reasonable. But there is one more thing, something that really makes a difference – for many centuries the Latin translation of Euclid's geometry in use had a flawed version of Book 5. In 1544 a new translation of Archimedes appeared which included the correct version of Euclid's Book 5, in which a clean account of Definition 4 is given. It is a definition that had been badly garbled by both Boethius and by the Arabic translators. Its correct form reads: “Magnitudes are said to have a ratio of one to another which are capable, when multiplied, of exceeding one another.” Galileo took his own definition of ratio from this relatively new translation of the definition and made it the basis for the derivation of most of his theorems. Because Galileo insisted on not compounding magnitudes of different types and because of his demand for complete rigor and proof (following Archimedes), Galileo thought he had the basis for a new mathematical method. Why did he not adopt algebra? Because he thought he had a new method of his own.

This example is instructive for several reasons. First, it helps make the point that geometry is used by Galileo in the same way that a hammer is used by a carpenter. In short, it is very much a technology. It is a tool that enhances human capacity for changing the world. Second, not every change in science is fruitful. In epistemology it is important also to explain how we make mistakes. No adequate epistemology can neglect to do that. In the history and philosophy of science it is equally important to explain failures and dead ends. It is not enough to merely account for the successes. And it is not sufficient to say that X failed where Y succeeded because X was irrational. (I find it somewhat rewarding to note that it takes work in the philosophy of technology to accomplish what philosophers of science have been unable to.) Third, it is worthwhile noting that despite the fact that Galileo’s methodology failed to attract adherents, geometry was not discarded as false or useless. It remains a viable tool.

Finally, this example puts us in a position to turn to the second criterion for selecting factors to define a context and subsequently a technological infrastructure. The determination of whether various factors should be included in the determination of an historical context must meet the criterion of explanatory coherence. If the things to be included do not contribute to the coherence of the explanation being offered, they should be eliminated. I think the role of the new translation of Euclid’s Definition 4 helps to explain why Galileo selected the method he did for his proofs and why there were no Galileans to take up his research program. The fact that he did not marry his long time mistress does not. Nor is it relevant that at this time Cardinal Richelieu held the de facto power in France.

Let us now turn to the question of how an historical context contributes to our understanding of a technological infrastructure for science. It becomes one when the factors selected can be shown to make an epistemological difference with respect to specific scientific developments, thereby explaining what happened in a manner which brings the relevant factors into a coherent story. That it is a *technological infrastructure* is a function of the fact that it identifies the players – human, artifactual, epistemological, institutional – and their interrelations in which the events in question took place.

A mature science is a complicated thing. It is not merely a theory. By concentrating on the logical structure of theories, philosophers of science have done some good things, but they have not made it possible to do the important philosophical job, which is, as Wilfrid Sellars put it, “To see how things, in the broadest possible sense, hang together, in the broadest possible sense” (Sellars 1963, p. 3).

Concentrating on the logic of theories does not tell us how science gets done. Before there are data to be used as evidence, there are laboratories and the places where the laboratories are located. And where they are located makes a difference. For example, different kinds of pressures apply in commercial labs as opposed to university laboratories. There are different objectives to be met. In some commercial labs, the emphasis is on commercially viable results. In some academic labs the emphasis is on securing grants to ensure the continuation of the research program (and the generation of overhead for the university administration to play with). In addition to the kinds of issues just noted, the doing of science includes laboratory assistants, experimental apparatus, the interactions among the members of the community (no, I am not talking about the social construction of scientific results) which fuel ideas and techniques. In short, if we play out the list of things we need to consider, we will find ourselves looking at the full scope of the working relations among those people involved in the investigation of nature. And if technology is humanity at work, then those relations and players constitute a technological infrastructure.

In closing, it seems appropriate to consider the down side of the view I am proposing. Cosmology is the science concerned with explaining the universe as a whole. It uses data gathered from a variety of instruments, telescopes of various kinds in varying locations. These instruments themselves embody numerous theoretical assumptions, from optics to electronics to the manufacture of ball bearings. The increased use of computers to manipulate data incorporates yet another wide ranging set of assumptions, some of them having to do with computer languages, others with the reliability of hardware. The kinds of explanations cosmologists generate do not, therefore, merely rely on the evidence pure and simple. The question, to my mind, is, how much of the theory is a function of the technology? In mature sciences, it appears that the more embedded the science is in its technological infrastructure, the more the infrastructure drives the science. Thus when we attempt to ascertain the cause of a change in theory, we will find it increasingly difficult to point to specific causal factors. I suppose we could simply say that it is the Hubble telescope that is forcing us to revise our cosmological theories. But that would simply be false. How that instrument is used, the kinds of support systems it requires, and how they influence the generation of images, cannot be ignored. If what I have been suggesting is correct, then we need to know a great deal more about the supporting systems and the environment in order to understand just what it is the science is telling us. And when the science is thus embedded in its technological infrastructure, changing scientific theories can only be accomplished by rejecting the technological infrastructure or by finding another theory which uses the same infrastructure, at which point the science is still captive to the technology. Thus, explaining scientific change will require a full account of the technological infrastructure of that science if we are to understand what kind of a change we are witnessing.

References

Franklin, A. 1997. "Calibration". *Perspectives on Science*, 5, 1.
Kuhn, T. 1962. *The Structure of Scientific Revolutions*. Chicago, IL: University of Chicago Press.

Lakatos, I. 1970. "The Methodology of Scientific Research Programmes". In Lakatos I. and Musgrave A. eds., *Criticism and the Growth of Knowledge*. Cambridge, MA: Cambridge University Press.

Laudan, L. 1977. *Progress and Its Problems*. Berkeley, CA: University of California Press.

Pitt, J.C. 1992. *Galileo, Human Knowledge, and the Book of Nature: Method Replaces Metaphysics*. Dordrecht: Kluwer.

Pitt, J.C. 2000. *Thinking About Technology*. New York, NY: Seven Bridges Press. 2006 Http://www.phil.vt.edu/Pitt/jpitt.html

Popper, K. 1959. *The Logic of Scientific Discovery*. London: Hutchinson.

Sellars, W. 1963. *Science, Perception and Reality*. London: Routledge and Kegan Paul.

Chapter 10
The Dilemma of Case Studies

Toward a Heraclitian Philosophy of Science

After Kuhn (1962) cast doubt on the usefulness of abstract positivist models by appealing to the history of science, many philosophers have felt compelled to use historical case studies in their analyses. Kuhn however did not tell us how to do this. Further, it is not clear exactly what appeals to case studies accomplish. We can frame this issue as a dilemma. On the one hand, if the case is selected because it exemplifies the philosophical point being articulated, then it is not clear that the philosophical claims have been supported, because it could be argued that the historical data were manipulated to fit the point. On the other hand, if one starts with a case study, it is not clear where to go from there – for it is unreasonable to generalize from one case or even two or three.

I will argue that even very good case studies do no philosophical work. They are at best heuristics. At worst, they give the false impression that history is on our side, sort of the history and philosophy of science version of Manifest Destiny. If historical studies are to be useful for philosophical purposes, they must be extended historical studies that contend with the life span of a scientific problematic. It is not enough to isolate a single experiment or to look at the activity of a lab under one director. One needs to place the case in the context of a problematic and to explain a problematic in terms of its origins and its fate. (Pitt 1992) But even if this were to be accomplished, it is not clear what philosophical work is being done. This may be, at best, history of ideas. The point here is simple: just as philosophical problems are not problems about the single case, historical issues are particular and must be seen in context. But seeing an historical issue in context does not by itself suggest any particular philosophical point. It may be that the problem here lies in our understanding, or lack of it, of what constitutes a context. The importance of understanding the appeal to historical contexts is to show how doing history in context limits the possible range of philosophical ideas and explanations. By way of example, I will consider the philosophical question of what constitutes a scientific observation. I will argue that a serviceable universal account of scientific observation is not possible, because the activity of making a scientific observation depends on, among other things the sophistication of the technology available at the time, hence what we mean by a scientific observation changes. What is allowed as an observation varies in time, place and with respect to changing criteria influenced by technological innovation.

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_10,

If I am right, this view provides a serious basis for rejecting Kuhnian paradigms. Problematics have histories, but that does not mean they are stable over time. Quite the contrary, the reason why it is important to appeal to problematics is that they change even as they serve to restrict research to certain topics. And it just may be this Heraclitian characteristic is a defining feature of science. As philosophers we seek universals, but the only universal regarding science is change. That seems to be a fact. But, it might be responded, as philosophers we are also interested in the normative – our job is to attempt to show what we ought to mean by x or y. While that is true, in our normative guise we also cannot ignore what in fact is the case. The hard job is to figure out how to do that. The lesson to be learned is that if philosophers wish to use historical cases to bolster their positions, then we will have to use very long studies and we will have to figure out how to relate the history to the philosophical point without begging the question.

The issue of not begging the question looms large. Let us start with a big question, which is continually begged: just *what* constitutes a case study? This goes to both horns of the dilemma, but particularly to the question of how to avoid appearing to manipulate historical date to fit philosophical theories. For without credible criteria for selecting or identifying a case as a case the charge can be legitimate.

Despite the currency of case studies, there are currently no criteria available to ascertain when we have one before us. We select the historical episodes we do for a variety of reasons with few, if any, operative guiding principles. I propose that we *can* develop a set of criteria for selecting a case study, but there are several costs. The problems involve the selection criteria. For example, if we want to start with the science and see where that leads us, then, without begging the question, we have to find the science. Identifying the science in question in a non-whiggish fashion is a delicate matter. We simply cannot assume that what we call physics today, is what the scientists practicing physics in 1830 would call physics. We can't find the case study because we can't find the science in which it is a case. But there is a way out.

The way out is to proceed by identifying a problematic. A problematic consists of a set of intellectual concerns that motivate a scientist or a group of scientists to pursue the investigations they do. I suggest that this characterization skirts the demarcation question because where a group of investigators can be identified we have a social fact as a starting place. For an example of such a group, I suggest Copernicus, Tycho, Kepler, Galileo, Clavius and Scheiner. Their interests need not constitute a 1 to 1 correspondence, but each had to consider what the others had to say as relevant to their research interests either singularly or in sets.

Now for the cost: problematics have their own history, they have starting points and end points, and in between they change, mutate, sometimes they evaporate, sometimes they metamorphize into something new. Further, in the course of working within the problematic, what emerges may not be what was expected. Finally, although this may seem obvious, to identify a problematic one must position it historically. This is to put the problematic in context, which is difficult, for in any historical setting there are many contexts, and we must avoid begging the question by selecting a context which conveniently supports our concerns. In short, if we start with case studies, we are assaulted on all sides by issues of question begging. Let us look more deeply at the notion of context.

What do philosophers expect to accomplish by appealing to history and historical contexts? We all know that, contrary to popular belief, Kuhn was not the first to wag an historical finger at us; Norwood Russell Hanson (1961) was doing history and philosophy of science in the nineteen fifties and his work was well received within the inner circle. Lakatos (1971), borrowing freely from Kant, asserted that philosophy of science without history was empty and history of science without philosophy was blind. In what sense is philosophy of science without history blind? Have we not been able to see clearly through the lens of logic to important structural characteristics of, for example, explanation and confirmation? If the claim is that what we have come up with doesn't match what scientists actually do, then it is not clear that that is a valid criticism since we philosophers have a normative, not merely a descriptive role to play. Determining the logic of key concepts and working that out is a perfectly legitimate activity. What is it that history is supposed to supply?

In part this is a question concerning what it is we think we are doing. Or to be brutally frank, what are the goals of philosophers of science who use history in someway or other, or more specifically, who see the appeal to historical context as important?

Let us begin by reviewing the evils contextualization is supposed to avoid:

1. Whig History; a term coined by Herbert Butterfield (1931); it refers to the attempt to impose current categories of analysis on past historical events.
2. Universalism – a corollary to (1); the idea that certain features of science are constant over time.
3. Modernism; the insistence that the most important developments of any epoch are science (conceived in contemporary terms) related.
4. Abstraction; the reification of key features of a period.
5. Internalism; the process of examining the work of a person by appeal only to his or her notes and texts without consideration of any social or external factors – falls prey to (3) or to (1), since to really know is already to understand the context in which an author writes.

Assuming it is possible to avoid the above, there remain serious dangers facing the contextualist. For what the historian concentrating on context does, having avoided these five cardinal sins, is to concentrate on individuals and to consider the influences on and the consequences of these influences for their work. What this means is the following. First, given (5), all that the historian can do is to reveal the social and intellectual factors that might be said to motivate the views expressed by the particular historical figure under consideration. For to provide a close analysis of the work of the person in question (it must be a person to avoid (3) or (1)), exposing its logic or even its content amounts to (5). To the extent that the views of some person or other are to be considered, it is only by virtue of his perceived audience or influences. But determining who are the audiences and influences falls to the historian to identify since, we are told, historical figures cannot be trusted to know whom they really are influenced by or to whom they are really responding. How the historian avoids (1), (3) or (5) – or how he or she knows whom to identify as the relevant audiences or influences remains something of a mystery. The problem here is fundamental. (a) It is not enough to say "*x* read *y*" – since that alone does not establish influence except in a trivial manner; (b) nor is it enough that *x* quote *y* or that *x* admits to either reacting to *y* or even attempting to extend *y*'s position – since *x* may not know what really motivates him or her. (The contextualist has opened the door to this objection by using it to reject Internalism, i.e., *tu quoque*); (c) it is equally inadequate to cite who read *x*, for it can be the case that *x* was read for all the wrong reasons – (c.f., the misuse of Nietzsche by the 3rd Reich).

The contextualist historian is now left in the position of arbitrarily identifying people in places and can only hope that the preponderance of the evidence and correlations account for what *x* said about *y*. The laudable intent of the contextualist is to show that great figures do not emerge from a vacuum. The problem, however, is that there is no obvious principle of selection which guides the identification of people who or events that allegedly transform the vacuum into a social context. The result can be that the figures highlighted can be minor or obscure; likewise for social factors. Without a well-articulated and defensible principle of selection, the attempt to construct a context is at best arbitrary; at worst it is self-serving. Why certain figures are identified is also not clear, since all the objections used above with respect to *x* apply equally well to these problems. The contextualist project, seen in this light, is hopelessly flawed.

As we have seen, if we pay too close attention to the standard justification for contextualization, the program collapses. And yet there is something positive to be said in favor of each of (1)–(5), i.e., the rejection of Whig History, Universalism, Modernism, Abstraction, and Internalism. It is just that taken together nothing much is left. Have we taken a wrong turn somewhere?

It might appear that we have been led to our unhappy conclusion by concentrating on only one aspect of the contextualization of history, i.e., the individuals. But the collapse of contextualism does not occur only when individuals are the subjects of discussion. For example, an anti-Whig historian will also justifiably reject talk of "science" in the sixteenth century, there being natural philosophy for the study of the natural world. Thus the reification of concepts also seems to be a problem.

So, what is the point of contextualization? What is the appeal to context supposed to accomplish? Minimally a context is supposed to provide an *explanatory framework* for specific historical developments, i.e., it sets the stage on which the historian's explanations will be seen to make sense when offered. The crucial mistake made by advocates of historical contextualization is to give the impression that there is only one appropriate context that satisfies the explanatory-allowing role. The writing of history is necessarily selective. However, the shift from individuals or activities such as history or art to context is no less selective or arbitrary, for (with apologies to Nelson Goodman (1953)) contexts are where you find them. For example, consider the contexts in which Galileo could be said to have operated. (1) The Renaissance, (2) The Scientific Revolution, (3) The Medicean Court (*pace* Biagioli (1993)), (4) The Archimedean tradition, (5) The Euclidean tradition,

(6) The Aristotelian tradition, (7) The Platonic Tradition, (8) The Medieval tradition, (9) the battle between the Vatican and the Italian secular states for political control of the Italian peninsula, (10) the Age of Exploration, (11) The Age of Elizabeth, (12) the sixteenth Century, (13) the seventeenth Century, (14) a personal struggle to financially support his family, (15) the personal politics of the struggle between theologians and natural philosophers (*pac*e Redondi (1987)), (16) The Counter Reformation (*pace* Shea (1972)). And so far we haven't even begun to explore whether we should approach Galileo as an engineer, a physicist, an astronomer, an instrument maker, an amateur musician, a father, a philosopher, a theologian, a good catholic or an irritation of the Pope's.

However, picking the relevant explanatory framework may not be as difficult as I appear to be suggesting. The trick lies in figuring out what it is about the person or the event you want to explain. The mistake to be avoided is assuming there is necessarily only one explanatory framework. Even so, there is something more problematic than determining which framework to pick, that is the problem of determining what constitutes an appropriate explanatory framework or frameworks for a topic, i.e., what constitutes an explanation in these contexts, or what constitutes an historical explanation *simpliciter.*

To ask this question assumes that there is one kind of historical explanation that fits all sizes. Clearly, this is not the case. We actually have two questions here – first there is the problem of selecting an appropriate framework. Second, once a framework has been selected, we still need to be able to sort out what kinds of explanations are appropriate and satisfactory and which ones are not. Answering these two questions is clearly beyond the scope of this paper. I will concentrate here only on one part of the second question and I will do so by trying to answer a slightly different question, namely "What do we want from an historical explanation?" – i.e., what is the point?

Rephrasing, if the question reads, "Why do we seek historical explanations", it sounds a lot like "Why do philosophers of science turn to history?" One tried and true answer is "To learn from the past". It is unlikely, however, that we seek historical explanation only to understand how we got to where we are now. We seek more from history – not merely an answer to the question "How did we get here?" but also "how can we avoid ending up in this situation in the future?" There is little doubt that that question cannot be answered for several reasons; first, the analogies between the past and the present are just those, analogies. Learning from the past is only as successful as the strength of the analogy between past and present, and in drawing the analogy we need to be careful not to fall into the trap of doing Whig history. (2) There is no single fact of the matter of the past – more information is constantly surfacing, depending on what we think we need to know. Ideologies, cultural fads, etc also influence the plasticity of our histories.

And yet the situation is not hopeless. The search for criteria by which to select frameworks to use in obtaining answers from the past depends as much on the perceived state of the present as on our perception of the options for the future. And it is in the latter that we will find out clues to the adequacy of historical explanation. The central idea is the notion of a *coherent story*. What makes for an adequate explanation is the sense that our account of why things happened in the past hangs together with what we know preceded and followed the event in question.

Appearances to the contrary notwithstanding, I want to argue that this is not Whiggish. Nothing in this proposal of a coherent story suggests that we necessarily must see what happened in the past in a direct, causal line with the future, which is our present. It therefore makes no sense to talk about the global importance of current events, theories, etc., since that verdict awaits the future. A relevant set of contexts can be identified in terms of their explanatory value, i.e., the coherence they contribute to the story accounting for why what happened happened. To the extent that the failure to include certain factors can be shown to be relevant to understanding what happened after the events in question justifies expanding the set of contexts. So, an historical context is a set of factors that provide an explanatory framework for an event, a person's actions or work, or a social tend, etc. The adequacy of the context is a function of its ability not only to account for the event in question, but also for its prior and subsequent history.

All that having been said, we still cannot account for the philosopher of science's appeal to history. The job of explaining why the past was the past is the historian's job. The philosopher who looks to the past as revelatory of the present is doing bad history, so that can't be the justification. Nevertheless, there *is* a philosophical job to be done with respect to the past. One of the features that need uncovering when we try to understand an individual's actions is the set of assumptions with which he or she is working. In particular, we need to know what were the expectations at play at the time in order to assess the quality of the work being done. Uncovering assumptions and exploring texts for hints to expectations are jobs philosophers are good at. But in so doing, we learn little about what is relevant for today. So, at the moment, it is not at all clear what the cash value of case studies is for the philosopher of science who starts with history.

Let us now turn to the other horn of the dilemma. Instead of starting with historical cases selected for the way they are assumed to illuminate contemporary philosophical issues, or for providing the data for building a philosophical theory, let us start from the side of theoretical questions. Unfortunately things don't get any better. The kind of question I have in mind is "what is a scientific explanation" or "what is a scientific observation" – when we look to history to answer such questions, we stumble in many ways over assumptions that at first seem innocent and yet eventually prove fatal. For example, when one asks what is a scientific observation, it seems that we are asking about the "observation" part, assuming that we know what "scientific" means. But even if we have a good solid understanding of what "scientific" means (which we don't) we can't simply assume that we can apply that understanding backwards in time – to do so is to engage in Whig historiography, which we all now know is inappropriate.

Now, let's assume that we not only know what "scientific" means, but also what "observation" means and what "scientific observation" means (which we don't) – now each of these expressions has a history and their meanings have changed over time. To look to Galileo's telescopic observations to enlighten us as to the meaning of "scientific observation" today is to run rough shod over good historiography and to assume that philosophical analysis has some sort of a temporal a priori intellectual legitimacy and that as philosophers we can appropriate history to our own ends, confirming our assumptions. What would it mean for Galileo to make a scientific observation of the moons of Jupiter? "Scientific" is not a term in play at the time. To claim that his observations were scientific is reading backwards from the present into history, which is unjustified. Second, it is not clear that at the beginning of the seventeenth century there was a formal understanding of what was meant by an observation as opposed to any number of other similar activities such as seeing, perceiving, sighting, etc. (See Pitt 2007)

Finally, with the advent of new instruments we can trace the transformation of the concept of an observation. We can agree on why certain highly constrained settings in a lab can yield observations. But what about the pictures of the surface of Io being sent back from the Galileo probe? There are a number of different kinds of steps in between the taking of a measurement of Io and our seeing the result here on earth. Transmitting devices need to be aligned, involving computers and computer programs. There is the encoding of the measurement and then the sending and the assumption that nothing happens to it while it makes its way from the orbit of Jupiter to Earth. Then there is the reception of the data, more computers and programs to transform the encoded data into a picture and Lo! An observation! To accept those pictures as observations requires an expanded understanding of what constitutes an observation from the simple naked eye seeing of nature and our report of that seeing to something considerably more complicated and sophisticated. The extent to which we have accepted the fact that we can use remote instruments to make observations is a far cry from simple seeing.

I propose that not just "observation," but all of the concepts we use to discuss science are in constant flux. Peter Galison (1998) makes that case with respect to the meaning of "experiment" in the twentieth century. What constitutes an explanation, evidence, data, observation, etc., all change over time and usually in response to some technological innovation.[1] That being the case – i.e., that the meanings of these concepts are in constant flux, it would seem impossible that we could learn anything about our present concerns from the past. And so once again, the question remains as to what we can gather from case studies.

So where does this leave us? We don't know what a case study is – if we shift to a problematic we open up a can of worms – problems are embedded in historical contexts, but selecting the right context without begging the question isn't obvious. On the other hand, if we assume that concepts associated with philosophical analyses of science have some sort of atemporality we violate legitimate historiographic concerns.

Does this mean that Kuhn's wake-up call to philosophers to pay attention to history was misguided? I think not. However, as philosophers we need to lower our sights or perhaps we need to raise them and consider more than only abstract

[1]Documenting this claim is the object of a project currently in progress, *Seeing Near and Far, a Hericlitian Philosophy of Science*.

philosophical criteria. Further, we need to develop a more robust sense of the sloppiness of our conceptual history. We seek precision, definitional clarity, analytic sophistication. These are good – but there is more to understanding: depth, flexibility, and a sense of the give and take and contingency found in history.[2]

References

Biagioli, M. 1993. *Galileo Courtier; the Practice of Science in the Culture of Absolutism*. Chicago, IL: University of Chicago Press.

Butterfield, H. 1931. *The Whig Interpretation of History*. London: G. Bell.

Galison, P. 1998. *Image and Logic; a Material Culture of Microphysics*. Chicago, IL: University of Chicago Press.

Goodman, N. 1953. *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.

Hanson, N.R. 1961. *Patterns of Discovery; an Inquiry into the Conceptual Foundations of Science*. Cambridge, MA: Cambridge University Press.

Kuhn, T. 1962. *The Structure of Scientific Revolutions*. Chicago, IL: University of Chicago Press.

Lakatos, I. 1971. "History of Science and Its Rational Reconstructions". In John W. and Gregory C. eds., *The Methodology of Scientific Research Programmes. Philosophical Papers Volume I*. Cambridge, MA: Cambridge University Press.

Pitt, J.C. 1992. "Problematics in the History of Philosophy". *Synthese*, 92, 1.

Pitt, J.C. 2007. "Seeing Nature; Origins of Scientific Observation," published as "La vision de la nature: emergence de l'observation scientifique". In Burian R.M. and Gayon J. eds., *Conceptions De La Science: Hier, Aujourd'hui Et Demain*. Paris: Ousia.

Redondi, P. 1987. *Galileo Heretic*. Princeton, NJ: Princeton University Press.

Shea, W. 1972. *Galileo's Intellectual Revolution; Middle Period 1610–1632*. New York, NY: Science History Publications.

[2]This leads me to believe that my colleague Richard Hirsh may be correct when he suggests that if you can't call the guy up and interview him it isn't history.

Chapter 11
Technological Explanation

The purpose of this chapter is to provide an account of technological explanation. The topic is relatively unexplored. Therefore, in many respects this is as much an attempt to lay out the territory that needs to be covered as it is a fully adequate theory of technological explanation.

The structure of the chapter is as follows: after a discussion of the need for a theory of technological explanation, I differentiate technological explanation from physical, teleological, psychological, and social explanation. Attention is then directed to answering questions as a means of providing technological explanations. A distinction between internal and external audiences is also introduced to provide a means for characterizing different kinds of explanations in terms of the audiences to which they are directed and the kinds of questions which when answered provide the appropriate explanation. Next the concept "system" is introduced. The idea to be developed is that a crucial component of a technological explanation is placing the artifact/mechanism/activity/function to be explained in a relation to other parts of the system in which it is embedded. The strong position that there is no explanation without relating the thing to be explained to something else is laid out. This idea is elaborated by showing how artifact specific issues such as the design, function or structure of an artifact can only be adequately explained by reference to the system in which they have a role. Inevitably talk of system will bring us to social systems as technological artifacts (technical aspects of social systems?) and the degree to which an explanation of some aspect of a technical artifact requires appeal to some aspect of a social system in which it functions.[1] Finally, following a discussion of some examples, there is a discussion of the lack of symmetry between explanations of technological successes and technological failures and the importance of that lack of symmetry.

Originally appeared in Handbook of Philosophy of Technology and the Engineering Sciences, Springer Academic Publishers, 2009. Reprinted with the permission of the publishers.

[1] A theme that follows from this line of thought and also further develops a position laid out in the last chapter of Pitt (2000) and elsewhere in this volume is that since mature sciences are embedded in a technological infrastructure, any adequate theory of scientific explanation requires a theory of technological explanation. Unfortunately a full development of this relationship would take us far afield from the topic at hand.

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_11,

11.1 Relevance

Why do we need a theory of technological explanation? The standard answer to a similar question, “Why do we need a theory of *scientific* explanation?” has a fairly straightforward answer: “Because science is supposed to explain the world and we need to know when such explanations are good ones.” But clearly a similarly phrased response to the question about technological explanation won’t do.

The task of justifying the need for a theory of technological explanation is that, unlike for scientific explanation, which only has to account for why things in the natural world do the things they do, there are lots of different kinds of questions we can ask about our technologies. That suggests that either there are many types of technological explanations or that the account we develop will be unique. While that doesn’t answer the primary question, it does suggest a strategy: first, identify what technological explanations are about.

I will not talk about Technology, but rather about technologies, specific technologies.[2] They are everywhere. And that is the very point. We are surrounding by, embedded in, dependent on, supported by, amused through our technologies. They make the way we live possible. They also have positive and negative impacts on us and on the ecology of the planet. They are the expression of the creative, inventive, and perhaps malevolent aspects of our collective character.

In short, our technologies and how we use them are what marks us out as human.[3] That means that if we want to know what we are and how we got to this point we need to explain how we created our technologies and how they assisted us and/or restrained us. This means that a theory of technological explanation is relevant to all forms of human activity, since they all involve dealing with technologies, including science. What a theory of technological explanation will provide is the means to explain how an artifact came to be what it is. This can be a causal story, but it will also be partially an appeal to a variety of social factors. A theory of technological explanation will also provide the means to explain the role of the artifact in our lives and the impact introducing the artifact had on our social structures, goals, and values. It will, finally, also provide the means to explain technological failures and to distinguish questions concerning system failure from issues of assessing blame and responsibility.

11.2 Technological Versus Scientific Explanation

As noted above, the key to developing an account of technological explanation is answering the question: what are we explaining? Unlike in science, where in the past it has been assumed, incorrectly, that the answer is fairly straightforward, as noted above the focus in technical explanations can be multifold. The traditional view has

[2] See Pitt (2000, chapter 1).

[3] Not everyone agrees with this claim, especially Ashley Shew. See her 2007.

it that the purpose of a scientific explanation is to help us understand why the world works the way it does in specific circumstances. The history of the development of theories of scientific explanation reveals consistent attempts to find a general theoretical account of what makes for a good explanation that applies across all the sciences.[4] From the first major modern efforts in this direction by Hempel in 1948, there have been numerous such attempts to construct a general theory of explanation. These efforts may now decline as work in the histories and philosophies of the individual sciences reveal major differences among them in terms of methods, that in turn seems to require the development of individual theories of explanation for physics, biology, chemistry, geology, etc.[5]

But even if universal theories of scientific explanation decline in usefulness in the natural sciences, it is not clear that they could have been of use for the technical sciences when considered apart from the natural sciences, that is, when technologies are not seen as mere applied science and the technical sciences, like the various disciplines of engineering, are considered in their own right. It is the very universal ambition of these philosophical theories that renders them inapplicable in the technical sciences. Consider, for example, Hempel's account.

11.2.1 The DN Theory

Known as the Deductive-Nomological Theory, Hempel's Covering Law Theory (DN) requires that in the premises of the deductive argument that constitutes the explanation there must appear the expression of at least one law of nature. If we were to apply this account to technological issues needing explanation, this would require that there exist natural laws governing technologies. However, there are no natural technological laws, except maybe the Law of Unintended Consequences. In the absence of laws of nature for technologies, this sort of theory will not work.

As noted, a significant component of a DN explanation is the requirement of the presence of a law. However, neither the model, nor the theory of explanation that supports it provides a decision procedure for selecting which law to choose. Laws, in this context, are formulated in the context of theories. Hempel's unarticulated assumption seems to be that in science there is only one theory in play at a time and hence there is no need for a decision procedure. But this assumes that even when there is only one theory in play that we know which of its many laws and generalizations to employ in *this* explanation. Further, since recent historical and sociological considerations have begun impinging on philosophical ruminations, we now know that scientific work rarely takes place in such a clean environment. In the context

---

[4]The literature on explanation is vast. It is far too large to discuss here. Fortunately, there are two major narrative histories of the debates. The first is by Wesley Salmon (1989) in his introduction to the edited volume, Scientific Explanation. The second, more recent discussion, is by Jeroen de Ridder (2007).

[5]Robert Cummins also alludes to this possibility in his 1975 article.

of doing science, there are generally multiple accounts vying for supremacy and much of the excitement in the sciences comes from the clash of theoretical explanations. In scientific, as in technological, contexts, to forget that these are very human activities, subject to the strengths and weaknesses of any human epistemological endeavor, is to set goals that are unattainable. To do so is to ignore what can actually be accomplished.

The DN model and responses to it are only one type of explanation and they were devised with physics as the model science to do the explaining.[6] It is a way of explaining why natural things in the world do the things they do by appeal to the structure of nature. And that is one reason they seem inappropriate for matters technological. Technological explanations are unique insofar as they concern technological issues, issues that emerge because of things human beings have built. Some of these items require knowledge of how nature works in order to be constructed, e.g., optics for telescopes, chemistry for drugs. But the request for a technological explanation will not be exhaustively satisfied by an appeal to the physics or chemistry of the matter. Why is this? One response, to be developed below, is that an adequate technological explanation must consider the audience to which it is addressed. But a second reason is that technologies are made by humans and at some point it is always appropriate to consider the impact on human living of a given technology. So a physical explanation is never an exhaustive technological explanation. Finally, while the design and production of technological artifacts is a complicated set of interlocking and overlapping processes, it is not the case that there is only one way to proceed in developing technologies. The shape, function, and components of an artifact take the form they have because of a variety of contingent circumstances, i.e., there is no one way to do this job. It is, perhaps, capturing that sense of contingency that is the most important and most difficult part of a theory of technological explanation.

11.2.2 Other Theories of Explanation

Other theories of explanation have been devised which try to account for the phenomena to be explained by appeals to other factors than the natural world. There are for instance, teleological explanations, social explanations, and psychological explanations. While each of these can provide good explanations for some things and/or events, they are all individually inadequate for the purposes of technological explanation because, like physical explanations, they tell only part of the story. Thus a teleological explanation accounts for the behaviors of a given phenomenon in terms of the final end for which it was constructed. However, that tells us nothing about why it has the design it does or why it is constructed out these materials rather than these others nor how its parts work together. Likewise, social explanations ignore the physical world in which artifacts are embedded. This is not to say

[6] See Pitt (1988).

that there are not social components of an adequate technological explanation, it is just that more is needed than just the social. Thus, if I want an explanation of why this dam was built here, it will surely not be enough to appeal to the politics involved. The geology of the site plays a role as well as the availability of materials, the appropriateness of its location etc.

Regarding the technological the task of constructing a general theory of explanation has not even been an objective, with one exception to be discussed below, since it is not clear that a general account of explanation in this domain is possible given the kinds of things that can be in need of explanation. It may also be the case that since there was a common misconception that since, it is alleged, technologies are merely applications of some science or other,[7] then the search for a technological explanation will naturally revert to a scientific explanation. In fact, it is not clear what a theory of technological explanation is supposed to explain. If we focus only on technological artifacts (leaving systems and social technologies aside for the moment), the design and function of the artifact can each be in need of explanation (Kroes 1998). Let us refer to these points as *artifact specific issues*. Attending to artifact specific issues is important, but we need also to explain the artifact's social impact and the values or value structure associated with its development and evaluation (Winner 1986). When you introduce the topic of values, things get very sticky very quickly. Some will argue that technological artifacts are foils of ideological systems (Winner 1986), others will claim that they are value-laden in other ways (Kroes, personal communication) and others continue to claim that artifacts are value neutral (Pitt 2000). In each of these scenarios it may be the case that it is the decision-making structures behind the development of the artifact that need to be explained in order to understand how an object came to be what it is and do what it does. To develop a theory of explanation for these latter subjects will take us further and further away from the task of explaining specific features of specific artifacts. In the long run this is what might be necessary; to explain some feature of an artifact may require that ultimately we have to explain the motivation for its coming into being, which will in turn require an explanation of the social and economic system from which it emerged, etc. However, there is a danger in taking only this direction. To move into this mode runs the risk of falling under the seductive spell of social constructivism and its mantra, "it is social all the way down". However, it is not the case that all is social, not, at least in some non-trivial fashion. Nevertheless, acknowledging the social allows for a distinction between artifact specific explanations and social explanations. This is a distinction that will prove useful as we explore the kind of explanation we seek with respect to the technological.

One non-constructivist area that has attempted to explain technical or technological development in general terms is economics. Like most economic accounts, the appeal to rational self-interest, market forces, evolutionary scenarios or class conflict, presents a narrow vision of the factors involved in our complicated technical world (see Elster 1983). For reasons explained below, technical explanation

[7] See Pitt (2000) for an argument against the technology-as-the-handmaiden-of-science view.

cannot rest on what amounts to ideological critique couched in economic terminology. More to the point, economists of various stripes have been concerned to explain technological *change* in economic terms, which is not the same as offering a technological *explanation*.

11.3 Questions and Internal and External Audiences

Returning to artifact specific issues, we find that a close examination of their *explanatory demands* takes us beyond the specific to a level of greater generality, even if short of universality. The issue of explanatory demand is crucial to unraveling the problem of technological explanation. If, as writers from Hempel to Achinstein agree,[8] scientific explanations are answers to why-questions, to form an adequate answer depends to a very large degree on who is asking the question and to whom the answer is directed; in other words, an adequate answer depends on the audience. But explanations are more than answers to why-questions. They also answer how-questions. Further, there are at least two very different audiences asking questions about technologies: internal and external and it is not obvious from the start which type of question they are asking and what is the best way to answer it.

The internal audience consists primarily of workers within a specific technological context, i.e., engineers, designers, etc. These are the individuals who are involved in developing the technologies in question. Their questions concern issues surrounding the design, the materials employed, the nature of the system into which the technology fits (more on systems below), meeting the design specifications and so forth. In short, they ask a lot of how-questions. The external audience consists of technology users, entrepreneurs, developers, politicians, critics, etc. Further, with an external audience there will be differing demands of generality.

The same question can be asked by different audiences but it can be answered by appealing to more or fewer specifics. For example, if the question is "Why did that light bulb turn on?" one specific answer could be that I flipped the switch, thereby construing the question as a how-question. That might be all that is needed. However, the simple question may mask a deeper one such as "where does the electricity to power the light come from?" which is a more complicated how-question. The answer to that deeper question might appeal to the concept of an electrical grid and how distinct places, like houses, get connected to the grid and how electricity is dispersed throughout a local site such as a house through a wiring system. In this context a why-question might be of the form "Why is the switch placed at that height?" The answer appeals to building codes and opens the door to social factors. A full explanation of why the light bulb turned on, therefore, requires a lot of ground to be covered, from the wiring of the house, the electrical grid, building

[8] See Pitt (1988).

codes[9] (please refer to Bucciarelli: Do you know how your telephone works? (1994, chapter 1). It involves a little bit of science, a lot about design and function and an acknowledgment of the role of the social. So it is beginning to look like a technological explanation is going to be a complicated thing. But the fact that it is complicated does not mean it is unstructured. The structure is supplied by invoking the concept of a system.

11.4 Terminology

A word on terminology – technical versus technological. "Technological" is used here to discuss systems, both mechanical and social, as means of controlling and manipulating the environment, writ large and small. "Technical" is used to refer primarily to artifacts and mechanical systems, but there is nothing crucial in using these terms in these ways. What is important is the realization that explanations in the realm of the technological/technical require appeals to systems of varying complexity. The relevant system or systems constitutes an explanatory matrix that must be demarcated with care.[10] Depending on which system is invoked a different explanation may be required. In this sense the construction of a technical/technological explanation may be more of an art than a science. Indeed, it is a function of the skill of the person asked to give the explanation in determining what will satisfy the questioner.

11.5 Systems

Common to adequate answers for internal and external audiences, however, is the notion of a system. A system is a set of relations among other things, places, artifacts, social institutions, and individuals. "System" is a broader notion than a similar concept, "context". A context is a specific set of relations in a specific space and at a specific time, even if that time is actually an extended period, such as the Scientific Revolution. "System," however, denotes a more general and abstract set of relations which can be represented as schema's, for example, as line drawings showing various types of connections without there actually being such a system in physical existence. Thus, every context can be seen as a system, but not every system is a context.

---

[9]In some respects, this emphasis on the different scopes of the issues depending on the question asked, is reminiscent of Larry Bucciarelli's discussion in chapter 1 of his 1994 classic when he ruminates on the question "Do you know how your phone works?" He concludes, "I conjectured that there could be no unique criterion for judging responses; there could be as many legitimate, that is to say accurate, ways to describe how the telephone works as there are respondents." (p. 4)

[10]In some respects this idea that we offer an explanation in the context of a system resonates with Cummins (1975) proposal that functional explanations are offered against a set of background assumptions and tacit knowledge.

The system may be a simple one, such as the relationship between a person, the light switch, and the light bulb, or it may be more complicated involving systems within systems such as an electrical grid and the wiring of a house. Upon reflection we find that no aspect of an artifact is ever explained in isolation, it is always with respect to its relationship to something else and wherever there is such a relationship, real or conceptually imposed, there is a system. At its most basic, a system is a structured relationship between two or more parts. Further, I will argue that understanding that some artifact is itself a system or embedded in a system is essential to being able to offer or understand a technological explanation. That is, a good technological explanation relies on the idea that objects, persons and systems are related to one another in differing ways and proper placement in the appropriate system of the thing (broadly construed) to be explained is crucial to being able to understand the explanation as well as to formulate one. This is also why understanding who the audience is for an explanation is so important. The person offering the explanation must be able to refer to a system that will be understood by the person to whom he or she is offering the explanation. In this respect then and for our purposes, "system" will be considered a fundamental concept and the factors bearing on adequate technical explanations will rely on considerations of systems.[11]

However, it is not enough simply to appeal to a system, for there are two major problems. The first concerns individuating systems. The second concerns the kind of information provided in the answer. As alluded to above, when it comes to individuating systems the question is how do we determine which is the most appropriate system for our purposes in this specific instance? One way to approach this problem is to return to the idea of differing audiences and why- and how-questions. If we also introduce the idea of a feed-back loop we can begin to see how the appropriate system can be identified. The appropriate system will be a result of the audience asking for the explanation. To determine which audience, and how complicated a system we need to appeal to, we ask questions. Let us return to the light switch example. If the initial question is "Why did the light come on?" and the initial answer is "Because I flipped the switch" and if there is no further question forthcoming – then the explanation, for that audience, at that time, is complete. At this point we can, if we wish, assume further that the questioner is a member of an external audience not interested in the further workings of switches and electric grids.[12] But, if the questioning continues, we must begin to explore the kind of answer that would satisfy the questioner and reassess our initial assumption as to what kind of answer would suffice. It might come to mind that the person asking the question might not be a member of the

[11] To a certain extent I am working off of the ground breaking treatment of Thomas Hughes (1983).

[12] Of course individuals can be members of both internal and external audiences. Thus an electrical contractor could be satisfied by the "I flipped the switch" answer when that is really all he wanted to know at that time in that place. On the other hand, when taking a busman's holiday he might pursue further questions just because he wants to see if the folks who wired this house did anything different from what he would have done.

external audience, he or she might be a member of the internal audience and wishes to know the source of the electricity and how the grid is constructed. In short, it will take some effort to determine what kind of an answer will satisfy the person seeking an explanation. However, at some point the questioning has to end – this process is not the same one as the 4 year old child asking why the sky is blue, unwilling to settle for any answer. The practical end point in a technical explanation is marked by an indication of satisfaction by the individual asking for the explanation. But, it might be asked, does this not end in a tight circle: the explanation is sufficient when the audience is satisfied and the audience is satisfied when the explanation is sufficient. No, it is not circular for the explanation is adequate when the audience is satisfied and the audience is satisfied when they have no more questions. However, it may be the case that the audience is satisfied for the wrong reasons. This is the second problem. But before we turn to it, we need to look at a similar proposal to the one just proposed.

In his 1975 paper "Functional Analysis", Robert Cummins develops an account of a functional explanation in contradistinction to Hempel's (1965) DN theory. Although developed independently, there are many similarities between Cummins' account and some of what is being proposed here. Cummins, for example, wants to explain a biological function in terms of its contribution to the proper working of a biological system, where that is accounted for in terms of an organism's disposition to do A in circumstances C. He employs what he calls an "analytical strategy", by which he means, roughly, accounting for a given function by way of analyzing it into sub-functions. He also talks about the organization of the organism in terms of a "program". Thus, you can explain the function of a sub-function in terms of its contribution to the program of the organism, where the program is a description of what the organism is supposed to do or how it is supposed to behave.

Now, technological explanations have to account for more than functions, but Cummins' view supports, in many respects, the intuitions behind the systems account of technological explanation being developed here. It is especially interesting to see how he handles the tension between seeking ever-finer sub-functions and more comprehensive programs. He offers no systematic answer to the problem of when to stop seeking finer details and when to rely on the "sophistication" of the program. It seems to be a balancing act that depends on the degree of acceptance of the given explanation, a process very similar to the situation in systems explanation for technologies, that is how to arrive at the stopping point in asking questions. (Cummins 1975, pp. 760–762)

How an answer is phrased also can make a difference with respect to its degree of acceptance. We are all familiar with the advertisers' gimmick of making its product appeal more desirable by using such phrases as "scientifically proven" or "as shown in a scientific study at a major university". According to Dennis Carlat in the June 2008 issue of *Wired*, a recent study at Yale University, reported that spurious explanations were deemed more satisfactory when preceded by the phrase "Brain scans indicate...". That suggests that appealing to the audiences valued set of experts or expertise can bring about satisfaction with an explanation without actually achieving a state of genuine intellectual understanding; satisfaction here is a psychological as opposed to an epistemic state. However, this is no more a problem that we face when investigators use false statistics or appeal to made-up data. Sometimes researchers lie. Sometimes fathers, tiring of the constant "why?" of their child, will simply make up an answer that the father knows will satisfy them. Sometimes a person attempting to provide a technical explanation will throw in an appeal to something he thinks might satisfy his audience, knowing that it is actually misleading. These things happen – but, as in detecting scientific fraud, we must constantly be on guard. That these things happen does not discredit the proposed account any more than lying scientists discredit all scientific investigations.

11.5.1 System and Design

Returning to the issue of explaining artifact specific issues we can see further the importance of system. Consider some of the factors involved in explaining the design of an artifact. Artifacts do not take the shapes they have by accident; they are designed with specific factors in mind. Some such factors include how the artifact is to be used – so one factor in design concerns how the artifact relates to a broader system of use – in its use it will interact with other artifacts, or the natural world, itself a system of systems. Another factor involves the marketability of a product – attractive and user friendly designs sell – here part of the system is actually outside the design process itself insofar as it is the system of sales and consumers, an Aristotelian final cause as it were (see Bucciarelli 1994). Another consideration involves the cost, availability, and reliability of materials, factors again involving appeal to a broader system and may in turn relate to marketability and use, which themselves require appeal to systems.

At this point we are in a position to answer a possible objection to the account being proposed. It has been suggested here that the determination of the adequacy of an explanation, by way of answering why- and how-questions, is a direct response to the audience asking the questions and our ability to relate the answer to appropriate systems. That, however, it might be argued, fails to distinguish an explanation from an adequate explanation. But we have already noted that satisfying answers to why-questions direct the questioner to the manner in which the artifact functions in a system, relating to other artifacts and other components of the system. So, if the question is "how do I get the light to come on?" directing the questioner to begin by lighting a votive candle before throwing the switch would not count as an adequate answer since there is no way to relate such an action, lighting a candle, to the electrical system in an satisfactorily explanatory fashion.

11.5.2 System and Function

Function is another system intrinsic feature of an artifact.[13] The function of an artifact can only be fully explained in terms of how it fits into a system.[14] It does not matter if we are talking about function as use or function as in how it works, we cannot escape appealing to a broader context. Even if we are only concerned with the mechanics of the artifact we still are employing a system for the explanation since the artifact itself is a system if it has more than one part. The larger question of why it does what it does must involve an appeal to something beyond the artifact itself. Consider an example, the internal combustion engine (ICE). We can ask many questions regarding the ICE. What does it do? If the answer is "it performs mechanical work," we have a clearly unsatisfactory answer. We can pursue satisfaction by moving in one of two directions or both. If we are a member of the internal audience, we can ask how it transforms energy, seeking a description of the mechanics of the artifact. If we are a member of the external audience, "Performing mechanical work" by itself doesn't tell us anything of value. So, let's change the question. What is the ICE used for? "It is used to propel a tractor." Now we are getting somewhere, but we have also appealed to an incipient system or two. The first system is that created by the relation between the ICE and the tractor. The second is the relationship between the tractor and something else. The tractor is a system of parts that are brought together to produce an artifact that can, among other things, plow a field. If you didn't know *anything* else you wouldn't see that as an explanation. But if you do in fact know something about agricultural production, and if you can begin to see the utility of the artifact in terms of producing a crop, harvesting the crop, getting the crop to market, (system of transportation) selling the crop (economic system), having it transformed into something usable as food, etc., then the explanation of the function of the ICE as a mechanism for propelling a tractor begins to make sense. A fully adequate answer requires also knowing how the IBE produces energy and what it is used for, both of which requires appeals to systems.

11.5.3 System and Structure

Structure, the final artifact specific issue to be discussed here, is also closely aligned with the design of an artifact, and an explanation of why an artifact has the structure

---

[13] The literature on functions is almost as enormous as that for explanation. See for example, Kroes 2001, Vermaas and Garbacz 2009, Vincenti 1990 and Wimsatt 1980, 2002 for the merest sampling. I am relying here on a common sense appreciation of what a function is, recognizing that the circumstances in which we appeal to "function" may in fact change the meaning of the term. Thus, asking for the function of a turn indicator on the steering column of an automobile is not same as asking how well a device functions.

[14] Although the account here was developed independently of Cummins (1975), the two approaches agree strongly on this point.

it does also requires an appeal to systems. A design is a design of some structure or other. We can speak of the internal or external structure of an artifact. For example, in designing a modern skyscraper, it is common to configure the building around a central service utility column containing elevators, power and water lines, etc. What the external structure looks like is a function of many factors, including location, building codes, materials, cost, aesthetics and, last but not least, the architect's ego and the desires of the individual or group commissioning the design. Each of the factors listed are themselves systems – different kinds of systems, to be sure, but systems none-the-less. In seeing how the structural and functional designs are related we can see a number of things. First, if we simply looked at the external structure of the building, we would not learn very much about it. Looking at the external structure only reveals the aesthetics of its design and how the building relates to its surroundings, if we are provided with a site plan. If we took a political approach we might speculate on the kind of power statement the building makes given its size. However, to appreciate the design fully, we need to know more than how the building looks and whether it fits in its location, we need to know how the structural and functional components work together. We need to know, for example, how people who work on the top floors are going to get there. This is where we appeal to the internal structure of the building – the role of the central core in proving the elevator conduit, along with the means for getting electricity to the rest of the building.

We can now turn to an account of what it means to say that the three artifact specific issues we have been considering, design, structure and function of an artifact, can be explained in terms of systems. And here we can clearly differentiate between conceiving of the building as a system and conceiving of it in its context. From a systemic point of view, we don't need to think about a particular building, we can do this in generic terms. We can, for instance, draw a large three-dimensional rectangle on a piece of paper. Then we can add some stick figure trees, to give some sense of its size, and perhaps dot in some other buildings to show how it its supposed to fit in a location. Then we turn to the building itself and pencil in the central core – indicating where the electric lines will go, how the plumbing can work – schematic floor plans. But if it is a specific building I am working on, then we need to know exactly how high it will be, where in the city it will be located and what putting it there will do to the area. We also need to know about the location infrastructure, does it have enough water for the building, how will the electricity get to the building, parking. When we turn to the internal structure of the building, it is not enough to say that each floor will contain offices, elevators, rest rooms, we need to show where these will be located, how large they will be, traffic patterns, and so forth. Here some social factors come into play as we have to decide, for example, between open floor plan or individual offices. In that case, how many corner offices and how many windows? It might be the case that the external design of the building does not allow for corner offices as the corners are structural. Here the architect needs to be sensitive to office politics and human psychology.

As already noted, an explanation is a response to, among other questions, how- and why-questions. Thus, why does this artifact have this design? The answer will be in terms of the systemic factors noted above. If we are talking about automobile designs, one possible answer could appeal to aesthetic fads, or the interchangeability of parts in different models – always to something external to the artifact itself. Likewise for function: an artifact does what it does in order to contribute to the successful operation of a system in which it has a role to play. To explain the structure of an artifact, we appeal to a variety of systemic factors that contributed to the artifact having the structure it does. The explanation consists in positioning the artifact with respect to the system in such a way as to answer the why question satisfactorily. Thus, if we are explaining why all (so it seems) so-called cross-over vehicles tend not only to look alike, but are equally ugly, we might refer to the attitude of social superiority owners of such vehicles cultivate. Thus: "we don't have to drive big gas consuming monsters like Chevrolet Suburbans or Ford Expeditions or Land Rovers to have four wheel drive – but to let you know that, we purchase cars so ugly you can't ignore them" – perverse, but once placed in a social environment, the ugliness of these vehicles can be explained in a satisfactory fashion.

So far we have been concerned with the explanations of what, for lack of a better term, we can call "hard" artifacts, automobiles, buildings, hammers. We have argued that explanations of the design, structure, and function of hard artifacts involve appeals to systems in which they are embedded in some form or other. These explanations are, minimally, answers to how- and why-questions. But, there are other kinds of questions that can be asked regarding hard artifacts. One disarmingly simple one is "What does it do?" To answer this question we also need to place the artifact in a system. Explaining what the artifact does involves relating one part of the system to another. Thus, flipping on an ordinary light switch connects the power grid to the light bulb. When explaining what an automobile does we immediately appeal to the broader system of the social and natural world when we say that it is a means of transporting people and goods from one place (which must be defined in the context of a system) to another. When we explain what a skyscraper does, we turn the building itself into a system in which people work, work which may involve communicating (via different systems) with people in other places (system required) in order to move goods and services around the world, thereby appealing to the broader social system of transnational commerce.

11.6 The Social

The explanation of hard artifacts increasingly involves appeal to the social domain as we get farther and farther away from explaining the mechanics of the artifact, i.e., how it does what it does, to an internal audience of, for example, engineers and respond to the concerns of an external audience asking about its impact on the society. But to understand the social in its explanatory mode is to understand it as a system, or a set of systems. Examples of social systems used to explain features of hard technical artifacts include economic markets, communication systems, legal systems, building codes, and standardized metrics. But there are also non-social systems we appeal to by way of explanation such as the environment. Thus the new movement called green architecture is a response to increasing awareness of the effects of the built environment on the ecology of the planet. Such explanations only make sense in the context of thinking of the ecology of the planet as itself a system, something we seem increasingly to do.

In addition to hard technical artifacts we must consider the nature of explanations in the context of various social technologies, i.e., so-called soft artifacts. Social technologies help human beings arrange their affairs. We often explain what people do by appeal to some feature of a social technology. Thus, a response to the question "Why did he slow down?" could be "Because the speed limit was lowered." Implicitly this is an explanation by way of appeal to the regulative powers of the law – one's behavior is often shaped by legal constraints. Legal systems are developed in order to provide orderly means for adjudicating conflict and for shaping social behavior – they are deliberate constructions that have a function.[15] To explain these functions we must appeal to the needs of a different system: society.

To call society a system is obviously problematic. Which society do we mean? Sometimes society is synonymous with nation state, sometimes with a broader historical/cultural group, e.g., Western society, sometimes with more specific groups such as religions, sometimes with geographic locations such as Southeast Asia. Then there are even smaller groupings such as the Mafia, a city, a sport. As international trade becomes increasing global in its structure and interdependence, we just may end up at some point meaning the world, which tells us very little. Thus, explanations that appeal to society must be society specific or run the risk of being empty.

Not all forms of societal control of behavior are the result of deliberate constructions such as a legal system. There are cultural constraints that may not be obviously derived from a system. For example, it is often difficult to understand economic factors, economic theories not withstanding, that influence behavior as occurring within a system – for there may be conflicting systems at work, such as, at least in the United States, economic security versus public service. This is where there is a need for careful work identifying conditions that can be used to isolate systematic factors to explain how social technologies themselves can play an explanatory role.

11.7 Explaining Failures

We now turn to a different issue that calls for technical explanation: failure. Technical failures occur at all levels, from O-rings on a space launch vehicle to voting machines to power grids to social services systems to education systems. The

[15]There is an air of paradox to the claim that legal systems are deliberate constructions when one considers what are often referred to as "common law" systems for these have developed over time and generally in a piecemeal fashion. However, it is enough for our purposes to point out that they can't have developed at all if the idea of a law was not generally understood. Once in place, the society could add laws as it sees need. Granted this is different from constructing a legal system and then imposing it on a society, but common to both is the acceptance of the idea of law.

analysis and explanation of failure comes under the general heading of *forensics*. Explanations of hard technical artifacts involve *engineering forensics.* Engineering Forensics can be seen as a form of reverse engineering, during which process a mechanism is taken apart to see what makes it work. In the specific case of engineering forensics, an investigation into an incident in which some artifact, mechanism, or system failed is undertaken with the objective of determining if an artifact or mechanism was the cause of failure and what specifically went wrong. Failures of social systems and social technologies require *social forensics*. Social Forensics go beyond engineering forensics in that these investigations examine the failure in the context of a social situation, looking for human failures and system failures. The point of engineering and social forensics is to explain why the artifact or the social system failed to do what it was designed or evolved to do. Hard artifact failure usually comes as a surprise and sometimes is accompanied by social misery or even disaster. The example of the failure of the O-rings in the launch vehicle of the shuttle Challenger gives us a good case of an explanation that deals with a particular part of a complex artifact system that didn't do the job it was expected to do despite the fact that the demands placed on it exceeded its specifications. It might be argued that the failure occurred not so much because of the failure of the artifact system, but because of the social system in which it functioned. The warnings of the engineers that the O-rings would fail were overruled for other, some say political (Vaughn 1996), reasons. This is a valuable example because it exhibits the complexity of the relationship between hard artifact systems and social systems and between engineering forensics and social forensics. The Rogers Commission Report (1986), the official report of the US government's investigation into the incident, identified the specific cause of the disaster as the failure of the O-rings, but also noted that there were other systemic failures to be considered as well. The failure of the O-rings was a structural failure in the sense that this part did not work in this structure under certain circumstances, circumstances outside the specifications for the part in question. It might be argued that this is a functional explanation. But that is not clear. For when we appeal to the circumstances outside the specifications for the O-ring, we are appealing to more than the function of the O-ring. The failure of the shuttle mission as a result of the failure of the O-ring can also be explained in terms of the failure of the social system that developed, managed and used the shuttle system. What we need to guard against here is making the social system the explanation for everything that went wrong – that way lies the rhetoric of social criticism, but it does little to explain the failure in a way that leads to corrections that actually make a difference. In short, appealing only to the social has the effect of black-boxing the engineering.

11.7.1 The Challenger Example

In the discussion of the Challenger disaster, the question of the responsibilities of the engineers involved is often raised. Is it enough for engineers to design artifacts that meet the specifications of the client or do they have further responsibilities to determine if those specifications are reasonable given the functions of the ultimate product? To raise such issues also allows a distinction between seeking causes for failure and assessing blame. On the one hand we want to know what broke and why. On the other hand, we want to know who was responsible for the situation. In the first case we are looking for technical explanations, in the second, for judicial or even moral ones. It may be the case that an ultimate resolution of the entire situation requires assessing blame and putting constraints in place to correct for whatever actions or inactions occurred leading to the problem. But it is not clear that doing so is a necessary part of a technical explanation, construed as an outcome of engineering forensics. One might argue, however, that it is a necessary end point for social forensics, for it is not enough to know why the system failed, but what needs to be done to fix it and sometimes that means identifying individuals as failing to be responsible, instituting review processes, laying down what amount to moral directives, etc.

11.7.2 The 2000 US Presidential Election Example

It can also be argued that we can explain the results of the 2000 United States Presidential election in terms of both a hard artifact technical failure and a set of failures of the social technologies in which it was embedded. The so-called hanging chads were a result of the failure of the voting machines in Florida to operate correctly. The resulting decisions by state officials and finally the United States Supreme Court can be seen as failures of the social and political systems. Laws were broken but no action was taken by state officials or even federal officials to identify responsible individuals and have them prosecuted, given overwhelming political ambition and arrogance. In the second case, we have the politics of judicial appointment at the highest level trumping legal precedent and the procedures specified by the US Constitution.

11.7.3 The Ladbroke Grove Railroad Crash Example

Finally we should look at a different kind of failure case, also leading to disaster. This is one in which political agendas did not overtly appear to play major roles, but compounded instances of human negligence led to a sad result and, it can be argued, there was no specific technical failure. The case in question is the Ladbroke Grove rail crash outside of London, England on 5 October 1999. At approximately 8 am two trains traveling at high speeds collided, resulting in 31 deaths and 523 other casualties. The immediate cause was identified by an official investigation as the failure of the engineer of one of the trains to obey a stop signal. Subsequent study, however, revealed a more complicated story. The operator of the train that missed the red signal had only been on the job for 2 months. The signal itself was non-standard – structured as a reversed “L.” The red signal was located to the left of the other lights, rather than at the bottom of a standard three light array. Further, the signal was obscured from view by overhead roads and bright sun. Finally, those responsible for the maintenance of the track and its signals had failed to take necessary action to correct the problems at this site, despite the fact that there had been eight instances of trains missing the red signal there in the previous 6 years – luckily all those trains managed to stop before an accident occurred. The frequency of such occurrences should have alerted someone to a problem. While the signal had worked properly, there apparently had been a series of events that contributed to the final tragedy that no one had bothered to connect. Yes, there was system failure, but not out of blind ambition or greed or stupidity. In this case it seems that a series of small changes over time resulted in a situation that no one had anticipated, despite a number of warning signs over the years. However, in this case, the official investigation lead by Lord Cullen resulted in record fines, suggesting that there was a clear determination of blame.

11.8 Conclusions and Objections

It may be objected that this account of technological explanation is inadequate; it is simply too soft, resting as it does on interpersonal skills rather than rigorous logical connections. As we have seen, the use of iterated why-questions provides a method for both individuating systems and locating the appropriate explanation. That is, by finding out what kind of framework will provide understanding on behalf of an inquirer, we insure that the technological/technical explanation actually answers the question and, as we have suggested, answering how- and why-questions is what explanation is all about. This also helps with regard to another, so far unexplored, issue: how to determine what exactly the question is. Earlier it was noted that the questioning develops in the manner of a feedback loop. The point here is as much to find out what the questioner is actually asking as it is to arrive at a satisfactory answer. In fact, until both the questioner and the person seeking to provide the explanation know what the question is no satisfactory explanation will be forthcoming.

11.9 Failure, Success, and Symmetry

There is yet one final objection, (for the time being): the apparent asymmetry between the technical explanation of why something works and why something fails. On the surface it appears as if the explanation for why something works is that it accords with general processes that we understand. In the case of why something fails, we appeal to the particulars of the case. There are two different issues here. Again we see the unwelcome influences of older discussions on our expectations for present concerns. In particular, we see the continuing influence of Carl Hempel. Hempel laid down the condition of symmetry for the logical structures of both explanation and prediction. That is, in DN the same logical form characterizes both explanations and predictions; they are both deductive arguments with their premises containing a law and statements of initial conditions, leading to a statement of fact *x*, (if one is explaining single facts) or the statement of a prediction that *x* will occur. If one takes, by way of analogy, prediction to apply to why things work and explanation to apply to failures, there might be something to work with. That is, it might appear that the objection to the asymmetry in technical explanations has some ground. This assumes that the objections raised earlier to the DN model don't apply to this question of symmetry. But they do – for in DN explanations we need to know which laws to use, and we cannot assume that there is only one theory in use at a time. If the problems produced by issues of individuating laws and competing theories are real problems, then DN fails in deeply serious ways, and the call for symmetry of logical structures would seem to fail as well. In short, it is not clear that the explanation of why something works must have the same logical structure as an explanation of why something fails to work.

Further, it is rarely the case that an explanation of why something works needs to appeal to the expertise of an operator or the coherence of a system. Those are taken for granted when we do a walk-through of an artifact. Likewise, the presumption from the start is that the parts are properly designed and manufactured. Thus when we explain why the light goes on in the house when I flip the switch, it is in the context of a very big *ceteris paribus* clause. All things being equal, if the objects are properly designed and manufactured and if the operators are properly trained and competent and conscientious, and if the system as a whole works well, then when I do *X*, I can expect *Y*.

The more difficult explanations regard the failures, for explaining failure requires that we figure out which components of the *ceteris paribus* we ought to have thought twice about and what it was in those components that needs fixing. In short, there is a built in *asymmetry* in technical explanations and that is not necessarily a bad thing; consider what would happen if this were not the case.

If we failed to operate with the very large *ceteris paribus* clause bracketing our explanations of why things work, we would have to build in all the constraints it is designed to wash over and the end result would be stasis. Thus, let us consider our light switch. The light comes on in the room when I flip the switch only when the switch is properly manufactured and wired, but we know electricians make mistakes in wiring things, and we know that sometimes materials used in manufacturing are flawed, so the switch, even if properly constructed, may not work properly. But we also know that we are unjustified in assuming that the switch is properly constructed because the operators at the manufacturing facility that make the switch also make mistakes, sometimes because they are improperly trained, or because they had a fight with their husbands or wives, or because they had two beers at lunch. Knowing all that can go wrong just with the switch, I will be unable to explain why the light comes on because there is too much that can go wrong in the system as a whole. An explanation with all those caveats is no explanation.

On the other hand, digging into those assumptions the *ceteris paribus* clause hides when we turn to explaining failure is just what makes such explanations so difficult to construct but also so valuable. Here we have to uncover the very things we needed to ignore before. Not because they are necessarily part of the final explanation, but because they might be, and to ignore those possibilities is to offer only a surface explanation, which, again, is no explanation at all.

In conclusion, the asymmetry between explaining how things work and why they fail is essential to providing technical explanations. In the case of explaining success we need to simplify, and in the case of explaining failure we must get very complicated. This is not to say that explanations of why things work can't be complicated – they certainty can be. The point here is that the kinds of complications are of a different order when we try to explain why things didn't work as we expected them to. The complications have to do with the human factors, individual or aggregated. This also explains why we don't blame the artifacts when things go wrong, we blame the people who use or misuse or abuse them. Thus, it is true that guns don't kill, people do.

References

Bucciarelli, L.L. 1994. *Designing Engineers*. Cambridge, MA: MIT Press.

Carlat, D. 2008. "Mind Readers". Wired, June 2008. 120–128.

Cummins, R. 1975. "Functional Analysis". *The Journal of Philosophy* 72, 741–765.

Elster, J. 1983. *Explaining Technical Change*. Cambridge, UK: Cambridge University Press.

Hempel, C. 1948. "Studies in the Logic of Explanation". *Philosophy of Science* 15, 135–175.

Hemple, C. 1965. "Studies in the Logic of Functional Explanation". *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York, NY: Wiley.

Hughes, T. 1983. *Networks of Power: Electrification in Western Society, 1880–1930*. Baltimore, MD: Johns Hopkins University Press.

Kroes, P. 1998. "Technological Explanations: The relation between structure.and function of technological objects." Society for Philosophy and Technology, 3, http:scholar.lib.vt.edu/ejournals/SPT/v3n3/KROES.html.

Kroes, P. 2001. "Technical Functions as Dispositions: A Critical Assessment". Techne, Journal of the Society for Philosophy and Technology, 5, http:scholar.lib.vt.edu/ejournals/SPT/v5n3/KROES.html.

Pitt, J.C. 1988. *Theories of Explanation*. New York, NY: Oxford University Press.

Pitt, J.C. 2000. *Thinking About Technology*. Originally published by Seven Bridges Press, New York, NY, now at www.phil.vt.edu/HTML/people/pittjoseph.htm.

Ridder, J. de. 2007. *Reconstructing Design, Explaining Artifacts: Philosophical Reflections on the Design and Explanation of Technical Artifacts,* Simon Stevin Series in the Philosophy of Technology, Vol. 4 Ph.D. Dissertation, Delft University of Technology, Department of Philosophy

Rogers Commission report. 1986. *Report of the Presidential Commission on the Space Shuttle Challenger Accident*.

Salmon, W. 1989. "Forty Years of Explanation". In Kitcher P. and Salmon W.C. eds., *Scientific Explanation*. Minnesota Studies in Philosophy of Science, Vol. 13. Minneapolis: University of Minnesota Press.

Shew, A. 2007. *Beaver Dams, Spider Webs, and the Sticky Wicket: An Investigation into What Counts as Technology and What Counts as Knowledge*. MS Thesis, Science and Technology Studies Graduate Program, Virginia Polytechnic Institute and State University.

Vaughan, D. 1996. *The Challenger Launch Decision: Risky Technology, Culture and Deviance at NASA*. Chicago, IL: University of Chicago Press.

Vermaas, P. and Garbacz, P. 2009. “Functional Decomposition and the Function Part-Whole Relationship in Engineering”. In Meijers A. Editor-in-Chief. *The Handbook of the Philosophy of Technology and the Engineering Science*. Berlin, Springer.
Vincenti, W. 1990. *What Engineers Know and How They Know It: Analytical Studies from Aeronautical History*. Baltimore, MD: Johns Hopkins University Press.
Wimsatt, W.C. 1980. “Teleology and the logical structure of function Statements”. Studies in the History and Philosophy of Science, 3: 1–80.
Wimsatt, W.C. 2002. “Functional Organization, Analogy, and Inference”. In Ariew A. ed., Functions: New Essays in the Philosophy of Psychology and Biology, Oxford, NY: Oxford University Press, 173–221.
Winner, L. 1986. *The Whale and the Reactor*. Chicago, IL: University of Chicago Press.

Part III
Design and Engineering

Chapter 12
Successful Design in Engineering and Architecture

Design is at the heart of both Architecture and Engineering. However, the factors that bear on design decisions and outcomes differ for architects and engineers. For engineers, the design of an artifact or a system is approached with questions of utility and efficiency foremost in mind. For the architect, function and aesthetics take center stage, with aesthetics sometimes overriding the ability of the object to perform its function. These differences raise questions of how to determine when a design is successful. Despite the fact that issues concerning architectural design are more prominent in the popular imagination than design issues in engineering, I will argue that the question of what constitutes successful design is especially troubling in architecture. The difficulty of establishing criteria for evaluating the success of a design raises questions as to the very meaning of "design" in the context of architecture, if not in the context of engineering. In what follows, I will explore the ramifications of understanding the meaning of "success" in two different ways, one appropriate to engineers and the other for architects. I will suggest that this is a difference that makes a difference in our understanding of creativity.[1] I will argue that the genuine mark of creativity is to be found when one is forced to operate within constraints. As I will show, this means that rather than turn to artists and architects for an understanding of creativity, we should look to engineers.

12.1 Engineering Design

Let's start with some arguably debatable examples of successful design:

- The George Washington Bridge in New York City
- The Coliseum in Rome, Italy

Originally appeared in Creativity: Technology and the Arts. Edited by Hans-Joachim Braun. Peter Lang, Frankfurt/Main. 2010

[1]With apologies to William James, "There can be no difference that doesn't make a difference elsewhere – no difference in abstract truth that doesn't express itself in a difference in concrete fact and in conduct consequent upon that fact, imposed on somebody, somehow, somewhere, and somewhen." (James 1907, 1942, p. 45)

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_12,

- The American Minivan
- The Guggenheim Museum in New York City
- The U.S. Space Shuttle
- The National Gallery East Wing in Washington D.C.

What makes the designs successful? Of course, it depends on what we mean by "successful". But prior to an analysis of "success", we need some distinctions and some definitions. Let us begin with engineering design.

Engineers design many kinds of things, from pens to space shuttles to sewage systems. As a starting point, I would propose that, for an engineering design to be successful, the design at least must result in a useful product for some audience. That sounds easy enough and reasonably intuitive. However, it hides a crucial distinction nicely elaborated by Kees Dorst (1997) between design methodology and design practice, or design theory and designing as a real world practice. Design theory in engineering is a well-developed area and engineering textbooks are full of elaborate multicolor flow charts with feedback loops. (cf Bucciarelli 1994, Dorst 1997, Gasson 1974, Vincenti 1990, Callister 2003, Douglas 1988, Manogon 1999). But since I am interested in *really successful* design, I am going to look at design *practice*.

How does it come to pass that a successful engineering design results in a useful product? To understand how this happens we need to look at the full process that takes us from an idea to a product. By so doing, we can distinguish a number of components or steps such as the beginning and the end. Thus, the process begins at the point where there is a perceived need or opportunity. For example, I need a tool to take care of this problem, such as a screwdriver with a shaft that is bent at a 45° angle for access to difficult spots. In the case of opportunity, if we made a clever new device that did Y we could probably make some money, for example, high definition TV.

The endpoint of the process is located in consumer satisfaction. When we speak of the design process in engineering, there are also several other parts of the puzzle with which it should not be confused.[2] "Success" can be claimed at any of several stages in the long road from the conception of an idea to consumer satisfaction. First, there is what we will call the *conceptual design process*, in which the need or perceived opportunity is turned into a conceptualization of a product. This is a complicated process because it is not clearly obvious that there is only one way to conceptualize the goal of satisfying the need or optimizing the opportunity. For example, to get to the moon should we build a single-stage booster or a multi-stage system? Guiding the decision-making are considerations of efficiency as well as external constraints such as size, expense, availability of materials and manufacturing means. Success at this stage means nothing more than agreeing that this or that compromise is the best we can do under the current circumstances.

[2]For a detailed discussion of the variety of factors that can come into play in design in the real world see L.L. Bucciarelli's Designing Engineers, (Bucciarelli 1994).

The conceptual design of a product is to be distinguished from the *production process*, in which the envisioned product is actually manufactured. However, it should not be assumed that what it will take to turn the design into a product does not itself bear on the design. Further, it may be the case that the manufacturing process itself needs to be designed – but, irrespective, the designing of the product and its manufacture differ from one another and are, in turn, different components of a bigger process. Furthermore, whatever we choose to call this bigger process, the manufacture of the product is not where it ends. The product is manufactured in order to be used. In some case, such as a new kind of tennis racket, the object must be marketed, and one criterion of successful marketing is sales. Our new tennis racket may very well be the "best tennis racket ever manufactured", but if no one buys it is not an engineering success. By and large, sales can be a significant determinate of how well the object does what it is supposed to do.

Often, however, it is not. Not to be contentious, MacIntosh computers are clearly superior to PCs and yet they have a minor share of the market.[3] In those cases where sales numbers is not the criterion, say, for example, in the case of designing and building a bridge, success can be said to be a function of how well the bridge does its job. The Verranzzano Bridge in New York City does its job very well. The Tacoma/Narrows bridge did not – it fell down. (cf. Rosenberg and Vincenti 1978, Scott 1956) Thus we have at least these three components of engineering design: (1) the design process, (2) the manufacturing process, (3) functionality. For each of these components we can specify "success" criteria. The design process is successful if it results in something that suggests what the thing is supposed to do and perhaps even points in directions for its manufacture. The manufacturing process is successful if the result is a quality product that can be marketed and sold. The product is successful if it has functionality, i.e., if it does the job it is supposed to do. An engineered product that does not work is not a success. Clearly more needs to be said on each of these, but, hopefully, these rough distinctions give us a starting point.

12.2 Architectural Design

When we turn to architectural design we face a different activity. For one thing, it would not be correct to say of an architecturally designed object that if it does not do what it was supposed to do it is not a success. To put it baldly, the job of making an architecture design work is often up to the engineer or the builder. By way of example, consider the design of the main Duke University campus. In Durham, North Carolina, USA. In 1966 Duke University embarked on the third attempt to make the neo-English looking campus buildings work. It is alleged that James B. Duke,

[3] For the purpose of this discussion, we ignore issues such as fad items, like pet rocks, whose "success" is more a function of clever advertising than performing some function to satisfy a genuinely perceived need.

the university's founder and major endower, ordered the architects to design a campus that looked like Oxford or Cambridge. The design they presented was strictly from the external point of view. It looked right. But inside the buildings was another story; classrooms were wider than long, or even in the shape of an L. Some rooms had pillars in the middle. No attention had been given to acoustics. Luckily, the three story stone buildings had two-foot thick exterior walls. Thus the redesign of the interior could proceed by completely gutting the building and starting over. This job was handed over to engineers.

Also, as when speaking of engineering design, when we turn to architecture, we need to realize that architects design many kinds of things, from simple buildings to pieces of furniture and even flatware. Therefore, for these purposes, let us restrict our attention to edifices – buildings, or complexes of buildings.

However, unlike in engineering, there do exist canonical criteria in architecture that have been used in the past used by architectural critics to evaluate the practical success of a building. That is a second crucial difference. The discussion concerning whether or not a building is successful usually takes places in the context of critical critique in a public forum such as a newspaper. Canonical criteria and critics not withstanding, there is much disagreement over what exactly the criteria mean and, especially since Venturi (1972) introduced the concept of the post-modern in architecture, the canonical criteria have been under fire.

The architectural canonical criteria come to us from the Roman architect, Vitruvius. The three criteria he laid down were, Utilitas, Firmitas, and Venustas. These have been translated to mean "Commodity, Firmness, and Delight" – Robert Bruegmann (1985) considers the exposition of these concepts by Geoffrey Scott in his 1914 *The Architecture of Humanism* to be the best, so let us look at what Scott has to say:

Commodity

> Buildings maybe judged by the success with which they supply practical ends they are designed to meet. Or, by a natural extension, we may judge them by the value of these ends themselves; that is to say, by the external purposes that they reflect. These, indeed, are two different questions. The last makes a moral reference, which the first avoids, but both spring, and spring inevitably, from the link which architecture has with life. (pp. 3–4)

On this account Commodity, or perhaps a more faithful translation is Utility, requires that the design of a building both be suited to the function it is supposed to perform and exhibit that function. The first seems reasonable enough – the second is a bit less obvious. Taken to extremes we might require that Post Offices look like a giant envelope and surely that is not what is entailed here. But it is not uncommon to expect, for example, governmental buildings to be larger than life, exhibiting the transcendent function of government over the interests of a single individual.

The second criterion is *Firmness*

> On every hand the study of architecture encounters physics, statics, and dynamics, suggesting, controlling, justifying its design. It is open to expression of material properties and material laws. Without these, architecture is impossible, its history unintelligible. And if, finding these everywhere paramount, we seek, in terms of material properties and material

laws, not merely to account for the history of architecture but to assess its value, the architecture will be judged by the exactness and sincerity with which it expresses constructive facts and conforms to constructive laws. (p. 2)

Bruegmann interprets this to mean that "Firmness. . . .is about structure and composition. A building should not only be sound and logical in is construction, but it should appear this way as well." (p. 18) It is not clear what it means for a building to be logical. Further, with the advent of newer construction materials and techniques, the appearance of the soundness of the construction has lost some of its force. Consider large enclosed sports stadiums. The supporting structure of the domes is often not clear and obvious. It is also not at all obvious that allowing the building to visually expose the source of its soundness necessarily is a good idea. This example is not completely on point, but it should highlight the issue. Corning Industries is a large U.S. firm specializing in products made from ceramics and class. When the Corning Plant in Christiansburg, Virginia, was built in the 1960s there was an expressed desire by management to use as many Corning materials in its construction as possible. So, some wise designer decided to use glass tubing for the plumbing and to have the tubes exposed overhead. When the plant was opened and tours were being given, the obvious mistake was noted and the tubes were quickly wrapped in ducktape.

The third criterion handed down to us by Vitruvius was *Venustas* or beauty or sometimes conceived as Delight.

We may trace in architecture a third and different factor – the disinterested desire for beauty. This desire does not, it is true, culminate here in a purely aesthetic result, for it has to deal with a concrete basis which is utilitarian. It is, nonetheless, a purely aesthetic impulse, an impulse distinct from all the others, which in architecture may simultaneously satisfy an impulse by virtue of which architecture becomes art. It is a separate instinct. It will borrow a suggestion from the laws of firmness or commodity; sometimes it will run counter to them, or be offended by the forms they would dictate. It has its own standard, and claims its own authority" (p. 4)

And therein lies the rub. What makes a building beautiful? Surely we want to resist the idea that beauty is simply in the eye of the beholder, but can we? Who is the arbiter of beauty? In what some call modernist architecture and then in post-modern architecture, the arbiter has become the architect. But there is a difference between the architect of the nineteenth century and the architect of the twentieth. The post-Enlightenment architect of the nineteenth century believed in the power of reason to reveal the nature of things. In this case, it was the nature of beauty. There was a deep-seated belief that there existed natural laws governing the beautiful and that the architect was best qualified to find those natural laws. That is, in dealing with this ineffable quality of beauty, while the modernist, nineteenth century architect took it upon himself to be the arbiter of taste, it was taste allegedly based on reason. As Bruegmann puts it,

Modernists believed the job of the architect, at least the genuine avant-garde architect, was to discover what these laws [of beauty] were and to insist on them even if they ran counter to society's expectations. In fact, as the nineteenth century progressed, the avant-garde moved further and further from the tastes of the population at large. (p. 22)

The search for and hoped for discovery of universal laws of beauty by the chosen few (i.e., avant-garde architects) was seriously under-minded in the 1970s by Robert Venturi who, thanks to his criterion for post-modern architecture, that the present must recapitulate the past, has helped spawn the ubiquitous large office buildings with various embellishments such as columns and arches that line the sides of the highway that leads from Dulles International Airport outside of Washington D.C. into the U.S. capital. With recapitulation of the past as the sole criterion, beauty becomes taste. And we all know *de gustibus non disputandum est.* Couple this with the architect's retained conviction, a holdover from the nineteenth century, that he or she is the arbiter of taste, this time, not based on reason but fad or ego, and you get the architectural plague of the current era.

And so, we ask, in this current milieu, what does it mean for a building to be successful? If the taste of the architect is the determinant and there are no other criteria to appeal to that make sense, then there is no criterion for success in contemporary architecture.

And yet that can't be correct. We continue to have schools of architecture and we continue to teach students how to design buildings, or is this mere pretence? Just as in engineering, there is a distinction to be drawn between the theory of design and its practice. It matters not what is taught in schools of architecture, the question before us is how to determine if an actual building is successful. What we have just seen is that the traditional criteria for evaluating the product have been undermined. They have been undermined by the development of new materials and techniques and by abandoning the nineteenth century modernist conviction that there are laws of nature governing beauty. Whatever criteria are provided have to do with the taste of the individual critic and that tells us very little about the building itself.

12.3 The Role of Creativity

So where does that leave us? We have multiple criteria of success for engineering design, but not very much of use for architectural design. Given this conclusion, it might be suggested that we have been concentrating on the wrong component of design. It well may be the case that our understanding of what constitutes successful design in architecture requires that we look to a feature fundamental in architecture, but not obviously as crucial in engineering: creativity. (cf. Gelernter 1998, Scully 1991) The difference between successful engineering design and successful architectural design, it may be suggested, is to be found in the creativity one finds in architecture. A successfully designed building is one that captures the imagination, it speaks to the values of the time and to our spirit. It is contemporary high art. By combining materials, function, and imagination the architect produces a statement whose purpose is to teach the viewer something. The problem here is, of course, whose imagination? Whose values? Whose spirit? Who has given the architect the role of teacher and why should we accept the architect as teacher?

There is, however, a deeper problem here. It is not a question of resisting artistic arrogance, it is the problem of the role of creativity in an increasingly technological world. By a technological world, I have in mind the popular image of the world of artifacts created by the application of formulae and the criterion of efficiency, dominated by the image of the heartless machine, i.e., the world, it is said, of the engineer. (cf Borgmann 1984, Feenberg 2002, Veseley 2004) The reason we have such a difficult time understanding success in architecture, it is argued, is that our current concept of success is dominated by the engineering efficiency model and of course that does not apply to architecture. The world of art, it may be argued, is increasingly irrelevant in an engineering dominated society. But, comes the response, just because we are so dominated by the engineering concepts of success, cashed out as they are in capitalist economic terms, that we need the artistic concept of creativity to humanize the world of efficient artifacts.

This is merely to rehearse an old and rather useless argument. Nor is it one that is profitable, for it relies on bogus distinctions, suggesting that architects and engineers are somehow different kinds of people, where "kinds" is be read in a fundamental sense, implying that they may even been different species. The fact of the matter is that architects are as much involved in the technological world as engineers. Architects work with materials to create artifacts. They are, together with engineers, humanity at work.[4] What we seek is a criterion by which to evaluate the quality of that work. Therefore, consider the following proposal: that we evaluate architectural output in terms of creativity. All we need is a definition and we can solve the problem.

To resolve our problem, let us turn the issue upside down by presenting a case that is basically counterintuitive. The case to be made here is this: in contemporary society, engineering projects are our best of example of creativity. That is, the most creative people in contemporary society are engineers. I would like to be able to say that of artists and architects but they resist standards of evaluation. What I want to do here is make the case for why standards of evaluation are essential to the very concept of creativity. The artist and the architect, and yes, even the *avant garde* in music, must realize the need for standards in order to do their work.

But, it might be argued, there is no need to make that case. The artist and the architect already have standards. Their role is to challenge the establishment. The contemporary standard is shock value. One has to be shocked in order to challenge one's own values and ultimately those of society. The role of high art is to be the conscience of society. But, I would reply, that is too simple. One can attempt to play the role of the conscience of society, whatever that may mean, in one of at least two ways. One can demean or one can uplift. It is easy to demean. And what does it accomplish? It is much harder to be uplifting and much more valuable, for in striving to be uplifting one strives for beauty. Yes, I want to argue that there is still a positive role for the concept of the beautiful in art and architecture. But to see that the case first needs to be made for the primacy of creativity in engineering.

---

[4] For an elaboration of this idea see my *Thinking About Technology*.

Some engineers are very clever. They manage to come up with most wonderful devices – big screen TVs, cars that do the driving for you, battery powered pepper grinders, ever faster computers with ever increasing memory, nano devices keep slacks from wrinkling, and the list goes one and on. But, is this creativity?

I will argue in the affirmative. But first I will propose a definition. Here is a first attempt: *To be creative is to produce variation given the constraints of the materials and other parameters within which you work.*

Although formulated independently, it turns that that this proto-definition echoes in crucial ways the 1996 proposal of Csikszentmihalyi,

> Creativity can best be understood as a confluence of three factors: a *domain*, which consists of a set of rules and practices; and *individual* who makes a novel variation in the contents of the domain; and a *field* composed of experts who act as gatekeepers to the domain and decide which novel variation is worth incorporating.

From Csikszentmihalyi's view, what is missing in my proposed definition is the role of the gatekeepers. However, not only is it implicit in my proposal, it is also more broadly construed – for in architecture, the gatekeepers include the people funding the project and they may have as little taste as the architect.

Csikszentmihalyi speaks of domains and I talk about constraints. What kinds of constraints? To begin with, there is the world. There are certain physical limits to the strength of materials. Second, there is the state of our knowledge. Third, there are economic limits, and limits to what is socially acceptable. Thus, given what we know and the resources at our command, can we design and produce a widget that will do X? Phrased that way, we see that our definition is not complete – something is missing. What I am about to suggest will be controversial for it is counter-intuitive. Any account of creativity must contain an appeal to goals. Consider the question above: given what we know and the resources at our command, can we design and produce a widget that will do X? This is what is missing from Csikszentmihalyi's account.[5] On his view, it appears that there is just random variation. But creative work is not merely variation for its own sake. Artists step outside accepted forms when they want to make a point about something in particular. That is, when they have a goal in mind. The inclusion of goals works across the artistic spectrum. For a composer: given what I know about music composition and the limitations of the instruments and the musicians, can I compose a piece that expresses Y? It works for an architect: Given what I know about materials, structures, buildings codes, the desires of my client, the objective the building, can I do this? But putting goals into an account of creativity seems counter-intuitive. One would expect to hear the following kind of objection: Isn't it the case that the creative spirit is most accurately characterized as unfettered? The more creative, the more unfettered – the creation is something that has somehow escaped conventions, the limits of the materials and other considerations. Is it not the case that

[5] There is a great deal more to Csikszentmihalyi's view than I have indicated and it is indeed worthy of explication and elaboration. I clearly cannot have done justice to his account in this short piece. However, this is not the occasion for that project.

to slip goals into our account of creativity is to put boundaries on it and hence is against the spirit of freedom implied by creativity? I will argue that this is not the case.

12.4 Creativity and Freedom

Creativity must be bound by limits, on this both Csikszentmihalyi and I agree, but I also insist on the limitations required by having goals. It is true that in being creative, the designer – using the term in its most generic sense here – produces something that transcends or transgresses some recognized boundary or other – but he or she does so for some purpose, if only to show what happens when that boundary is removed or overcome. Otherwise, it is not clear why she would challenge the boundary. To be creative then is to act so as to show something, to prove something, to make a point. As we shall see, it is also more than that. But unbound action is not creative – it may be free, but creativity demands more than freedom.

At this point it is clear that the definition requires reformulation. The definition originally proposed was:

> To be creative is to produce variation given the constraints of the materials and other parameters within which you work.

However, we do not need to do more in principle with the definition since the goal-oriented dimension of creativity is covered by the concept of "work". In *Thinking About Technology* (2000) I defined work as "the deliberate design and manufacture of the means to manipulate the environment to meet humanity's changing needs and goals." (pp. 30–31). Putting that account together with our proposed definition of creativity,

> To be creative is to produce variation given the constraints of the materials and other parameters within which you engage in the deliberate design and manufacture of the means to manipulate the environment to meet humanity's changing needs and goals.

Assuming that this is an acceptable definition, it should now become apparent why engineers are our most creative people.

12.5 Engineering and Creativity

Engineering is not one thing. It is not just that there are many different types of engineers, mechanical, aerospace, chemical, biological, etc., it is that there are different types of engineering activities. Among others, there are two in particular we should distinguish. Let us tentatively call them the *discovery activity* and the *application activity*.

During the discovery activity engineers engage in something akin to Kuhnian normal science. (Kuhn 1962) These are the activities engineers engage in to discover what the physical parameters are. Here is where we find the research that yields tables and formulae for calculating stress and resistance of materials. Here is where they develop the tests to see if this assay or that is at the correct level of purity. Here is where we find engineers creating the knowledge of the physical world that ends up making some of what engineers do at the applied level look like cook-book engineering: i.e., the application of tried and true methods and techniques for doing this or that. Because a lot of engineering amounts to little more than this sort of thing, it may be the basis for the appearance of a lack of creativity in engineering. But there is surely a lot of creative work, in just our sense, that goes into generating this knowledge which in many respects it is not like scientific knowledge. It is not accounting for why this or that does this or that, but rather discovering how *doing* this y*ields* that and then systematizing it in one fashion or another so that others can use that knowledge to do something else. It ends up saying something like: in order to do this, you have to make sure that the materials you are using meet these criteria under these circumstances. Generating this kind of knowledge is tedious work, but that is what engineering graduate students often spend a lot of time doing.

When we turn to the second kind activity alluded to above that we get to the heart of engineering creativity. In applying engineering knowledge to a design problem the engineer must lay out the parameters within which her or she will be working. Thus, if I use these materials, I will face these limitations. I can't use the materials I want to because of cost considerations. Further, the object has to be light enough to carry around all day. If I choose this configuration I will face these obstacles. Further, given Bucciarelli's work, the conception of the object or system varies depending on who is considering it – marketing has one idea, the designer another, etc. That is why Bucciarelli introduces the notion of an object world – a phrase that helps capture the multi-faceted aspects of design work; each perspective represents a different object world. And since design is rarely the work of one person, another set of limiting parameters come from the very existence of the other members of the design team. Thus, working within the limitations presented by the materials, the economic and social environments, and the other members of a team, the creation of an artifact can only be considered a triumph of creativity.

Now it might be objected that in the definition of "creativity" the reference to the production of variation might not seem like enough to warrant the characterization of being creative. For, the argument goes, building a slightly more efficient version of something already around is producing a variation, but it does not represent much in the way of being creative. This is true, however the definition was not intended to give an account that would assert that only the most creative activity counts as being creative. However, it is the case that the proposed definition *appears* to require a metric by which we can measure degrees of creativity in order to be complete. That, however, is not clear. We already have a metric and it is to be found in the account of what is a successful design. To introduce a variation is to be creative, but it may not be a successful variation. Not every act of creativity produces a work of genius. Something that looks like a good idea may not result in a successful product – for a variety of reasons consider the Edsel. So, while every successful design requires a creative act, not every creative act results in a successful design.

To return now to the claim that engineers are the most creative people around, let me begin by introducing a few caveats. First, not every engineer is involved in product design, and we are talking about design. Some engineers are primarily involved in number crunching – which is legitimate work, but it may not involve much creativity. In fact we may not want them to be creative since we expect them to come up with the facts, the physical limitations we need to know about before we can be successful. Second, creativity is not limited to engineers. I am, however, arguing that they are the most creative. We are now in a position to explain why that is the case. Basically it comes down to knowledge. The more you know about the physical, economic and social parameters, the fewer the genuine options for increasing variation. The more you know about the limits within which you must work and how those limits interact with each other, creating further limits, the fewer degrees of freedom there are. This is a case of the more you know, less appears possible. Therefore, given increasing limits on possible action, production of variation becomes more and more a valued act, just because it is so hard to accomplish.

However, it might be observed that this account also characterizes the architect. Isn't it the case that the more the architect knows about materials, the site of the proposed building, the intentions of the client, etc, the less freedom he or she has and the more impressive the design with variation? Yes and no. Here is the bottom line and it speaks volumes. The more successful architects today are part of large firms that offer complete service, meaning by that they are usually architecture and engineering firms combined. The architect produces the design and then the engineers have to figure out how to make it work. And if they can, then the creative credit should go to the engineers. It is easy enough to sketch a few ideas, but to actually bring those ideas to physical fruition in a way that works is a work of art.

Where does that leave us? We began by examining some criteria for success in design for engineers and architects and concluded that while we could articulate criteria of success for engineering, we could not for architecture. The reason for our failure with respect to architecture was that the traditional criteria for evaluation have been undermined. Further, it was argued that we need standards for evaluation and that it was not enough to pronounce a building a success or at least not merely on the basis of personal taste. In light of that, a different criterion was proposed: that we evaluate architectural success in terms of creativity. I further argued that we could make the case not only for the role of creativity in engineering, but for engineers being the more creative designers. The standard for creativity was shown to be variation based on knowledge of constraints, not freedom. So, if we now turn to architecture and apply that same criterion, what do we get? The successful architectural design is one that creatively reflects the constraints of the design situation, geographical, social, economic, and artistic.

12.6 Conclusion: Architectural Failures and Successes

Michael Graves' Portland building, on the criterion elaborated above is a failure.

To begin with, the building appears to disregard its surroundings. The tiny windows create a kind of visual dissonance with those of the buildings around it. It does not harmonize with its location; it just sort of sits there. That is, the building appears

Fig. 12.1 The Portland building

to be an impediment to moving around the area, whether that movement is visual or physical. It shares little architecturally with the surrounding edifices. It is an example of excessive variation. It is located in what would seem to be a square area that would otherwise be a park – and yet it seems to mock the idea that there might be a park here instead – it is a heavy building whose parts seem arbitrarily thrown together. It is a bully.

In contrast, consider Philip Johnson's AT&T building in New York.

It fits; it makes sense in its context. It differs from the other buildings, but it does not reject them or they it. It shows what else a building can be in that context without sneering at the context. Yes, there are differences between it and the other buildings with which it is situated. But the AT&T building appears to celebrate those differences, not ignore or dismiss them. Consider again the dual skyscrapers Graves designed in Den Haag, Netherlands. The result here is not clearly as successful. From the train, as you pass by them at a distance, they appear to be almost perfect. Tall, massive, and yet the exaggerated traditional Dutch rooflines that make their placement appear natural. Unfortunately, the buildings contribute to a kind of artificial demarcation of parts of the city, between the lived-in city and the governmental city that empties into the evenings. The governmental complex, of which they are a part, forms a clump in the middle of a vibrant part of the city that you have go around to get from one part of the lived-in city to another. The complex interferes in

Fig. 12.2 AT & T Building

the life of the city – it has negative and dark aspects to it. Since no one lives there, it is dark at night except for the glare of street lights, empty and brooding, even threatening. Thus, despite the pleasing visual effect from a distance, the actual impact of the builds appears to be negative.

In short, by example I have been attempting to lay out some criteria for architectural success. Clearly these efforts are but a very first and meagerly sketch of an idea. Much more needs to be done. But it is hoped these ideas are suggestive for further work. Variation is important, but not variation that negates everything else. Harmony is important, but not harmony to the point of boredom. Celebration of the site is crucial but not by way of degrading what else is already there. Visual excitement is important, but not to the point of contributing to an overall failure of the impact of the building.

In one sense both engineering and architecture share a common thread when it comes to understanding what constitutes successful design – the object has to work. What this means, however, differs between them. For the engineer the bottom line is does it sell? For the architect, I would argue that among other things, a successful building should make us want to see more, but not more of the same.

Acknowledgement My thanks to Carla Corbin and Dr. Thomas Staley for their assistance and to the participants of the Center for Science and Technology Studies at Virginia Tech for helpful feedback on an earlier version.

References

Borgmann, A. 1984. *Technology and the Character of Contemporary Life*. Chicago, IL: University of Chicago Press.

Bruegmann, R. 1985. "Utilitas, Firmitas, Venustas, and the Vox Populi" In Marder T.A. ed., *The Critical Edge: Controversy in Recent American Architecture*. Cambridge, MA: MIT Press, pp. 1–24.

Bucciarelli, L.L. 1994. *Designing Engineers.* Cambridge, MA: MIT Press.

Callister, W.D. Jr. 2003. *Materials Science & Engineering*: An Introduction, 6th Edition. New York, NY: Wiley.

Csikszentmihalyi, M. 1996. *Creativity: Flow and the Psychology of Discovery and Invention*. New York, NY: Harper-Collins

Dorst, K. 1997. *Describing Design; A Comparison of Paradigms, Doctoral Dissertation* at de Technische Universiteit Delft.

Douglas, J.M. 1988. *Conceptual Design of Chemical Processes* . New York, NY: McGraw-Hill.

Feenberg, A. 2002. *Transforming Technology, a Critical Theory*. New York, NY: Oxford University Press.

Gasson, P.C. 1974. *Theory of Design.* London: Batsford.

Gelernter, D. 1998. *Machine Beauty – Elegance and the Heart of Technology*. New York, NY: Basic Books

James, W. 1907, 1942. *Pragmatism.* New York, NY: Meridian Books.

Kuhn, T. 1962. *The Structure of Scientific Revolutions.* Chicago, IL: University of Chicago Press.

Manogon, P.L. 1999. *The Principles of Materials Selection for Engineering Design.* New York, NY: Prentice-Hall

Pitt, J.C. 2000. *Thinking About Technology.* New York, NY: Seven Bridges Press.

Rosenberg, N. and Vincenti, W.G. 1978. *The Britannia Bridge: The Generation and Diffusion of Technological Knowledge*. Cambridge, MA: MIT Press.

Scott, R. 1956. *In the Wake of Tacoma: Suspension Bridges*. Reston, VA: ASCE Press

Scully, V. 1991. *Architecture; The Natural and the Manmade*. New York, NY: St. Martin's Press

Venturi, R. 1972. *Learning from Las Vegas.* Cambridge, MA: MIT Press.

Veseley, D. 2004. *Architecture in the Age of Divided Representation: The Question of Creativity in the Shadow of Production*. Cambridge, MA: MIT Press

Vincenti, W.G. 1990. *What Engineers Know and How They Know It*. Baltimore, MD: Johns Hopkins University Press.

Chapter 13
Design Criteria in Architecture

13.1 Introduction

In Chapter 12 I contrasted criteria for successful design in architecture with that in engineering. I argued there, among other things, that with the advent of "postmodern historicism" in architecture, beginning in the 1970s with the work of Venturi, there ceased to be operative criteria to evaluate architectural design and I made a first step towards outlining what such criteria might look like in the current age. I suggested that:

– Variation is important, but not variation that negates everything else. The Pompidou Center in Paris is an example of this.
– Harmony is important, but not harmony to the point of boredom. An example of a harmonious but boring architectural creation is the Levittown type suburban housing development in the United States.

In this chapter I elaborate those ideas, contrasting them with traditional canonical criteria, and offer some additional criteria in an effort to capture this fundamental idea: that architectural design must strive to make architectural projects work in context, given their functions. In short, I will develop a design objective called "Common Sense Design", based in part on some of the suggestions William James makes in his 1907 *Lectures on Pragmatism*. In part this involves developing the idea that certain designs have managed to survive relative to the domain in which they were developed and that we should learn from them. This is an argument against universalist principles of design, focusing on not just the locality of the site, but, rather, on the insights we can glean from the indigenous culture. As an example I will end by considering the Michael Graves complex in The Hague again, which, from a distance, is a success, but, in context and in impact, appears, on one interpretation, to be a failure. Seen in another light, Graves' complex can be favorably compared to Frank Lloyd Wright's Guggenheim Museum.

Originally appeared in *Philosophy and Design*, edited by Pieter E. Vermass, Peter Kroes, Andrew Light and Steven A. Moore. Springer Academic Publishers. 2008. Reprinted by permission of Springer Academic Publishers

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_13,

13.2 Architectural Design and Philosophy of Technology

First, why the emphasis on architectural design? Or, more bluntly, what does architectural design have to do with the philosophy of technology?

To speak of living in a technological society is to speak of a society in which human activity seamlessly engages artifacts of one kind or another, from computers to houses to shuttles to legal systems, etc., in the processes of living and seeking a better life. Those artifacts are designed. Sometimes they are designed for one purpose and used for another, but they remain designed. Thus, at the heart of the concept of an artifact is the concept of design. And since the philosophy of technology is concerned in many ways with artifacts, many questions about architectural design can be seen to fall within its purview.

Put simply, architects design spaces as well as the constructional systems that enclose and mediate them. These are spaces that we use for living, working, recreation, etc. Sometimes they contribute significantly to achieving the goals we seek to accomplish in those spaces and sometimes they do not. Therefore, before we design the space we ought to have some criteria to guide our design. We need such criteria to maximize the probability that we will succeed in accomplishing the goal of creating a space that contributes positively to the activity for which that space was designed. These criteria should serve two purposes:

1. they should guide design, and
2. they should be the criteria by which we judge the success of the design.

To say this is not to commit to a vicious circle, i.e., we judge the finished product in terms of whether it meets the criteria we used to design it. It is more complicated than that because in the time line from initial concept to a design to finished product it is quite possible, in fact, I would argue, almost inevitable that the meanings of some or all of the criteria undergo subtle but important changes. That is, we may think we know what we mean by harmonious when we start the design process, but when we look at the finished space, it may not have turned out to be harmonious, in which case either we did not know what we meant by the concept when we began, or the concept of harmony we employ in evaluating the end space has changed from when we started and we now have two different interpretations of the same word. This can happen for a variety of reasons, but my explanation is that when we think of a concept like Harmony, given that it is part of our *criteria* for a successful space, we jump to the conclusion that as a criterion it must be universal and fixed in its meaning, when in fact there are no such fixed meanings.[1] To take this one step farther, I am willing to defend the view that in each application of, for example, the concept of harmony, we add to or subtract from what we

[1] This analysis is firmly related to Goodman's (1955) new problem of induction and his concept of projection.

thought we meant when we started the design process. Meanings change in application, or to put it in Peircean terms, meanings change when reality pushes against language.

13.3 James and Common Sense

Basically James' account of common sense claims that the categories of common sense thinking are historically contingent, certain categories emerge because employing them in that context at that time increase survivability and success, however defined (James 1907, 1981). What may be an example of something or other in one context may not be in another.

Consider the following story, a real life example. I had asked some friends from the university if they wanted to help my wife and I load hay bales that were out in the field onto a truck and then unload them into our barn, they, all Ph.D.s, agreed and thought this would great fun. The hay field in question is on a hill and reasonably steep and visible from the road that winds down into the valley below. While we were near the top of the hill I saw the pick-up truck of an old framer who lived down the valley stop and turn around and make its torturous way up the mountainside to where we were. My wife was driving the hay bale truck, she grew up on a farm but it was in the flatlands. I was up on the truck stacking the bales as they were tossed up onto the truck bed. The old farmer, Dan, got out of his pickup and stared at us and just shook his head. "How many Ph.D.s involved in this operation?" he asked. I replied there were six of us. He snorted and then he asked "Any of you ever heard of gravity?" and then he laughed and laughed, got in his truck and started back home continuing to shake his head. It seems we had the truck pointed uphill – and the guys tossing bales had to throw them uphill against the pull of gravity. It was much easier to throw them downhill onto the truck bed, getting an assist from gravity. He knew that instinctively, well, he grew up riding along side his daddy from the time he could walk, absorbing so much of the common sense knowledge of how to get things done on a farm that it seemed like instinct.

This is the sort of thing that James means by common sense. Through a variety of means, some ways of doing things in a certain place for a certain purpose come to be common sense as they share acceptance in the community that does not require justification, they have been *vindicated over time*. Yet, the old French saying, *Plus le change, plus le meme chose*, is false. Consider the same scene 20 years later. The hay field has been sold and the new owner no long makes the small "square bales" of hay, but he still makes hay. However, now he makes hay in huge round bales. In order to get them down off the hillside he has to transport them one at a time on a spike on the back of his tractor, and to load one of the round bails on a spike on a hill you have back the tractor uphill, spike end pointing up so you can impale the bail and then move down the hill without it falling off. The use of gravity has changed.

13.4 A Common Sense Proposal

Architects design spaces, but not all spaces are designed. As an undesigned space consider a forest, although there are designed forests in the Netherlands, France and elsewhere. Furthermore, spaces are always to be found in other spaces. And it is to the spaces within spaces I would direct our attention. I am not concerned with questions of the intention of the designer, for his or her intentions have their own problems. Instead I want to focus on the space itself. If spaces are always to be found in spaces, then the relationship between and among spaces seems a logical starting point for a new discussion of design criteria. I should also like to note that spaces have histories. A particular space is what it is because it has come to be that space over time. This applies to a building, a city, or an environment. The forces that create the spaces differ, some are through human intervention, like zoning, some are forces of nature. But spaces have histories and the interesting thing about these historical spaces is that there seems to be something like an evolutionary success story to the spaces that have sustained a certain continuity over time.[2] That is, some types of spaces work better in some spaces than in other spaces. And when it comes to building new spaces, I would suggest that we apply something I want to call *architectural common sense*. This is basically the normative claim:

– the space should fit the space it is in, *ceteris paribus*.

In talking about spaces in spaces, it should be clear that I am talking about the location and external look of a space. There are other issues as well to be considered, but time and space make these topics for other times. However, two seem especially important to at least note them. The first concerns the notion of function. That is: Does the space do what it is supposed to do? Having raised that issue, another immediately springs to mind: Who determines what the space is supposed to do? The ready answer, the person or institution that issues the commission, is problematic since the users of the space often have interests in conflict with those who commissioned the space and with those who designed it. Who determines whether the space in fact accomplishes what it is supposed to is another question like the first to be left unanswered.

The point to be established is that the criterion for claiming a space is a good piece of architecture is that the space fits. I like this idea for many reasons. One is tempted to ask what it means, however. And that would be reasonable. So, as a first stab consider the following:

– A space fits in a space if it is in harmony with the space it is in.

---

[2]This idea that spaces have histories and that knowing that history is important in design derives in part from some earlier ideas. In (Pitt 2007) I introduced the notion of explanatory contexts. The mark of an explanatory context when dealing with historical material is that it tells a coherent story. In (Pitt 2001) I elaborated the notion of a coherent story into a philosophical problematic, where the point is made that to understand a philosophical problem in an historical context one must know its past history and, if possible, its resolution or its projected resolutions. Echoes of these ideas are to be found in the ideas of common sense design criteria.

To understand what it is to be in harmony with a space is best approached negatively, that is, it is easier to explain when a space is not harmonious than to explain what harmony means. This approach has many drawbacks. In particular, by saying what harmony is not is not to say what it is. However, there is no need to nail down a set of necessary and sufficient conditions, since, as I argued above, the meanings of the criteria change in the act of application.

Nevertheless, there are several things we can say about harmony that should at least set us on the track to, if not a definition, at least a characterization. To begin with, there seems to be a scale on which different degrees of harmony can be mapped. At the extreme end of the scale is the religious sense of harmony found, for example, in Buddhism. Closer to our theme is the harmony of the Japanese Tea Ceremony. At the other end of the scale is the lack of harmony we find in a space that startles us or which continually draws our attention back to it because of a sense of inappropriateness. At this end I would place Michael Graves' Portland building., discussed earlier.

Assuming that this example has provided us with some sense of what it is for a space not to be in harmony with another space, let us now take another look at what appear from a distance to be dual skyscrapers which Graves designed in The Hague, the Netherlands, also discussed earlier in a rather negative mode. From the train as you pass by them at a distance, they appear to be almost perfect. Tall and massive, they have exaggerated traditional Dutch rooflines that make their placement appear natural. They appear to be wonderful examples of the common sense architecture of which I spoke earlier. Graves has managed to bring forth a traditional design that has withstood the test of time and yet given it a clearly modern presentation. There are historically good reasons for the style of roofline mostly having to do with the weather. In addition, the style has acquired a kind of emblematic nationalistic character. These are clearly Dutch.

Unfortunately, as noted earlier, the buildings contribute to a kind of artificial demarcation of parts of the city. What appeared from a distance to be two separate buildings are in fact part of a single complex grounded in a massive base. In this respect it behaves very much like the Portland building. The complex interferes in the life of the city, it has negative and dark aspects to it.

Is the space marked by the The Hague Graves complex harmonious or not? The answer is not, as you might think, "it depends", rather I would claim that it is not given that, in one clear sense, it really does not "fit", since it, like Graves' Portland building, does not contribute positively to the environment it is in, rather it disrupts it. Yet the lack of harmony is not exactly the same with respect to the two spaces, and this is part of what I mean when I said that the meaning of the concept is modified by its application. On the positive side, the Graves building asserts "Dutch" in a Dutch environment. On the negative side, it has a negative impact on the social life of that space. The Graves Portland building, however, could also be said to have a disruptive social affect since it sits in a space that probably would be better served as a park. But who knows, the possible park could become a major location for drug dealers and other undesirables. Irrespective of its social impact, it remains the case that it is visually not a fit. There is nothing in the design that says it belongs there, that it has anything in common with the neighborhood, that it has a historical linkage with the area. It is just an ugly building plopped down in the middle of a city to which it has no relevance.

The more one thinks about it, the more the notion of relevance becomes increasingly important in evaluating a space. To see my point, let us return to the Graves' Hague complex. Surely, one would say they are relevant. They are governmental buildings, their monumental size is traditional in government buildings, speaking to the transcendent nature of government. They are clearly Dutch government buildings, so there is a second relevant feature. However, if Graves had put these two buildings where he put his monstrosity in Portland, they would have clearly been out of place and clearly not a fit. The interesting question is ”Why not?” It seems that when talking of relevance, we have to look at additional features of the space. Are they, for example, relevant to that city, conceived as an historical space? Not as they stand. If the city decided to build a new governmental center at the outskirts of town, that might have been a different story. In fact, it would have been a wise thing to do, it could have been an opportunity to showcase the modern Netherlands and highlight its vibrancy and dynamism. As it currently stands, those buildings are disruptive of the space they are in and you cannot be both disruptive and harmonious. Celebration of the site is crucial but not by way of degrading what else is already there. Visual excitement is important, but not to the point of contributing to an overall failure of the impact on the space.

13.5 Common Sense Design

Let me conclude with a few comments on common sense design. My appeal to “fit”, and “harmony” has as much to do with creating a space in which to live and work as they do with history. And harmony seems to require even more. Having a sense of the historicity of the space is part of what is needed to live in harmony in it. On the surface it makes no sense to put a modern 60 story glass and steel skyscraper in the middle of an ancient village of 200 people. That does not require a fully developed aesthetics, just, it might seem, common sense. It would be an insult to the generations of inhabitants of the village and the values and way of life they have contributed to the culture. Yet our Jamesean sense of common sense brings with it this very sense of historicity, in that there is a definite case of cognitive dissonance that emerges when we try to project the image of a 60 story glass and steel skyscraper into Delft's town hall plaza. But why should this be so? It is, I submit because given our past experiences of cities like Delft, we do not expect to see such a space in that space. Goodman, in speaking of Hume's account of induction puts it this way.

> Regularities in experience, according to [Hume], give rise to habits of expectation; thus it is predictions conforming to past regularities that are normal or valid. But Hume overlooks the fact that some regularities do and some do not establish such habits. . .(Goodman 1955, p. 81)

Goodman's solution is his theory of projection. My solution is to say that certain expectations, in the form of standardly used but thoroughly unexamined inferences bring with them the history of those expectations. And they do so by way of there having been developed over time acceptable inferences which we are taught to make because they have been successful in guiding action.

Yet, when we invoke the power of history we must be careful. History is a complicated mistress. While she grounds us in the past, we must not, at the same time, consider the past as something concrete. In short, to be grounded is not to be stuck. I am not denying that there were events that transpired over time in a certain order, etc. Let's call that "what actually happened" or History 1. Nor am I talking about history as *the narrative* we construct about what happened: History 2. Furthermore, in constructing such a narrative we need to be alert to the *historiography* we employ, History 3. Thus we might employ certain terms in a manner that suggests they are constants. An example could possibly be my use of the term "Dutch" in describing the Graves complex. On the other hand, if I am true to my earlier comments, terms like "Dutch" ought to change over time due to a variety of historical contingencies. Thus it would be inappropriate to refer to the people living in the area around the Netherlands as Dutch in 1250 BCE since, according to the Oxford English Dictionary, the term was first used in the ninth Century BCE to refer to Germans (hence, Deutschland) and only gradually restricted to what we now know as the Netherlands, beginning in the sixteenth century. So, in a sense we can say that history changes, that is, History 2 changes. The narrative changes as we learn more about the past and as we change our criteria for how to construct an adequate narrative (History 3). Keeping that thought in mind, we can offer a different, and even a positive assessment of the Graves complex in The Hague.

13.6 Conclusion – Graves Reconsidered and the Mystery of the Guggenheim Finally Solved

In their attempt to hold back the sea and increase its usable land mass, the Dutch have become increasingly concerned and identified with the technology of dikes and pumps, and with their constant battle with nature to secure their limited space. The meaning of being Dutch has changed from being identified with a sea faring colonial empire to that associated with a highly technologically sophisticated culture directly confronting nature. In the light of that evolving history, Graves, in his The Hague complex, instead of what I had suggested above, could be seen as looking to the future of the Netherlands, with its increasing dependency on massive and sophisticated technologies and how it might solve past problems in a technologically futuristic fashion. A closer look at the The Hague complex reveals a complicated set of interconnected buildings and elevators that might be construed as a futuristic dam, pointing the way to the next stage in the evolution of Dutch culture. Hence its massive and forbidding base can now still be seen as massive, but because that kind of a dam needs that kind of base. Further, what could be described as threatening the park on one side of it, can now be seen as defending it from the intruding ocean. Likewise, constructing a 60 story skyscraper in the middle of Delft's central square might also suggest the future by way of providing a means for providing living space in the face of decreasing opportunities for land expansion and the need for alternatives to the traditional Dutch way of living in single family houses. In so doing, what, on one view, could be seen as an affront to Dutch cultural sensibilities, might, on this one, be a means for suggesting solutions for historical problems.

One final example: the Guggenheim Museum in New York City. To put it mildly, when first unveiled it raised a significant fuss. In a line with traditional town houses facing Central Park, it presents not a traditional flat face but a curved space clearly descending in a spiral from top to bottom. In one sense it can be seen as totally out of place in that environment. It breaks the line one's eye follows as you look up the avenue. It sticks out and disturbs its surroundings. What was Frank Lloyd Wright thinking?

Let me suggest that he was thinking about the history of art and demanding that we reconsider how we think about it as well. Traditional art museums present their displays in disjointed rooms. In this way we can look at seventeenth Century Dutch painting in one room, and nineteenth Century American Romanticism in another, thereby allowing us to capture a snap shot of art history. But what if that is the wrong way to view the history of art? Is it really the case that we can draw clear boundaries between the sixteenth and seventeenth centuries, or between American and Dutch art? What Wright said to us *via* the Guggenheim is that the history of art is a continuum and to see it that way you need a different type building, and the rest is history, so to speak.

In sum: Common Sense is a set of responses to the challenges of an environment based on an historical appreciation of that environment and what counts as successful action in it. To be successful means you need to be thinking not just about the history, but also about the problems that history has confronted, some of which remain unresolved. Common Sense is, then a way of thinking about decision making which leads to actions that take into account the successes, failures, and values of the past and builds the future in light of those successes, failures, and values.[3] Finally, I would add that one of the hallmarks of Common Sense is its appropriation of new techniques as they are developed. It is not commonsensical to reject new materials, technologies, and techniques when they provide the means to solve problems we have been unable to resolve in the past. So, if Common Sense principles of architectural design insist the space must fit, what it takes to fit includes more than some kind of visual harmony; fitting also includes fixing problems. In so doing, we may be forced to acknowledge what we have been unwilling to do before, that older values have been supplanted. In that respect, Common Sense is not nostalgic, it always looks to the future.

[3] For an elaboration of this view see the decision-making model developed in (Pitt 2000).

Acknowledgements I wish to thank Carla Corbin and Thomas Staley for many helpful comments and corrections on earlier drafts and especially to Steven Moore and Thomas Staley for pushing me to the necessary conclusion. Of course, remaining errors remain my responsibility.

References

Goodman, N. 1955. *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.

James, W. 1907 and 1981. *Pragmatism*. Indianapolis, IL: Hackett.

Pitt, J.C. 2000. *Thinking About Technology.* Seven Bridges Press, New York, http://www.phil.vt.edu/HTML/people/pittjoseph.htm.

Pitt, J.C. 2001. "The Dilemma of Case Studies". *Perspectives on Science: Historical, Philosophical, Social* 9(4): 373–382.

Pitt, J.C. 2007. "Seeing Nature: Origins of Scientific Observation". In *Conceptions de la science: Hier, Aujourd'hui, Demain, Hommage à Marjorie Grene, sous la direction de Jean Gayon et Richard M. Burian, avec la collaboration de Marie-Claude Lorne*. Ousia, Bruxelles, pp. 272–289.

Venturi, R. 1972. *Learning from Las Vegas.* Cambridge, MA: MIT Press.

Chapter 14
Philosophy, Engineering, and the Sciences

Philosophers don't like details when it comes to facts. They are perfectly happy to worry to death the myriad meanings of "meaning". But when it comes to working though the factual components of what are fundamentally empirical claims the work is slim. And because some of what looks like a philosophical claim, but for what is really an empirical claim, the results of not finding out what is really the case can result in some philosophical claims appearing rather stupid. One I have in mind concerns the relation between science and technology, or more specifically between science and engineering.[1] In this chapter I look at something that has not been looked at by philosophers: the real world interaction between doing science and engineering. What I argue is that contemporary science cannot be conducted until some serious engineering is already in place. This may not be news to scientists and engineers, but it is news to philosophers, especially to a distinct group of philosophers of science who tend to think of science in isolation from the real world. These philosophers are concerned with such issues as the logical structure of explanation or the role of probability in the logic of confirmation. But the results have nothing to do with understanding how science really works, meaning by that how scientists go about their research.

We all know the *old* story: the scientists do basic research and technologists (for our purpose here, specifically engineers) apply it. This is a troublesome account because something about it doesn't ring true. In particular, *how* does the move from basic science to applications take place? The results of basic scientific research are published (when they can be published[2]) in very specialized venues using very

---

Originally appeared in *Philosophy and Engineering*, edited by Ibo van de Poel and David Goldberg. Springer Academic Publishers, 2010. Reprinted by permission of Springer Academic Publishers.

[1]Talking about "technology" as if it is a thing in itself is unhelpful. I argue this case in my 2000 book. Likewise for "science". In that work I argue for the need to look at some category of practitioners comparable to scientists if we are to learn anything of value. I lay out some criteria that lead me identify engineers as the technological counterpart to scientists. But this ultimately is inadequate.

[2]Restrictions on the dissemination of research results often are found when scientific research is conducted for the military or by private laboratories funded by industrial or pharmaceutical companies.

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_14,

specialized language not readily accessible to most mere mortals. I am not even suggesting that engineers can't read this literature, but simply asking if, given their other responsibilities, can they find the time to do so? It is not clear to me that having made some discovery or other that the scientist picks up the phone and tells his engineering colleague "now you can do this or that". Nor is it clear that engineers keep close track of the burgeoning scientific literature to find out what's new and have immediate "ah, ha!" moments, or even later, "duh" moments. There is also the complex problem of intellectual property rights, finding interested investors, manufacturers, distribution routes, etc. So, in the end, the old story is not only a false story, but highly misleading. I want to tell a different story. The point is this: if the technological infrastructure of science is, in part, the product of engineering research and hands on design and inspection, and the research that makes it possible to build both labs and instruments comes out of engineering research, then engineering research is just as fundamental as scientific research. But there is more, for the result of looking at the relation between science and engineering through these lens results in seeing that conceptualizing the issue in terms of who is subservient to who, science or engineering, is wrong from the start.

Although he hasn't said this explicitly, the argument I am proposing is congenial to the views Peter Galison (1997) develops in his *Image and Logic*. In particular, I have in mind his distinction between the inner and outer lab, especially the outer lab. I take Galison's outer lab to be amenable to the notion I introduced in my *Thinking About Technology*, the technological infrastructure of science. What the technological infrastructure idea is supposed to capture is the range of things that make the doing of science possible: funding agencies, universities, private corporations, technicians, labs, graduate students, etc. I will proceed by returning to my motivating issue, which is whether or not it is correct to think of the relationship been science and engineering as one of subservience. I am going to begin by looking at a couple of examples and argue for a more encompassing view.

Some cases are easier to understand than others – i.e., some discoveries more readily suggest applications than others. This can happen is many ways, but looking at two examples will yield the general idea. In the first case a scientist can be looking to achieve a specific end which itself is an application. Thus, consider a microbiologist working on transmission of micro-organisms through ground water, who sets out to construct a bug that will eat oil. Let's say he started on this project after hearing of a particular disruptive oil spill when a tanker went aground. After successfully creating the oil-eating bug, he sets up his own company, rents some space at the Virginia Tech Corporate Research Center, hires some graduate students to make the things, contacted oil companies to inform them of the product and is now making money hand over fist.[3] What is missing from this picture is an engineer. Moreover, while a close examination of the process employed to create the bug deeply resembles a classic engineering design process, complete with feedback loops, our researcher is a biologist, not an engineer. In this case, the line between

[3] This description is based on a real episode.

scientist and engineer is clearly blurred, given the standard story. But, the positive result of this story is that it opens up the possibility that in some cases, the so-called scientific method is more like an engineering design process than some idealized and false view of how scientists do their work. Here we started with a product in mind and after considering the restricting parameters – the end product must be inexpensive to produce, must pose no danger to the environment, must be easy to transport, etc. – proceeded to propose a mechanism, test it, refine it, retest, etc.

A second example of how scientists connect to applications is the "opps!" case. This occurs when a mistake is made or an accident occurs in a lab and an unintended result comes up that has immediate applications because of the result itself. The process by which this discovery makes its way into the public domain may or may not involve engineers down the road, but the awareness of its applicability does not. For example consider the case of penicillin. The following account is from the Discovery Channel web page.

> Alexander Fleming discovered penicillin in 1928. Of course he wasn't actually looking for it at the time- he was researching the 'flu'. He noticed that one of his petri dishes had become contaminated with mould. Other scientists may have recoiled in horror at this result of shoddy work practice, but not Alexander. He chose to investigate.
>
> Whatever this intruder was, it was killing off the Staphylococcus bug – a bug causing everything from boils to toxic shock syndrome. Eventually he identified it as the fungus Penicillium notatum and it put the knife into Staph by means of a chemical that destroyed its ability to build cell walls. Being a scientist, he thought long and hard about what to call this new chemical, a chemical released from the fungus Penicillium notatum.
>
> That's right he called it penicillin. Nice one Alex. Unfortunately naturally occurring penicillin isn't very stable and thus not very useful. Fleming had found a wonder drug, but couldn't do much with it. Luckily just three years later two Oxford researchers created a stable form and today it's one of our most important tools in the fight against disease.

Consider now an example of a discovery that was delayed in its application and why.

This account is taken from Wikipedia – giving acknowledgement to all my student's resources.

> In 1968, Dr. Spencer Silver, a scientist also at 3 M in the United States, developed a "low-tack", reusable pressure sensitive adhesive. For five years, Silver promoted his invention within 3 M, both informally and through seminars, but without much success. In 1974, a colleague of his, Arthur Fry, who in a church choir in North St. Paul, Minnesota, was frustrated that his bookmarks kept falling out of his hymnal. He had attended one of Silver's seminars, and, while listening to a sermon in church, he came up with the idea of using the adhesive to anchor his bookmarks.[1]. He then developed the idea by taking advantage of 3 M's officially sanctioned bootlegging policy. 3 M launched the product in 1977 but it failed as consumers had not tried the product. A year later 3 M issued free samples to residents of Boise, Idaho, United States. 90% of people who tried them said that they would buy the product. By 1980 the product was sold nationwide in the US and a year later they were launched in Canada and Europe.[2]. Post-It Notes are produced exclusively at the 3 M plant in Cynthiana, KY. In 2003, the company came out with Post-it Super Sticky notes, with a stronger glue that adheres better to vertical and non-smooth surfaces.

The point of these two examples is to suggest that to understand the move from scientific discovery to practical application needs more than hand waving at science and technology as such. We have already observed that the process of going from a discovery to a practical application is more complicated that the standard story would lead us to believe. While complicated, it nevertheless seems possible to spell it out using the standard story. However, I want to argue that even doing so will still give us a skewed picture.

The picture is skewed because it starts with the scientist. It suggests that the scientist does research and comes up with discoveries, but it does not fill out the picture as to what is entailed by saying the scientist does research. To resolve this we need to pursue a classic Kantian transcendental argument: what does the scientist need in order to do what he or she does?

In order for a scientist to conduct research, he or she generally needs a lab. It can be as simple as a computer, or as complicated as a radio telescope, but to say a scientist conducts research entails that there is a context in which that research is done, even field scientists who study the behavior of the great apes treat the environment in which the apes live as their lab. Once we open that door, the entire picture changes.

In *Thinking About Technology* I introduced the notion of *the technological infrastructure of science* as "an historically determined set of mutually supporting artifacts and structures that enable human activity and provide the means for its development." (Pitt 2000, p. 129) As noted above, parts of this complex are the labs, graduate students, technicians, instruments, universities, and funding agencies that make modern science possible.

Consider what is involved in hiring a new scientist at a typical American university. I am not talking about the hiring process, i.e., the means by which the individual hired is selected – but rather the rest of the process that must be completed before the offer is accepted: the support package offered to the new potential hire as an enticement to accept the offer. No active researcher would think of accepting a position without being guaranteed a lab, i.e., a particular space and start up money to equip the lab with the appropriate equipment needed to conduct his or her research, to hire a technician or two or three and to support at least a couple of graduate students. The typical "start-up package" at my university for a new Ph.D. in one of the sciences or in one of the areas of engineering, coming out of school and off a 2 year post-doc is approximately $400,000. It obviously gets way more expensive for senior researchers.

Now let us unpack this a bit further. Laboratory space is expensive. Depending on the research to be done, there will be a water supply and sinks, exhaust hoods, computers, isolation spaces, etc., all housed in buildings meeting more stringent building codes (meaning costing more to build) than your typical classroom building. The differential here just for the costs of the buildings is $\$50/\text{ft}^2$ for a classroom building versus $\$150/\text{ft}^2$ for an unequipped laboratory building. Doing science is expensive.

Second, part of the start-up package involves the money needed to fund the research. But it is also money that provides the time for the researcher to develop a research program and to write grant proposals to support further research once the start-up monies run out. That means there have to be sources for that funding. I would argue that the sources of funding, like the American National Science Foundation and the National Institutes of Health on the public side and various foundations on the private side also constitute part of the technological infrastructure of science. They are enabling systems.[4] Moreover, by virtue of having the money and issuing calls for proposals in certain research areas on certain topics, they not only enable scientific research, but to a large extent they control its direction. In the United States federal sources of research funding may not fund stem cell research on strains of stem cells recently developed. This is having an interesting effect in two directions. (1) It is forcing certain kinds of research to be suspended or terminated for lack of funds. (2) It has pushed individual states like California to appropriate the funds themselves for such research, thereby putting them in the position to attract researchers in these areas away from states where they cannot do their work and making the universities and research centers in California a major force in this area. So funding sources make a difference in how science is done and what kinds of scientific research will be done and where it will be done. The picture of how scientific research is done and why is getting messy.

Let us return to the lab – for convenience sake let's make it a university lab. The picture sketched above is too simple – we don't just give the new researcher a lab and some money. The buildings have to be designed, built, and inspected to meet building codes and certain specifications. Instruments have to be designed and built. In short, not only are the funding agencies needed, the engineers who translate architects' designs into buildings and who make sure they meet building codes, as well as the engineers who design and oversee the building of instruments are essential infrastructure components for scientific research. The materials that are used in the buildings are the product of engineering research for the most part, and that research requires the same kind of support as scientific research does – labs, technicians, graduate students, funding agencies, etc.

To be even a bit more specific, the spaces where scientific research is done do not simply appear out of nowhere. It is designed space. And then it is built space. I am deliberately making a distinction here between designing the space and building it. It actually needs to be a threefold distinction: designing the space, figuring out how to build it, and building it. Architects, if they figure into this process at all in a significant way, work in the first part, designing the space. For the most part, the most significant part of the work involves figuring out how to make the proposed design work – and that is an engineering job. Architects are notorious for drawing lines that appear to connect and leave it up to engineers and builders to figure out how to actually make them connect. It is of no small note that the most successful architectural firms today – what are called full service firms –involve both architects and engineers in the process of getting a building from plan to fact, sometimes

[4] See "Research Space: Who Needs it, Who gets it, Who pays for it?" By Ira Fink (2004) for some hard data on these issues.

they also include the builders. So, the very spaces in which scientific research is conducted is heavily influenced by engineers. But there is more, for the materials used to build these spaces are constantly being improved thanks to engineering research into materials. And that research is conducted in much the same way scientific research is – in specially designed spaces, and so the cycle spirals upward and beyond. Peter Galison's account of laboratory design in *Image and Logic* speaks directly to this point.

What engineers do and how they do it, to coin a phrase, is fundamental to what scientists do and how they do it. So far I have only addressed the spaces where scientific research occurs; if you will, the building of the spaces. But if we also look inside the science lab, we find the footprints of the engineers all over the place. Maybe not in labs of the gorilla researchers, but in the labs in the buildings we have been discussing we find instruments. Sometimes instruments are designed and built by scientists. If you will allow the anachronistic use of the term "scientist", when we consider Galileo the scientist, then we also have to contend with Galileo the instrument maker. One of the sources of income he relied on was the sale of instruments he not only invented, or made popular, but also built and sold, such as his military compass and his telescope. (Drake 1978)

And it is well known that many contemporary scientists build their own experimental apparati, pulling this and that off the shelf, which is one of the things that makes replication of experimental results so difficult.

Nevertheless, when it comes to buying equipment from commercial suppliers to equip your science lab, engineers are involved up to their elbows. For in the production of standardized lab equipment engineers play a major role, for these instruments are their provenance. (See Baird 2004) In short, the contemporary scientist could not do her job without the engineer. There would be no appropriate space in which to work. The development of quality materials would be greatly delayed. If anything, there would be fewer and more poorly made instruments as well as whatever else is needed to fill out a functioning lab, instruments needed to conduct that work without engineers working independently.

This is not to say that from the beginning of time, engineers were central to the doing of science. The thesis I am reaching for is this: *modern* science relies on this technological infrastructure, in which large components involve work in which engineers play a major role. As historical backdrop it would be an interesting doctoral thesis to trace the historical development of the split of *scientia* into science and engineering. Something obviously happened in the sixteenth to seventeenth century. The *media scientia* were already recognized as doing something applied – both Da Vinci and Galileo were often employed as what we would today call engineers – working on military fortifications while doing multiple other things, like painting and writing music, etc. But to talk that way may be too simplistic. Why should we assume that there was a split into something like science and engineering from something like the *media scientia?* Maybe things don't happen that neatly. If you are looking to draw straight lines ignoring what is actually going on, you can probably do so. But straight lines are boring.

So I guess the bottom line here is that simple generalizations about the relation of this to that need a more nuanced historical analysis that goes deeply behind the surface to uncover what really is going on. I hope I have provided a schematic for the kinds of details that need to be examined – it is not presented as the full story by any means. There are a couple of problem areas here that we need to be sensitized to: (1) the reification of human activities – i.e., science as somehow something that doesn't take place in a time and place being done by people; (2) Galison's idea of how science changes, not all at once, but different parts changing at there own pace, works here; (3) the politics of priority – this has not been addressed in the current paper, but it is worth raising, even in passing: as any sociology undergraduate major will tell you, there is a competition in society among groups for some kind of social recognition. In our story it is alleged to be between science and engineering. But that just may be the wrong way to frame the discussion. It assumes there are these *things* that are called *science* and *engineering*, when if fact they are complexes of great complexity. Simplifying the rhetoric makes it easier to present a case for superiority or priority, but presenting the case does not make the case. It is one thing to talk about how the miracles of scientific discovery improve our lives than to go into details about the particulars of research into the biochemical structure of stem cells. There is much to be said about the rhetoric employed in the politics of the funding world. And we should make no mistake about it, it is all about money when the fancy language is put aside. In the end questions of priority and subservience seem to boil down to who gets the money to fund their favorite research projects. And what people do and say to get that money is a topic of endless fascination. That is one reason why we should address the particulars, the people and what they do. The other is that science and engineering simply don't do anything, people do.

Acknowledgement I wish to thank Ashley Shew and Nikolas Sakalarious for very helpful comments on an earlier draft.

References

Baird, D. 2004. *Thing Knowledge: A Philosophy of Scientific Instruments*. Berkeley, CA: University of California Press.

Drake, S. 1978. *Galileo at Work*. Chicago, IL: University of Chicago Press.

Fink, I. 2004. "Research Space: Who Needs It, Who Gets It, Who Pays for It?". *Planning for Higher Education*, 33(1), 5–17.

Galison, P. 1997. *Image and Logic*. Chicago, IL: University of Chicago Press.

Pitt, J.C. 2000. *Thinking About Technology*. New York, NY: originally published by Seven Bridges Press, now: http://www.phil.vt.edu/HTML/people/pittjoseph.htm

Chapter 15
What Engineers Know

To say that what engineers know constitutes engineering knowledge, just as what scientists know constitutes scientific knowledge, is a misleading way of expressing what ought to be a truism. For surely what constitutes scientific knowledge exceeds not only what one scientist knows but not even the sum total of what all scientists know – since there are scientific truths that no scientists may remember at any given time. Thus, Mendel's laws were forgotten until they were "rediscovered". On the other hand, it may be the case that the total of scientific knowledge is less than the sum of what all scientists know since what scientists know is not uniformly consistent. That is, what some scientists know is sometimes at odds with what other scientists know – perhaps even contradictory – hence a reduction in total knowledge.

Interestingly, the sum total of engineering knowledge does not seem to suffer from this problem. Contradictions do not seem to appear within the confines of the epistemology of engineering. There may be disagreements among engineers as to what is the most efficient solution to a problem but – given certain assumptions about the contingencies involved – it is not the case that two engineers similarly educated and experienced could be armed with sufficiently different perspectives that they would flat out contradict each other.

In this chapter I examine some aspects of engineering knowledge in order to determine what it is that engineers know. A lot will depend on how we construe "knowledge". I will argue for a pragmatic account of knowledge, in which, based on the very grounds on which the claim of superiority is made for scientific knowledge, engineering knowledge is shown to be far more reliable than scientific knowledge – thereby exposing the lie in the traditional view that science is our best and most successful means of producing knowledge. I will begin with a quick sketch of a pragmatic theory of knowledge, followed by a look at scientific knowledge before turning to engineering knowledge. I conclude with a look at the fate of some tradition philosophical problems.

Originally appeared in *Techné; The Society for Philosophy and Technology Quarterly Journal*. Vol. 3, 1998. Reprinted by permission of the editor.

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_15,

15.1 A Pragmatic Theory of Knowledge

Epistemology is an old topic and it remains stuck-in-a-rut. Since at least Plato, theories of knowledge have concentrated on one crucial factor – the inner mental state of a single individual. Prior to the work of David Hume that mental state was certainty. After Hume, empiricists abandoned certainty for some modified form of justified true belief. Nevertheless, the stress remains on what a single person knows. The view I am urging was first expressed in the work of Charles Saunders Peirce. The tradition Peirce founded extends through William James, John Dewey, C.I. Lewis, Nelson Goodman, W.V.O. Quine, Nicholas Rescher and, of course, Wilfrid Sellars, just to name a few. The simple idea they endorse in one form or another, is that to qualify as knowledge a proposition or set of propositions must be endorsed by an appropriate community. In *Thinking About Technology* (2000), I put it roughly this way: individuals produce candidate claims for knowledge, and these candidates become knowledge once they are endorsed by the appropriate community using agreed upon standards. This gives nothing to the Strong Programme sociologists, nor to the relativists – after all, Peirce was a realist. But it does relieve us from the fruitless tedium of devising doomed criteria by which we can determine whether an individual uttering a proposition with X, Y, and Z properties can be said to know something. The criteria are doomed because they ignore contingency, historical and otherwise. A pragmatic account, on the other hand, shifts the emphasis to, for example, the criteria that the scientific community has devised. But, even here, the criteria must meet some bottom line condition. For the pragmatist the bottom line is successful action. According to C.I. Lewis, "the utility of knowledge lies in the control it gives us, through appropriate action, over the quality of our future experience" (Lewis 1950, p. 4).

15.2 Scientific Knowledge

The nature, structure, and justification of scientific knowledge have been topics of central importance for most of the twentieth century. While it is still not clear that there is complete consensus on the criteria for scientific knowledge, nor should there be since science is an evolving activity, several key features have emerged from the discussion. These have grown out of a reassessment of criteria initially proposed for scientific knowledge in the course of the Scientific Revolution, when the kind of knowledge the New Science was proposed to deliver was alleged to differ fundamentally in kind from what had been previously accepted as knowledge – Aristotelian in character, proceeding from esoteric definitions of fundamental concept.

From the New Science tradition, there are several treasured characteristics of scientific knowledge that recent discussions have forced us to abandon or significantly modify. Given the New Science's emphasis on the role of mathematics, scientific knowledge was described as "universal," "true," and "certain." As the special features of the different sciences – most notably the social sciences – became more pronounced, however, the universality claim had to be modified and carefully bracketed. In the social sciences the development of social relativism made this inevitable. Scientific claims to "truth" and "certainity" suffered a similar fate. But in these cases the problems were not due to specific aspects of the individual sciences. Rather, they resulted from the difficulty of demonstrating the truth of scientific claims in a non question-begging manner – on the one hand – and – on the other hand – from the recognition of the fundamentally underdetermined nature of the relation between any scientific claim and its evidence.

Given emendations formulated in light of criticism, which arose in response to these newly reconstituted problems, the traditional account offers some features that remain viable. For example, it characterizes scientific knowledge as produced by researchers exploring the domain of a theory who aim to provide an account of the relations among the objects and processes of that domain, an account which provides the basis for an explanation of phenomena generally observed or detected in another domain. If I were tempted to isolate one crucial characteristic of scientific knowledge, it would be this: Scientific claims derive their meaning from the theories within which they are associated, hence, scientific knowledge is theory-bound.

The theory-bound nature of scientific knowledge presents additional problems beyond those noted above for some traditional assumptions about scientific knowledge – in particular the view that scientific knowledge, if true, is true for all time. If scientific knowledge is theory-bound, and if – as we know from the history of science – theories change, then scientific knowledge changes. Hence, what is accepted as scientific knowledge is not true for all time, at least not all of it, not yet. But this should not be a startling claim. The development of human knowledge is a process of continuous exploration in which we re-evaluate what we know in the course of new findings, and we jettison that which no longer remains consistent with the latest body of information.

We should note further that the tentative nature of scientific knowledge does not mean that knowledge is merely relative – especially in any sense that gives comfort to those opposed to the epistemic priority we traditionally give to scientific claims. The dynamic process in which scientists continuously revise what they are willing to endorse – and by which they examine their assumptions and their methods – is at the very heart of the strength of the sciences. Thus, despite the theory-bound nature of scientific knowledge, the self-critical process of scientific inquiry insures that the knowledge it claims is the best available at that time insofar as it is judged "best" according to community standards.

The ultimate aim of scientific inquiry is explanation. Thus, in the context of a pragmatic account, the ultimate success of the use of scientific knowledge is explanation. We use a theory to explore a domain of objects, sorting out their various relations for the purpose of explaining what can't be explained otherwise by appeal to the activities of the objects in that domain. Why is a tabletop hard? To answer that question we have found that we need to appeal to a scientific theory which proposes that there is a domain of smaller objects which are held together by a series of forces and that it is because of the forces and objects in that micro-domain that our phenomenological report of a hard table is possible. The aim of science is to help us understand the way the world appears to us, and it accomplishes this aim by constructing and testing theories that appeal to features of the world which are not immediately obvious.

There are other aspects of scientific knowledge that are essential to its vitality, but they need not be of concern here. In order to have a fruitful starting point to investigate the nature of engineering knowledge we need only concentrate on these two features; (1) Scientific knowledge is theory bound, and (2) scientific knowledge is developed to explain the way the world works. Unfortunately, while the process of trial and error and reappraisal characteristic of scientific activity seems to reveal its strength, this process also serves to undermine its claim of epistemic superiority over engineering knowledge. Likewise – as we shall see – the theory-bound nature of scientific knowledge creates a number of problems that do not plague engineering knowledge.

15.3 Engineering Knowledge

In What Engineers Know and How they Know it (1990), Walter Vincenti identifies and develops a theme first introduced by Edwin Layton in his landmark paper "Technology as Knowledge." Vincenti provides an account of engineering knowledge from the point of view of a practicing and deeply reflective engineer. Both Layton and Vincenti endorse the view that engineering knowledge – and technological knowledge in general – constitutes a discrete form of knowledge that is different from scientific knowledge. In a later piece, his classic 1987 Society for the History of Technology Presidential Address, "Through the Looking Glass or News from Lake Mirror Image," Layton endorses the findings of A.R. Hall, and claims that "technological knowledge is knowledge of how to do or make things, whereas the basic sciences have a more general form of knowing." (Layton 1987, p. 603) Vincenti echoes this, invoking Gilbert Ryle's famous distinction between knowing how (technology) and knowing that (science).

Both Layton and Vincenti are concerned to defend the view that – while both science and technology may borrow from or rely on each other in various ways – they constitute two distinct forms of knowledge since they aim at different ends. Science aims to explain and technology/engineering aims to create artifacts. Vincenti puts it this way, "technology, though it may apply science, is not the same as or entirely applied science" (Vincenti 1990, p. 4). He defends this claim in part with an intriguing and highly suggestive proposal. As he sees it, if we start with the proposition that technology is applied science, then there is no possibility of considering the view that technology could involve an autonomous form of knowledge that could account for those technological achievements which are science independent – such as the pyramids of Egypt and the roads of ancient Rome. Given the existence of highly visible science-independent technologies, we have good reasons to believe that we should not characterize technology as merely applied science. It is does not follow from the fact that science and technology each has occasion to rely on the other, nor that one is a subset of the other. Assuming its quasi-autonomous form, what can we say about the distinctive nature of engineering knowledge as a specific form of technological knowledge?

Starting from a wonderfully succinct definition of "engineering" by G.F.C. Rogers – which is highly reminiscent of Emmanuel Mesthene's definition of "technology" (Mesthene 1970, p. 25). – Vincenti identifies three main components of engineering and then concentrates on the notion of design. According to Rogers (as quoted by Vincenti and augmented somewhat by me),

> Engineering refers to the practice of organizing the design and construction (and I (Vincenti) would add "operation") of any artifice which transforms the physical (and, I (Pitt) would add, "social") world around us to meet some recognized need (Vincenti 1990, p. 6).

One of the commendable aspects of Rogers' definition is his characterization of engineering as a practice. That is, engineering – like science – is an activity with specific objectives. Given Rogers' insight and Mesthene's definition of "technology" as "the organization of knowledge for the achieving of practical purposes" – by a series of substitutions we see that, appropriately enough, engineering knowledge concerns the design, construction, and operation of artifacts for the purpose of manipulating the human environment. Vincenti proceeds to further narrow the focus of engineering knowledge to the topic of "design knowledge," by concentrating on design. It is worth quoting Vincenti's description of the design process at length because it immediately introduces an important distinction between the design as a set of plans and the design process.

> "Design", of course, denotes both the content of a set of plans (as in "the design for a new airplane") and the process by which those plans are produced. In the latter meaning, it typically involves tentative layout (or layouts) of the arrangement and dimensions of the artifact, checking of the candidate device by mathematical analysis or experimental test to see if it does the required job, and modification when (as commonly happens at first) it does not. Such procedure usually requires several iterations before finally dimensioned plans can be released for production. Events in the doing are also more complicated than such a brief outline suggests. Numerous difficult trade-offs may be required, calling for decisions on the basis of incomplete or uncertain knowledge. If available knowledge is inadequate, special research may have to be undertaken (Vincenti 1990, p. 7 – emphasis added).

The process Vincenti describes is "task specific" and essentially characterized by trial and error, but that still doesn't reveal the general nature of the contents of design knowledge. This is because to capture the nature of the knowledge required for any kind of task, Vincenti must invoke a detailed model which breaks that process up into both vertical and horizontal components, thereby allowing for a precise identification of what is needed when and where in the total design process. This schema is proposed for what Vincenti, calls normal design, as opposed to radical design. Normal design has five divisions beginning with the crucial aspect of any problem-solving process, the identification of the problem. Vincenti, an aeronautical engineer, draws from his own discipline for appropriate examples, but the schema is general enough to encompass a large number of design processes. For example, the design of an architectural project including sighting of the building, electrical systems, plumbing, etc., or the design of a space-based, orbiting telescope.

1. Project definition – translation of some usually ill-defined military or commercial requirement into a concrete technical problem for level
2. Overall design – layout of arrangement and proportions of the airplane to meet project definition.
3. Major-component design – division of project into wing design, fuselage design, landing-gear design, electrical-system design, etc.
4. Subdivision of areas of component design from level 3 according to engineering discipline required (e.g., aerodynamic wing design, structural wing design, mechanical wing design).
5. Further division of categories in level 4 into highly specific problems (e.g., aerodynamic wing design into problems of platform, airfoil section, and high-life devices). (Vincenti 1990, p. 9)

The process Vincenti outlines appears simple enough. One defines the problem, breaks it into components, and subdivides the areas by problem and specialty required, as needed. What is not obvious at first glance is the way in which the levels interact. Upon further reflection, one can see that what happens at level three will have ramifications for the overall design and visa versa, but recognizing this requires some work. In short, any design project must allow for a good deal of give and take throughout the process. In this respect, if one focuses only on the give and take, the design process sounds reminiscent of the scientific process. But there is a more and it clearly marks out a crucial difference between the process of scientific inquiry and engineering design. As Vincenti says it,

> Such successive division resolves the airplane problem into smaller manageable subproblems, each of which can be attacked in semi-isolation. The complete design process then goes on iteratively, up and down and horizontally through the hierarchy. (Vincenti 1990, p. 9, emphasis added)

If – by way of example – we apply this way of thinking to an architectural problem, we can easily determine what kind of a building to design (level 1), e.g., specific or multi-purpose, as opposed to the kinds of bathroom fixtures to have (level 4), although the one will ultimately bear on the other.

At this point we can pause and take stock of this comparison of scientific and engineering knowledge. First, the characterization of scientific knowledge as theory-bound and aiming at explanation appears to be in sharp contrast to the kind of knowledge Vincenti seeks. Engineering knowledge is task-specific and aims at the production of an artifact to serve a predetermined purpose.

There is a second important difference between the two forms of knowledge that is revealed by Vincenti's account of engineering knowledge. With engineering cast as a problem-solving activity (not in itself a characteristic which distinguishes it from other activities such as biology or even philosophy), the manner in which engineers solve their problems does have a distinctive aspect. The solution to specific kinds of problems ends up catalogued and recorded in the form of reference works which can be employed across engineering areas. For example, measuring material stress has been systematized to a great extent. Depending on the material, how to do it can be found in an appropriate book. This gives rise to the idea that much engineering is "cookbook engineering," but what is forgotten in this caricature is that another part of the necessary knowledge is knowing what book to look for. This a unique form of knowledge that engineers bring to problem solving. But there is more: We read the phrase "cookbook engineering" usually in a derogatory way. But what is wrong with it? If the knowledge in the book represents information we can use in a variety of circumstances, nay, in circumstances wherever certain contingencies hold – then isn't this knowledge that comes close to being universal, certain and, must we say it? – true? Could it be said that those who refer to engineering knowledge as stored in books as cookbook knowledge are employing a bit of rhetoric, in order to hide the inadequacies of scientific knowledge?

Contrast this cookbook knowledge with theory bound knowledge. When the theory is shown in some way or other to be flawed fundamentally, it is replaced. That means that what we thought we knew to be the case, isn't – which hardly sounds like knowledge to me. However, a good cookbook providing stress calculations can be used anywhere, anytime, as long as you factor in the appropriate contingencies. Just reflect on the basis of the metaphor – a good cookbook makes it possible for anyone to prepare a good meal.

Let's go one step further and contrast Vincenti's account of the engineering design process with the activity of science. I think it has been shown in sufficient detail in a number of places, by a number of people, that there is no such thing as the scientific method, i.e., that there exists one method which insures objectivity and guarantees the production of universal, certain and true knowledge. One appeal to the theory-based nature of scientific work should dispel any lingering illusions. In light of the fact that a scientist working within a theory is exploring the domain circumscribed by that theory, the direction of his or her research, i.e., the kind of research he or she will undertake, will be theory – determined. On the other hand, while the domain of the theory is necessarily where the research will be directed, there is no guide supplied by the theory as to what should be investigated and how. Further, there is no one method that works for all sciences. Consider Astronomy. Given the kind of one time only observations that we find in astronomy – replication, traditionally a cornerstone of scientific method, at least in principle, is impossible. Does this make astronomy not a science, hardly. On the other hand, Vincenti's account of the engineering design process provides specific and definite structure to the process of proceeding through the design process.

We can also go beyond Vincenti and look at the work of Larry Bucciarelli (1996), who denies that there is one single design process in engineering. Bucciarelli observes that no single unique design is dictated by the nature of the object being designed or the problem to be solved. But his objection stems not from the denial of design in engineering, but rather from a fine-grained understanding of the nature of the contingencies associated. That is, with Bucciarelli, we can find processes whereby the give and flow of ideas and the importance of the relevant contingencies follow the kind of pattern that Vincenti suggests, only in a more complicated way, when you consider the different types of communities interacting. The important point here is that in engineering design, there is at least a beginning point, for Vincenti, it is the problem, for Bucciarelli it is the object. Both see that whatever processes are at work they are dynamic and interactive, but they have a task-oriented beginning point, but no such beginning point is given for scientific research.

15.4 Philosophical Problems

Two possible consequences of the cookbook nature of engineering knowledge are: (1) That such knowledge can be transported across fields and (2) it can be used anywhere – the fundamentals of dam building do not change – the contingencies of the particular circumstances may dictate one approach over another, but the basics will remain solid. In contrast, scientific knowledge is not clearly "transportable" across fields in the same way as engineering knowledge. One crucial obstacle presents itself: The problem of incommensurability.

The problem of incommensurability is a philosophical problem that came to the forefront in large part with Kuhn's characterization of the nature of scientific change. For Kuhn, fundamental change in science occurs through paradigm replacement, with his view of incommensurability applying, primarily, across paradigms. A paradigm for Kuhn is many things. However, for the process of this discussion let us consider it as a complete system of thought, including methodological rules, metaphysical assumptions, practices, and linguistic conventions. Two paradigms are incommensurable, it is alleged, because claims in different paradigms cannot be compared so as to determine which claim from which paradigm is true.

For this view to be plausible, a particular theory of meaning must be assumed and a very dubious meta-linguistic assumption must be activated. First, let us look at the theory of meaning. Basically, the theory of meaning, behind the assumption of incommensurability, presumes that expressions receive their meaning contextually, within systems, i.e., paradigms, governed by unique sets of rules. This by itself is not so troublesome. The difficult part comes through the meta-linguistic assumption that there is no point of view common to both paradigms from which it would then be possible to compare claims from different paradigms. Such a common neutral point of view is necessary, it is argued, since the meanings of expressions are governed by the rules of the paradigm. If we shift an expression from one paradigm to another, its meaning will change since it will be determined according to different rules.

Among other difficult problems to sort through here is the apparently unjustified twofold assumption that there is one fundamental theory of meaning which applies to all paradigms, i.e., the meanings of expressions within any particular paradigm are determined by the rules of the paradigm, but, by contrast there is no single theory of meaning that allows for comparison of expressions across paradigms. However, if we can assert that all paradigms provide meanings for the terms which occur in that paradigm through the specification of rules, then why can we not, in the same meta-language in which we pronounce this dictum, then create another paradigm with the express purpose of allowing for the comparison of expressions? It is, for example, not at all obvious that the ways by which terms are made meaningful is through the specification of rules. That is, however, the account we are considering, and it is the source of Kuhn's problem of incommensurability. That much has been stipulated through Kuhn's account of a paradigm. But, unless something further prohibits us from doing so, surely we can say something like this: for the purpose of comparing two expressions, each drawn from a different paradigm: if the results of applying those expressions in the meta-language, according to the rules of the meta-language, is the same, in the meta-language, then for all accounts and purposes those two expressions mean the same thing. In short, if two expressions drawn from two different scientific theories yield the same result when transported into a third theory, then they can be said to make the same claim.

The solution is based on our account of engineering knowledge. If something formulated in the context of one paradigm can be used successfully in another, then deep philosophical problems about obscure theories of meaning recede. To treat the problem of incommensurability this way is not to solve it as much as to ignore it. This too may not be a bad thing. There are many philosophical problems still around to which we no longer pay attention since they seem beside the point, for example consider the pseudo problem of how many angels can dance on the head of a pin? It is not clear that this problem was ever solved, but who cares? And so too the problem of incommensurability. If the problem as stated was never solved it appears not to matter. This lack of concern is a function of having shifted our ground from worrying about providing an abstract philosophical justification for something that only philosophers worry about to a pragmatic condition of success: consider the consequences of using this claim from this theory in this context. If it solves our problem, then does it matter if we fail to have a philosophical justification for using it? To adopt this attitude is to reject the primary approach to philosophical analysis of science of the major part of the twentieth century, logical positivism, and to embrace pragmatism. This is a good thing to do, especially when we are concerned with technologies that have real world effects.

Finally, I noted that engineering knowledge was transportable, not just across fields but throughout the world (and perhaps beyond). Anticipating an objection from my colleagues concerned with various manifestations of cultural imperialism – let me attempt to forestall such issues. I am not saying that we should transport such knowledge. The appropriateness of such activities is a matter for policy considerations. That is not what I am talking about here.

Returning now to the issue I proposed at the beginning – that engineering knowledge is a more secure form of knowledge than scientific knowledge, on the very grounds by which it is alleged that scientific knowledge is our best form of knowledge. However briefly, we have noted that scientific knowledge is transitory – that it changes as theories change. We have also noted that scientific method is likewise not only transitory, but unstable, depending on the area of science being discussed, not only is there no method that will work across the sciences, within a science, the nature of the domain of objects being investigated may suggest different methods; compare biochemistry with botany. Finally if scientific knowledge is to be appraised through a pragmatic theory of knowledge, and given that the objective is explanation, then as theories change, explanations fail.

The history of science then becomes the history of failed theories and unsuccessful explanations.

In contrast we have engineering knowledge, which is task oriented. If the application of engineering knowledge, consisting of information in books and task specific methods and techniques results in the production of objects and the solutions of problems which meet the criteria of those for who the jobs are done, then it is successful. Because it is task oriented, and because real world tasks have a variety of contingencies to meet – e.g., materials, time frame, budget, etc., we know when an engineering project is successful or not. Further, those cookbooks represent the accumulated knowledge of what works. It is universal, certain and, if it works, must be true in some sense of "true". So, on the criteria we advocate for science, engineering knowledge seems more secure, more trustworthy, with longevity. What engineers know, therefore, is how to get the job done – primarily because they know what the job is.

References

Bucciarelli, L. 1996. *Designing Engineers*. Cambrdge, MA, MIT Press.

Layton, E. 1974. "Technology as Knowledge". *Technology and Culture*, 15 31–41.

Layton, E. 1987. "Through the Looking Glass or News from Lake Mirror Image." *Society for the History of Technology*, Presidential Address, Chicago: The Open Court Publishing Company. Court Press.

Lewis, C.I. 1950. *An Analysis of Knowledge and Valuation*. La Salle, Ill.: The Open Court Publishing.

Mesthene, E. 1970. *Technological Change: It's Impact on Man and Society*. Cambridge, MA: Harvard University Press.

Vincenti, W. 1990. *What Engineers Know and How They Know It*. Baltimore, MD: Johns Hopkins Press.

Chapter 16
Design Mistakes

The Case of the Hubble Space Telescope

This chapter is an exercise in what is being called the empirical turn in the philosophy of technology. As such it attempts some serious empirical work with respect to philosophical claims about technology. This may appear to suggest there is something novel here. To my ears, this latter claim presents a problem that needs to be addressed before we can proceed.

The problem is this: the effort to take an empirical turn in the philosophy of technology sounds like something novel and, in one sense, it is not. There has been lots of empirical work on philosophical topics dealing with technological issues. For example, one can point to Langdon Winner's investigations with respect to the electrical power industry or several of the studies on the impact of various technologies on planetary ecology. So how is this empirical turn something new? That is, how is this turn different from what has already been going on?

To answer this question, we have to understand the kind of empirical work done so far. Almost uniformly, at least in the United States, it is work with a political/ideological agenda in mind. That is, to the extent that philosophical discussions of technology have employed empirical data, it is employed in the service of an ideological perspective. The people who do this I call The Social Critics. Langdon Winner is such a social critic and his work presents us with a good example of how this is accomplished. (Winner 1986) His political position, that large power generating companies disenfranchise the little people because the companies, or rather their boards of governors, control the means of production, is patently Marxist. Now there is nothing wrong with having a political ideology to work from, it may even be impossible to avoid. What is wrong, and what I have been arguing for years (Pitt 1987, 1995), is to condemn science and technology for sins real or imagined on the basis of that or any other ideology. Or to approach it slightly differently, we make progress in the philosophy of technology when we do not allow our metaphysics to frame the discussion. If we start from a blatantly metaphysical point of view, then we have eliminated the possibility of meaningful discussion, since differing metaphysical frameworks tend to be incommensurable. Now if philosophy is a dialogue

Originally appeared in *The Empirical Turn in the Philosophy of Technology*, edited by Peter Kroes and Anthonie Meijers. JAI. 2000. Reprinted by permission of the editors.

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_16,

in which we engage in a mutual effort to examine problems of significance, then we must be able to speak to each other, hence the methodological injunction to tend to epistemological matters first and let the metaphysics emerge as it may. In this case it will be to examine a real world situation without any particular ideological eyeglass, i.e., to find out what the fact are and what they tell us about our theories of technology.

Two objections to this position immediately emerge: (1) isn't that position itself the expression of an ideology? and (2) the facts don't speak for themselves. In the case of (1), without rehearsing the fairly detailed explanation I have developed elsewhere (Pitt 2000), let me attempt to offer a quick denial. People use ideology to develop positions to attack other positions and to support their own political agendas. The attempt to understand, in all its complexity, the world of technology, for the sole purpose of understanding it, and not to argue for or against any particular restrictions of or additional funding for any particular project is what makes the approach developed here different and non-ideological. It is the absence of a secondary objective beyond understanding that removes the ideological sting.

With respect to the second objection, more needs to be said. What we see in the world is framed by our theoretical assumptions (metaphysics?). Philosophers of science have been telling us since Bachelard (1934) and Norwood Russell Hanson (1968) that observations are theory-laden and Richard Rudner (1953) reminded us early on that scientists make value judgments, hence they would seem to employ an ideology as well, it is just in the service of truth and understanding, but an ideology none-the-less since they also argue for public support on the grounds of the merits of science. How can we expect to have the empirical turn in the philosophy of technology produce any less biased results than those of the social critics if the observations, i.e., the facts, are also contaminated?

The difference is that here we are engaged in attempting to develop criteria for evaluating the empirical claims made about technology, which criteria ought to result in an improved ability to discern where and how ideology is employed. In short, it is not enough to understand a technology, however contextualized. Rather, we need to know how that knowledge is constrained by the methods, assumptions, and values we or others bring to the investigation.

In that spirit, my project here looks at three different claims about the design process, Vincenti's, Bucciarelli's, and my own, with an eye to determining how well these claims stand up in the light of real world activities. In the interest of time, only one case study is examined and even that is only partially developed, but it has the merits of having at its heart a flawed project, the Hubble Space Optical Telescope (HST), i.e., a product that embodied a design/manufacturing mistake. Here we have an example of a project gone wrong despite being eventually corrected. I argue that neither Vincenti's reiterative model nor Bucciarelli's claims about the social nature of the design process bear up when examined in light of what happened to the Hubble.

Now it should be asked why it is so important to accomplish this goal, i.e., to explain what went wrong. More is at stake than simply finding out why a project didn't meet the constraints of some theoretical model or other. It has, rather, to do with the very question of assessing the adequacy of theoretical models in general and, in turn, exposing what was referred to above as the ideological element in our models. In this case, the ideology is what I will call the Myth of the Engineer, in which engineers are portrayed as the paradigms (pace Kuhn) of rational and project-oriented problem solvers.

In labeling this ideology the Myth of the Engineer, I do not aim to offend. I seek to understand how things happen. This approach has been developing in the philosophy of science and in science and technology studies. It aims to understand how science works, Putnam's miracle, warts and all. In its infancy, the philosophy of science started with the assumption that science is our best means for achieving knowledge about the natural world and it set about attempting a justification for this claim. It just well may be true that science is our best knowledge generating machine (but maybe not, see Chapter 15), but that needs to be discovered, not assumed. The fully contextualized historical and socially informed philosophy of science that is now emerging justifies the early assumption of the epistemic superiority of science, but it also exposes the mess that is real daily scientific activity and it thereby provides a deeper appreciation of the genuine accomplishments of scientists when they do emerge.

There is one more reason for looking at a project gone wrong, rather than concentrating on engineering success stories. In epistemology, one of the criteria of adequacy for any theory of knowledge is that it be able to account for how we make false knowledge claims, as well as how we come to have knowledge and how that knowledge is warranted. The processes by which technological projects are accomplished is analogous in many ways to the knowledge producing process of science as well as being disanalogous in other ways (see Pitt 2000). To the extent that engineers are successful in producing products and in codifying the means for generating equal success in other projects, they develop something analogous to scientific theories and protocols. If that is correct, then philosophers of technology should be able to lay down criteria for determining if a project is a success or not. I will be attempting to show that Vincenti and Bucciarelli are producing helpful models or theories which account for successful design processes, but not for ones that fail. And for those that fail we need something stronger than falling back on the position that if the project failed the participants weren't very good engineers. I am not claiming that this what Vincenti and Bucciarelli do, but it is what their models make possible. Attacking the professionalism of the participants shifts the blame away from the inadequacy of our models in an ad hoc manner, which is a cop-out since we seek to understand what happened, not to assess blame. More to the point, if we can *show* that a project failed because it didn't follow our idealized models, and then turn around and *explain* the failure by blaming the participants for being defective in some way, we have in fact bought into The Myth of the Engineer. We are in effect saying that no well trained engineer would have allowed that to happen, and since it did happen then these folks can't be good engineers, hence this failure does not count against this particular model and its idealized account of how engineers ought to behave. To proceed in that fashion is just as flawed methodologically as Langdon Winner's claims that all big business is bad because of his Marxist orientation.

Part of what I have been arguing is that before we engage in assessing blame, we need to know what really is or was going on. Properly done social criticism requires a stage two activity. Not only does it require good detective work at stage one, but we also need to know what to look for in order to take the second step. I propose that *we look at who made what decisions and why.* Only by concentrating on the decision-making dimension of technological development can we proceed to responsible stage two work. (warning: do not confuse the talk of Stage One and Stage Two here with the different components of my model outlined below)

A couple more preliminary notes before proceeding to a discussion of the three models of the design process. First, I have located my discussion of technological issues in the context of engineering. There is a general problem with talking about technology per se, just as there is a problem with talking about science per se. Generalizations about science have been shown to fail in a variety of ways, so it should come as no surprise that the same should happen to generalizations about technology. To make a long story short, talk about science needs to be talk about scientists, and particular groups of scientists and particular scientists in those groups. That is the only way to make sense of what actually happens. To accomplish similar results when discussing technology, I have tried to identify a counterpart to scientists: engineers. Like scientists, engineers have professional societies, mirroring their specialties – journals to publish in, special educational programs for credentialing, etc. I don't know what a scientists is, but I do know a biologist when I meet one. I don't know what a technologist is, but I do know what electrical engineers do. And so the following discussion of technology is contextualized within an engineering framework, with the caveat that not all engineers are designers and not all designers are engineers.

The second point concerns the focus on design. If there is a similarity to science, then the design process is to technology what the scientific method is to science. But just as there is no scientific method *simpliciter,* there is no design process pure and simple. There are many, and they are the methodological counterpart to methods of experimental design and testing in the sciences. In what follows I will be looking at models of the design process to determine whether or not they stand up to what engineers actually do – are they explanatory, and if so, in what sense?

According to Vincenti,

> Design, apart from being normal or radical, is also multilevel and hierarchical. Interacting levels of design exist, depending on the nature of the immediate design task, the identity of some component of the device, or the engineering discipline required. For airplanes, which are typical devices that constitute complex systems, the levels run more or less from the top down:
>
> 1. Project definition – translation of some usually ill-defined military or commercial technical problem for level 2.
> 2. Overall design – layout of arrangement and proportions of the airplane to meet project definition.
> 3. Major-component design – division of project into wing design, fuselage design, landing-gear design, etc.

4. Subdivision of areas component design from level 3 according to engineering discipline required (e.g., aerodynamic wing design, structural wing design, mechanical wing design)
5. Further division of categories in level 4 into highly specific problems (aerodynamic wing design into problems of platform, airfoil section, and high-lift devices).

Such successive division resolves the airplane problem into smaller manageable subproblems, each of which can be attacked in semi-isolation. The complete design process then goes on iteratively, up and down and horizontally throughout the hierarchy. (Vincenti 1989, p. 9)

Despite the fact that Vincenti couches his analysis in the framework of airplane design, the general process seems generalizable to any other design project, be it architectural, landscaping, curriculum design, etc. Further it allows for movement back and forth between the levels and interactive readjustment. What it leaves out, in the form in which it is expressed, are the people and the social environment in which this all takes places. Without explicitly injecting the participants into the process, then the mechanism for moving between these various stages is missing. This is what Bucciarelli attempts to correct.

According to Bucciarelli,

There is no science of design process in the way the participants understand that term. This is not to say that the process is irrational, that a story can't be developed and told that makes sense, or that one cannot, on the basis of this story, infer improvements in the process. It *is* to claim that to be "scientific" about the study of design process one must admit the possibility that the object – as either physical principle or economic necessity – is only part of the picture, and a very fuzzy part at that. If we want to understand the design process, we must remain sensitive to the full breath and depth of social context and historical setting. (Bucciarelli, 1994, p. 18)

Bucciarelli is correct to emphasize the social and the historical and to recognize that the object is only part of the picture. Unfortunately, Bucciarelli continues a little further on to strengthen the claims about the social in ways that render his project suspect.

My working hypothesis is that the process is not autonomous, that there is more to it than the dressing up of a scientific principle, more than the hidden-handed evolution of optimum technique to meet human needs, and more than the playing out of the bureaucratic "interests" of participants seeking power, security, or prestige. In the affirmative this hypothesis takes the form: *Designing is a social process.* (Bucciarelli 1994, p. 20)

So far so good, Bucciarelli is taking the move to the social and distancing himself, so it seems, from the more radical social constructivists in science studies. But there is more:

Executive mandate, scientific law, marketplace needs – all are ingredients of the design process*, but more fundamental are the norms and practices of the subculture of the firm where the object serves as icon*. . . . . In the simplest terms, design is the intersection of different object worlds. No one dictates the form of the artifact. Hence design is best seen as a social process of negotiation and consensus, a consensus somewhat awkwardly expressed in the final product. (Bucciarelli 1994, pp. 20–21, emphasis added)

As we shall see, the claims about the lack of outside interference on the form of the artifact are a bit too strong; likewise, the overwhelming emphasis on the norms of the firm; and it is perhaps too strong to claim that “No one dictates the form of the artifact” for, especially where government contracts are involved, there are specifications to be met. Thus while the exact form of the artifact may not be dictated, its performance measures are. I do like the emphasis on the internal workings of the firm and the notion of compromise that comes out of a realistic understanding of the negotiation process; it is a powerful antidote to armchair theorizing. Elsewhere (Pitt 2000) I have spoken of the stark ontological primacy of compromise in the domain of the social. But the compromises take place among a large number of players, sometimes the most important ones are internal to the firm, but not always. In the case of the Hubble, as we shall see, governmental funding was a key issue that wasn’t merely reducible to the internal question of the firm’s profits.

In the design process, there is no one necessarily key player who plays the major role in determining the outcome known as the product. The model, MT, I have been working with starts with the definition of technology as *humanity at work* and unpacks that definition in terms of a multilevel in-put/out-put transformation process. The basis for this model is Glendon Schubert’s (1965) analysis of the structure of decision-making in the United Sates Supreme Court. Consistent with Schubert’s view of the nature of the decision-making/policymaking process is the consequence that there are input/output transformation processes whose function is to develop other input/output transformation processes. Let us distinguish between these as first- and second-order transformations. Decisions are first-order transformations. The result of a first-order transformation may be either another first-order transformation, i.e., a decision to make another decision, or a second-order transformation. i.e., a decision to create a tool of some sort.

A second-order transformation involves the construction of a device to meet specific goals An oil refinery is a second-order transformation. So is a legal system or a geometry. They are the result of first-order transformations in which a decision(s) was made using available knowledge, etc., to build, for example, a space telescope. Thus, decision-making procedures are first-order transformation processes. In the case of the construction of a space telescope we have a nice complicated example because the decision to build the telescope actually amounts to authorizing another series of first-order transformers which are the processes to be used for planning, designing, testing, and construction of the project. The building of the telescope involves further decisions as well such as the manipulation of materials. The completed telescope is itself a second-order transformer since it transforms raw materials by, roughly, mechanical means. So, using the basic notion of an input/output process we can still distinguish between mechanical and social processes and decision-making processes, thereby allowing such institutionalized decision-making processes as bureaucracies and funding agencies to be characterized as technologies.

But to characterize technology as humanity at work, meaning by that the set of first and second-order input/output transformations, doesn’t mean that we can’t analyze this account further. To complete the model we need more than the notion of levels of input/output transformations. There is a crucial third ingredient we must include if our model is going to reflect the most important component of humanity at work, where that work is acknowledged to take place in a social context with a number of players both internal to the decision making team and external to it interact. The final ingredient is assessment feedback.

Technology assessment is a special kind of decision-making in which the effects of implementing decisions of the first kind are illuminated by means of a feedback mechanism that makes it possible to upgrade the knowledge base for further decision-making. The scientific process, as one aspect of humanity at work, may be its best general example. In a science the constant reassessment of theoretical assumptions in the light of new results, both empirical and theoretical, is essential to the development of viable theories.

But more to the point, the most important aspect of contemporary technologies is the extent to which an assessment feedback mechanism is formally incorporated into the decision-making procedures surrounding the development and implementation of plans for new ventures. In some cases this is mandated by government. But more interesting is the situation today in which the importance of assessments and feedback loops is being insisted on because of the magnitude of some technological ventures and their potential consequences. It even appears that assessment has become an important value governing humanity's work. A full-scale discussion of this phenomenon (not to be undertaken here), would reveal the means by which changing goals and values affect the development and implementation of new and innovative techniques for transforming raw materials into suitable results in both the physical and social domains. Let me just mention one such change to illustrate the point. In the Netherlands, when a new oil refinery is proposed, it must include in its plans not only the usual environmental impact statements, but a step by step process for dismantling the plant when its working life is over and for the disposal of its various components. This is taking environmental impact assessment to the next level.

At this point it should be apparent that the proposed model is intended only to schematize the complexity and pervasiveness of technology rather than to be a definitive description of its structure. On any detailed analysis it will become clear that wherever a decision is involved, so too are a variety of other considerations, among them the assessment process, the nature of the second-order transformer, i.e., mechanical or social, the social circumstances, the goals of the individuals as well as the goals of the institutions, and so one cascades down a virtually unclosed spiral. The success of this model, then, should not be judged in terms of whether it simplifies our view or not, it doesn't; nor was it intended to. The world is a very complex place and we do ourselves serious disservice by thinking that we can always get around in it better by simplifying things. The merits of the proposed account are to be found in the manner in which it allows the complexity of the actual situation to be exhibited, while still providing the means to isolate and analyze the relevant components and their interaction in the light of constantly changing circumstances.

We now have our three models. Let us turn to the problem at hand. How did NASA manage to launch a defective telescope? Or, not to appear to put the blame on NASA, how did a defective telescope come to be launched? Most of the material I use here comes from the United States Congressional Record (November 16, 1993). The US Congress, following the discovery of the blurry images coming back from the Hubble, ordered an investigation. This resulted in a report presented to Congress by William Colvin, Inspector General of NASA. Several things are not in dispute: (1) The primary mirror of the HST was flawed. It suffered from a spherical aberration. (2) The flaw was the result of a manufacturing error and several failures on the part of management. (3) The manufacturer of the mirror, following the investigation, agreed to repay the government $25 million. The willingness of the manufacturer, Perkin-Elmer, to agree to this settlement, suggests something more than a mistake was involved.

The problem was that the main mirror of the Hubble suffered from a spherical aberration. According to Eric Chaisson:

> . . . spherical aberration is an undesired optical condition induced by a circular mirror having a deformed figure – that is, a curvature of the wrong "prescription." It has nothing to do with the smoothness of the mirror, but depends only, yet critically, on the shape of the mirror. Unable to focus all the light to a single point, such a misshapen mirror actually displays a range of foci. Alignment of the optics can emphasize the light captured either by the outer or by the inner parts of the mirror, but. . .the resulting image, regardless of these attempts, is a bright point surrounded by a disk of diffuse light. (Chaisson 1994, p. 174)

The revelation that the telescope suffered from this specific problem and not from a failure to focus the primary and secondary mirrors was first proposed by scientists at the Space Science Institute. The NASA engineers and the engineers at the company which manufactured the mirror, Perkin-Elmer, initially objected and sought to find the solution in adjusting the two mirrors for focus. After all efforts were exhausted, the engineers agreed to the diagnosis of a spherical aberration.

This was an expensive project, paid for by tax payer dollars – so naturally there had to be a Congressional Inquiry – The result of the House Inquiry was an extensive investigation with a report from NASA following. The investigation revealed the following six significant irregularities (and here I will quote directly form the testimony of William Colvin.)

> 1. **Non-approved Reflective Null Corrector Washers.** In the process of adjusting the spacings of the reflective null corrector for the HST Mirror, technicians discovered they could not move the field lens into the prescribed position. Instead of calling in the designer, the contractor inserted an ordinary washer under each of the bolts holding the field lens retainer to the adjustable plate. The insertion of these non-approved washers in an instrument whose precision is measured in tens of millionths of an inch required a nonconformance report. There is no written evidence that a report was generated.

A nonconformance report is a report describing how what was done deviated from the specifications of the government design. Why did the Perkin-Elmer engineers fail to file such a report? In an article in *Science,* the explanation rendered is that officials at both NASA and Perkin-Elmer "allowed themselves to be overwhelmed by the massive cost overruns and schedule slippages in other parts of the project. As a result they neglected the mirror work, which seemed to be relatively well, and failed to enforce their own quality assurance procedures". This would seem to argue against both the Vincenti and Bucciarelli models. For not only do we not have evidence here of the reiterative process Vincenti proscribes, which should have intervened when the technicians couldn't move the mirror into the right position and should have called for the designer, but didn't, but Bucciarelli's confidence in the firm's concern with the object seems to be missing as well. My model focuses on the decision making and it can deal with this situation by concentrating on why the decision was made not to call the designer back in: cost overruns and schedule delays, these are externalities with a major impact on the final product. Vincenti emphasizes the structure of the design process. Bucciarelli stresses the social context, especially that of the firm. My emphasis is on the decision-making and the factors that bear on it. In the remaining five items in Colvin's reports, you will see that each point comes back to the decision made by the individuals involved. The factors playing into those decisions varied, some of them were internal to the firm, some were not. However, only by directing attention to which people made what decisions with respect to the materials used and the objective to be obtained can we understand what happened. Unless it requires commentary I will just go through the last five findings of the investigation.

The second item Colvin relates is

> 2. **Unexpected results from Inverse Null Corrector.** The inverse null corrector was used by Perkin-Elmer as a part of the reflective null assembly. It could be used as a check on the reflective null corrector's alignment and stability. It could show whether there was a gross flaw in the reflective null corrector as well as measure its stability. Since the inverse null emulated what the primary mirror surface should be when finished, it should have produced an interferometric pattern with straight lines, or in other words a null condition.
>
> Instead, a pattern of way fringes was produced by the inverse null corrector at each testing cycle using the reflective null corrector. Although the inverse null corrector was producing results which did not agree with expectations, no nonconformance report was generated, and the designer was not consulted. While no qualitative analysis was performed, the inverse results were dismissed as being attributable to error in the null corrector's manufacture.
>
> As a result of this condition, Perkin-Elmer managers decided to make a design change which revised the usefulness of the null corrector from its initial utility as a "double check" on the health and stability of the reflective null. Through testimony we ascertained that Perkins-Elmer managers decided to use the inverse null corrector solely as a stability check.
>
> 3. **Refractive Null and Reflective Null Tests Do Not Agree.** Within a week of the first reserve null results, Perkin-Elmer received other unexplained testing results. A refractive null corrector was used by Perkin-Elmer in Wilton, CT, for initial grinding of the mirror blank to get the mirror surface shape roughly to the desired finished shape. The mirror was then transported to Danbury, CT, where testing was first conducted using the reflective null corrector.
>
> The test results from the first reflective null corrector test and last refractive null corrector test did not agree. Yet the expectation was the interferograms would be similar. In assessing the differences, Perkin-Elmer commented on the poor quality of the interferograms – caused by the rough surface of the mirror. They also asserted that some hand polishing of the outer edges of the mirror had occurred after the mirror left Wilton. No further quantitative analysis was performed.
>
> 4. **Recommended Gross Error Test Not Performed.** At the conclusion of the polishing phase, a Perkin-Elmer Vice President and General Manager formally requested an internal

> review of the primary mirror certification by senior scientists employed by Perkin-Elmer. This group, called the technical advisory board, held a review of the mirror fabrication and test results. On May 21, 1981, they recommended to the Perkin-Elmer Vice President that an independent test be performed on the mirror. The recommendation stated that "another test of the figure using an alternative method such as a Hartman test or a bean should be made. The purpose of the test would be to uncover some gross error such as an incorrect null corrector." No such test was ever performed.

If no such test was made, we can infer that the Perkins-Elmer Vice President decided not to make it. Why he decided that most probably has to do with the factors alluded to above, cost overruns and schedule delays.

> 5. **Vertical Radius Test Anomaly.** In May 1981, the most significant of the irregular events occurred when the refractive null corrector was used to measure the center of the curvature of the mirror. . . . the refractive null interferogram showed wavy lines – clear evidence of an error in one of the measuring devices and possibly the mirror. . . . the chance that two separate measuring devices, the inverse and refractive null correctors, would find matching errors and both be wrong and the reflective null right – is infinitesimal. . . . Yet Perkin-Elmer. . . . failed to resolve the discrepancies in a quantitative ways. . . . . had Perkins-Elmer attempted to determine the source of the error, with the analytical and measurement tools in place at the time, they could have determined that the flaw was in the reflective null corrector in one or two days. Perkin-Elmer personnel assumed erroneously that there were large "as built" errors in both the refractive and inverse null correctors. They ignored the results of the intended sanity checks on the reflective null corrector – . . . the results of the last. . . test concerned Perkin-Elmer managers, but they did not disclose the results or their concerns outside of the optics fabrication group. . . . to our best determination, Perkins-Elmer did not share the discrepant results of the vertex test with NASA. The NASA Plant Representative. . . was provided a copy of the center portion of the vertex radius test interferogram. This cropped version did not disclose the curved fringes which would have indicated a problem with the test results.

Here we are at a crucial point. Someone deliberately cut out the center of the interferogram in order to not reveal evidence of a problem. The decision to hide relevant evidence is at the heart of the resulting problem with HST, since as we shall see, other members of the Perkins-Elmer group were left in the dark. Colvin's report continues:

> Perkin-Elmer quality assurance personnel told us that they were present at daily meetings on the HST and they were never made aware of aberrant test results, nor were there any such discussions at the meetings.

If this last point is true, then it suggests that communication within the company was flawed. It is not clear if there was competition or merely an effort to not look bad. It begins to sound like upper management charged with overseeing manufacture of the mirror dug itself into a hole which further decisions simply deepened. I am not sure if this supports or undermines Buccarelli's view.

> Finally, there was rapid close out of the Perkins-Elmer optical fabrication team. The program manager had eleven items he wanted to check out, but was given approval to do so. And in conclusion, "[T]he Head of Manufacturing Optical Analysis stated that he and his manager would have couched the vertex radius anomaly in terms of a "need to recertify the reflective null corrector" "This is nothing more than an admission of intended obsfucation, which tends to support my conclusion above.

What are we to make of this incredible list of failures to report bad test results, of failure to adhere to design specifications and failure to adhere to protocols? It might be attributable to the internal ethos at Perkins-Elmer. But it seems there were other factors at work as well. Robert Smith (1989) details the myriad of political pressures that also were at work. Economic considerations were paramount. Virtually every decision in the early days through the mid-80s was made to lower costs. Everyone wanted a space telescope, no one wanted to pay for the real MacCoy. The decision to launch the HST on the shuttle forced further design changes that aggravated the complexity of the project. At every turn, as we uncover layer after layer, we find that it is the decisions that were made, motivated by varying considerations from political clout to economics to engineering design, that give us the clues to what actually happened. This is sketchy – the case itself is complicated because the technical nature of the design needs to be part of our conversation, as does the political context and the economic situation. For example, budgetary considerations forced a cut-back in the number of NASA inspectors assigned to the project. At every turn we find crucial decisions based on non-uniform considerations that point the way. Finally, we find that failure to assess decisions in light of the larger scope of the project are what ultimately led to failure. Why bad tests were not reported and discussed up and down the line is not clear. Why designers were not consulted when production problems emerged is not clear. What is clear is that the resulting engineering design failure was a failure to utilize the feedback function of my decision-making. model, MT.

It is important here to pause to consider a possible objection. It might be argued that the launching of the defective Hubble was not an engineering failure so much as one of management. Clearly the final decisions not to test and to ignore the results of tests that were already on record was a management decision. But the discussion here is not about engineering versus management. It concerns *the design process*. This was a government contract and it had a number of specific protocols built into it. These included consulting with designers when problems occurred, using the null refractor results as intended, etc. What happened was clearly a breakdown in the design process, a process that included both line engineers and management. To claim that this was all a management problem ignores the fact that it all began with some technicians failing to follow protocol. My concern is how the design process actually works. What we have here is an example of a process that is set out in writing and then ignored by virtually all parties.

In looking at Buccarelli's and Vincenti's models, with this case in front of us, several things can be noted. First, while Vincenti's structured account of the design process puts us on the road to systematic analysis, by leaving out people he omits the mechanism for moving back and forth between his various stages, i.e., the decisions. People are presupposed and that is not enough. Without specifically including people in the structure, we cannot focus on their decisions, and, it is clear that it is the decisions that form the loci for understanding.

As before, Buccarelli's move to the social takes us in the right direction. But I fear his work is less than successful by virtue of objectifying the object world in a way which delimits human agency. But, as it should be clear from the HST case, human decisions are at the heart of what happens to the object.

Finally, it is also clear that we cannot argue that these were just bad engineers who were responsible for the problems with the HST; that is, unless you build honesty, apple pie and the American Way into your definition of a good engineer. These folks knew enough to engage in systematic behavior designed to hide their mistake and they succeeded, until someone turned on the telescope.

What can we glean from this brief look at this case? If there is a moral it is this: the main focus of efforts to understand technological projects needs to be the people doing the work. It is their decisions that hold the promise of understanding how what happened happened. True, those decisions will be contextualized by a variety of factors, but few of them can be determined ahead of time except in the most general terms – people, materials, institutions, etc. – but that tells us little. The devil is in the details, grand schemes seem increasingly less appealing. This also applies to MT. All we really need to work on is a better understanding of why people do what they do. If MT helps at all, it is in breaking down the process of human activity into steps for analytic purposes. Perhaps this is what the empirical turn in the philosophy of technology is all about.

[Assignment to class: reconcile this chapter with the argument in Chapter 10]

References

Bachelard, G. 1934. *Le Novel Esprit scientifique*. Paris: Alcan

Bucciarelli, L. 1994. *Designing Engineers*. Cambridge, MA: MIT Press

Chaisson, E.J. 1994. *The Hubble Wars*. New York, NY: Harper-Collins

Hanson, N.R. 1968. *Patterns of Discovery*. New York, NY: Cambridge University Press.

Pitt, J.C. 1987. "The Autonomy of Technology". In Durbin P. ed., *Technology and Responsibility*, Philosophy and Technology, Vol. 3, pp. 99–114. Dordrecht: Kluwer.

Pitt, J.C. 1995. "Discovery, Telescopes and Progress" In Pitt J.C. ed., *New Directions in the Philosophy of Technology*, Philosophy and Technology, Vol. 11. Dordrecht: Kluwer.

Pitt, J.C. 2000. *Thinking About Technology*. New York, NY: Seven Bridges Press.

Rudner, R. 1953. "The Scientist *qua* Scientist Makes Value Judgments". *Philosophy of Science*, 20: 1–6

Schubert, G. 1965. *Judicial Policy-Making*. Glenview, IL: Scott, Foreman and Company.

Smith, R. 1989. *The Space Telescope*. New York, NY: Cambridge University Press.

Vincenti, W. 1989. *What Engineers Know and How they Know It*. Baltimore, MD: Johns Hopkins University Press.

Winner, L. 1986. *The Whale and the Reactor*. Chicago, IL: University of Chicago Press.

Part IV
Nano

Chapter 17
The Epistemology of the Very Small

17.1 Introduction

The world of nanotechnology is the world of the very small. According to Eugene Wong in his testimony to the Subcommittee on Basic Research of the US House of Representatives Committee on Science in June 1999,

> One nanometer is 1-billionth of a meter. To get an idea of the size, we can compare some familiar things. The diameter of an atom is about 1/4 of 1 nanometer. The diameter of a human hair of 10,000 nanometers. The protein molecules, which are so important, so critical to life, are several nanometers in size. Moving to man-made things. The smallest devices on commercially available chips are about 200 nanometers, whereas the smallest experimental chips are approximately 10 nanometers in their smallest dimension. (page 3, *Nanotechnology*: *The State of Nanoscience and its prospects for the next decade*.)

The question I want to investigate here is "how can we come to know what is going on in this domain of tiny things?" There are a couple of issues to be examined: (a) what do we mean by "know"? and (b) how do we access this domain? Some would argue that the two are separate – that we can come to an agreement on the meaning of "knowledge" independently of settling the question of how we can access the nano-world. I want to argue that this is not the case. What we come to know about the nano-world is very much a direct function of how we access it and the criteria we bring with us that allow us to evaluate that access. This claim is part of a larger thesis: that we also modify our conception of knowledge as we develop criteria for calibrating our instruments.

17.2 Seeing the Unobservable

One would think that there really isn't a problem here since, for the last 60–70 years in the philosophy of science there has been an on-going argument over the status of objects smaller than what we can see with the naked eye. Basically the question

---

Originally appeared in *Discovering the Nanoscale,* edited by Joachim Shummer, Alfred Nordmann and Davis Baird. ISO Press. 2004.

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_17,

to be answered is this: if you can't see it, is it real? The question actually is somewhat more complicated than the formulation just provided. It is usually couched in the context of determining whether or not non-observable objects proposed by scientific theories exist. This question cannot, however, be reduced to the question of observability alone, for not all theoretical entities are unobservable, e.g., galaxies, and not all theoretical unobservable entities are very small, e.g., black holes. Further, scientific theories are not the sorts of things that we can assert are true with absolute confidence. They are constantly being challenged, modified, changed, and revised. Further, given the constant state of flux of theories in use, no one really knows what the theory is at the time on which it is being worked out. We finalize the content and form of a theory only after we have rejected it and moved on to something else – we finalize these versions of theories in textbooks.[1] All this being the case, it is no wonder that the status of theoretical entities, entities proposed by scientific theories that have yet to be completely proven to be to true, is in question, in particular, those very small entities, the ones we can't see.

But, it might be objected, we *can* see them – by way of various microscopes – devices by their very names designed to scope (see) the very small (micro). Here, however, is where things get sticky. The crux of the matter has to do with the meaning of "to see". The meaning of the verb "to see" has changed over time. Further, I would argue, what it means to see something has changed precisely because we have developed instruments to help us "see" more and more in different ways. And, furthermore, we have come to call this "seeing" without attending to the fact that it is not "seeing" in the usual way. Further, because we are usually inattentive to that fact, we fail to capture the nuances of the conceptual difficulties we should encounter when we talk about seeing things through a microscope. Let me explain.

17.3 The Role of Metaphor

The sense in which we "see" though a microscope is different than the sense in which we "see" a tree or a coffee cup. Or to put it another way, we have extended the meaning of the verb "to see" to accommodate our use of microscopes. Or, to put it a third way, to talk about "seeing through a microscope" is to employ a metaphor. A metaphor is a way of easing our way into an understanding of the unknown by applying the familiar to the unfamiliar. We call a number of things "seeing" today because we metaphorically equate what we are doing with seeing as we naturally understand it.

For example, seeing through a microscope differs from seeing a tree with the naked eye because we don't have to learn how to see a tree. We may have to learn

[1] It is an interesting feature of undergraduate science education that undergraduate students are rarely, if ever, taught the latest, most up-to-date theories. The textbooks, I would argue, are out of date by the time they are published. This is one reason why getting undergraduate science students involved in research in an active laboratory is so important to the future of the scientific enterprise.

that that thing there is a tree, which is learning how to use our language – but you can run into a tree and hurt yourself and know that that thing there is what hurt you and not know that it is a tree. But can you do that when looking through a microscope lens? I would argue “no”. It is not because you cannot run into microscopic entities – it is rather that you can’t see them at all until a couple of things happen that aren’t required for seeing in the macro-world (i.e., the world of tables, chairs, trees, etc.; the world in which we live): (i) you have to learn how to use an instrument; (ii) you have to learn how to see what is there.

17.4 Learning to See Through Microscopes

Learning how to see through the microscope for the first time is difficult. You have to learn how to do a number of things, for example, not to get your eye too close to the lens, and keeping your head still and turning the focus knob at the same time Those things take a little while to master. But the truly hard part is learning to see what is on the slide. This problem was with us from the start. Consider what Hooke had to say in his *Micrographia* in 1665.

> What each of the delineated Subjects are, the following descriptions annext to each will inform. Of which I shall here, only once for all add. That in divers of them the Gravers have pretty well follows my directions and draughts; and that in making of them I have endeavored (as far as I was able) first to discover the true appearance, and next to make a plain representation of it. This I mention the rather, because of these kinds of Objects there is much more difficulty to discover the true shape, then of those visible to the naked eye, the same Object seeming quite differing, in one position of Light, from what it really is and may be discov’d in another. And therefore, I never began to make any draught before by many examinations in several lights, and in several positions to those lights, I had discover’d the true form. For it is exceeding difficult in some Objects, to distinguish between a prominency and a depression, between a shadow and a black stain, or a reflection and a whiteness in the colour. Besides, the transparency of most Objects renders them yet much more difficult then it they were opacous. (Hooke 1665, Preface)

In a letter to Oldenburg, Leeuwenhoek, sometimes called the father of the microscope, complained of the same problem “. . . some of the forms I see are so fine and small, that I don’t know how even a good draughtsman could trace them, unless he make them bigger.” (Quoted in Dobell 1932).

But yet we have learned how to see using a microscope – partially it required the development of cell-theory and later, the theory of crystals. That is, once we had a way of understanding the sorts of things we were looking at, we had the means to see them as separate and distinct items, possessed of various properties, shapes, and appendages. This requires theory. It is not enough to know how to use a microscope, one must know what to look for. What to look for is dictated by various theories about the domain of the small.

But even the possession of theory is not enough, we also must develop the means of individuating individuals one from another, and there we rely on, in the case of biological organisms, staining techniques. And, further, we had to learn to rely on the credibility of staining techniques. This is not a trivial matter. Let me relay a true story. When I was a graduate student at Western Ontario we lived in an old farmhouse that had been converted into two apartments. The upstairs was occupied by another young couple. Mike was a MS student in biology working on the eye of the Hackfish. He was having trouble staining his slides, so when he had the opportunity to attend a conference where he could ask for some help he leapt at it. At the conference he managed to corner the acknowledged expert on staining slides and explained his problem. The expert reportedly told Mike the secret to success: "first, turn off all the lights in the lab and make sure the windows are darkened. Then close your eyes and raise your left foot. Then, hopping on your right foot, make a 360 turn to the left. Then lift your right foot and do a 360 to the right. Then stain your slides." Mike was crushed. After he returned from the conference we had numerous discussions about what kind of a message the great man could have thought he was conveying, but never figured it out. Mike finished his degree, but he had lost his faith in science and left to go work for British Petroleum. He now lives in Calgary.

The moral of the story I take to be this: some of what we do in the process of seeing the very small involves a skill that cannot be taught by rote. That being the case, you would expect the results of using stains on slides to be doubtful, but, interestingly, they are not. Part of what is involved in seeing with a microscope involves accepting the fact that some people are better at staining slides than others, and we rely on them to prepare the slides. In a crucial way we have extended the concept of seeing by accepting the fact that it may take more than one person for a seeing to occur and, further, they both might not actually do the seeing.

In addition to learning to rely on staining techniques to provide us with access to the very small, we also have to accommodate what I will call the problem of focus. Prior to 1702, focusing was done the old fashioned way: you brought the object to be examined into focus by holding the object in one hand, the lens through which you were looking in the other, and adjusted them until something recognizable came into view. According to Gerard L'E. Turner, Leeuwenhoek's microscope was "A tiny lens contained in a metal plat, with a spike to hold the specimen close to the lens; the instrument was then handheld immediately in front of the eye." (Bud and Warner, 1998). James Wilson, an Englishman, developed the screw-barrel roughly 40 years later in 1702. The screw-barrel allowed for mechanical focusing. With the development of mechanical focusing, stability became a factor that could be mastered. So, as we have seen, learning how to see through the microscope involved a number of steps, advances in theory, skill, and in the mechanical arts themselves.

I would like to look closer at the problem of focusing. Learning to focus an instrument is now an accepted part of seeing. But consider how strange this is. You don't have to be taught to focus your eyes to see macro objects like tables and mountains. What occurs is a natural phenomenon. Our biology takes over. And when you think of it, it is a rather amazing feature of our bodies. Focusing a seeing instrument, however, is an unnatural act. And yet, because it is integral to seeing with that instrument, it has become accepted as part of what we do when we use an instrument to see. And it is all part of the extended metaphor we now employ when we talk about seeing through a microscope or a telescope. It includes staining slides (or in the case of a telescope, computer enhancing photographs or using colored filters), focusing instruments, theory, etc., all by way of accommodating what we do as similar to what our eyes do.

17.5 Learning to See with Electron Microscopes

In a fascinating study entitled *Picture Control, The Electron Microscope and the Transformation of Biology in America,* 1940–1960 (Rasmusen 1997), Nicholas Rasmussen, examines in great detail a number of these issues as they pertain to the electron microscope. In particular, he focuses on how what I have called the criteria for acceptance are established, i.e., the social domain. Allow me to offer a lengthy quote:

> . . . early biological electron microscopy involved a struggle for picture control on a number of levels. . . . . . a picture control figured in a biologist's subjective experience of the electron microscope as one of three relevant readouts, and along with focus, one of the two open to intervention. Of course, there was no such thing among the seven indicators and nineteen switches and knobs on the console of the Radio Corporation of America (RCA) EMU microscope, the instrument depicted in Cecil Hall's 1951 cartoon. All the more reason to strive for an understanding of what that knob was for. Control of who could make pictures with the electron microscope, how pictures should be made, what pictures would be printed, and how those pictures ought to be used in establishing biological facts were the dominant issues when the new instrument was introduced to biologists at the onset of the Second World War. . . . By the end of the war, a community of scientists in whom expertise was vested – authorized microscope users who for the most part agreed about who should use the instrument and in what manner for which purposes. . . . was established, and assumed a basic level of regulatory control. But for individual microscopists, control of the characteristics and interpretation of pictures remained a problem, and one that was divergently addressed in different biological subfields, even in different research programs within them. (p. 1)

Now Rasmussen is talking about the social evolution of standards in the same breath as the social evolution of consensus over who had access to the machines etc, and it sounds very social constructivist. But let's face it, the battles and issues he identifies, to the extent they are addressed in the social realm, are appropriately discussed as issues of power, access, and interpretation. Perhaps key among them is power. For what we are talking about is who sets the criteria and on what grounds. But no matter what the politics may be, there is a world out there that sets the bottom line. Or does it?

It is at this point that we need to distinguish between optical and electron microscopy. With optical microscopes we are actually looking at something. We prepare a slide by putting something on it. Further we are aware of the fact that when, for example, we stain a slide, we have introduced something to the slide and we can test to determine how that affects the specimen. What exactly we are seeing is a function of how we interpret what we see using theory, but that there is something there to see is clear. With an electron microscope, on the other hand, we do not "see" the specimen. The machine uses a extremely fine point on a stylus to reveal the contours of a surface without actually touching the surface. Instead of dealing with the physics of light and the properties of specimens as we do with an optical microscope, with the electron microscope we get a "picture" of that surface through the use of various computer programs which take the input from the stylus running over the surface and using the physical theory of the properties of matter "interpret" the results, producing an image.

The question here is the extent to which the machine creates the phenomena. There is a weak and a strong version of this claim. The weak version holds that without the machine we would not be able to see what we see. This would suggest that the things we see with the machine are there in the world, but we don't have the means to access them without the machine. That claim is fairly innocuous. The problem arises because of the stronger interpretation of the claim that the machine creates the phenomena, which is: what we see is an artifact of the machine itself – if doesn't exist in the real world until we have the machine. If that is true, then the next question becomes "well, what kind of a thing is it? Does it exist or not?" To address this let us consider in slightly greater what it is that an electron microscope does.

Rasmussen and Hawkes' give a rather succinct account that will assist us:

> An electron microscope produces a magnified image through a specimen's interaction with a beam of high energy electrons, usually 50–200 kilovolts. There are two principle forms of this instrument. In a transmission electron microscope (TEM), an electron beam at least as large as the imaged area passes through the specimen and forms an image on a fluorescent screen or photographic film. In a scanning electron microscope (SEM), an electron beam that is small compared with the imaged area passes over the specimen in a regular pattern, and a picture of the specimen surface is reconstructed on a video tube. Image contrast is formed in many ways. In the TEM, electrons are deflected by atoms inside the specimen, without absorption, creating a shadow pattern of greater and lesser electron transmission. In SEM, interaction of the beam with the specimen surface produces varying intensities of backscattered and secondarily released electrons for each position in the scan, and these are registered by a detector placed appropriately near the specimen. (Bud and Warner 1998, p. 382)

In each type of electron microscope, we end up with an image. But it is not an image directly obtained by seeing. The image is the result of a process in which the object under examination is not "caught" but rather reflected. But it is not reflected as a mirror reflects your face. It is a secondary reflection, almost like trying to draw the right hand wall of a hand-ball court by observing where the ball lands on the front court after angling it off the right hand wall. The assumption is that the image represents the object. But it is not a representation such as we find when we draw a picture or produce a painting, say, a still life. And yet, we are content to say that the images are reasonable pictures of the objects – even though we can't see the objects directly. Under normal circumstances, common sense would contest the claim that an image produced by an electron microscope is an accurate representation of a very small object that cannot be seen. But we accept the claim. Why? The question becomes more demanding when we consider some further complications. Rasmussen and Hawkes' lay out some of the problems for seeing biological specimens:

> The electron beam demands a vacuum, so specimens cannot be alive and require drying in some minimally destructive way. Since electrons interact strongly with matter, the beam penetrates only very thin specimens, Moreover, *the beam heats specimens, and so can alter volatile biological materials.* Contrast is another obstacle, since the different substances in living things vary little in opacity to electrons. (Bud and Warner 1998, p. 384, italics added for emphasis)

So an early major problem was the modification of the specimen by the electron beam. The solution was find a way to fix the specimen. In the biological sciences the solution was initially chemical, then supplemented by freezing. In the physical sciences this involved the development of techniques for coating the specimen with a thin film.

What is of interest to us is the fact that the development of means to stabilize the specimen did not alter the initial problem of the manner in which the electron microscope produces an image. The reliability of the image was not the issue, the stability of the specimen was. Essentially, we find the same situation as with the optical microscope: an evolving set of techniques and standards that fundamentally change our conception of seeing. But, what is interesting is that both the sense of seeing and the standards and techniques evolve together, with the end result being consensus on what a good image looks like, even though it is not an image in the earlier, pre-electron microscope, sense.

17.6 The Nano Scale and Nano Technology

So let us now return to the nano scale and nano technology. Nano technology is the construction of very small artifacts and systems of artifacts. It is miniaturization taken to the max.[2] And our question is how do we know that the things are working at the nano level as they are supposed to? One way is to look and see. And this is what we cannot do with electron microscopes without begging the question. A second way, much more economical and intellectually sound, is to wait and see if what these mini machines are supposed to do actually happens. It is a pragmatic solution. William James' most notable contribution to philosophy was the aphorism: For a thing to make a difference, there must be a difference. I do not believe that we will have a problem knowing if the nano machines are doing their job.

However, our understanding of our interaction with the nano world shares similar characteristics with what we mean when we see through a microscope. I quote again from the Congressional hearings on Nano Technology, and ask you to attend to the language carefully. Richard Smalley, Nobel Laureate, is discussing the impact of carbon nanotubes. He is discussing a slide he has put up on the screen.

[2]It is important to note that this is as far as we can go in miniaturization given our current state of technology since the next level down is the quantum level, where stability of the material is itself in doubt.

> As individual nanoscale molecules, these carbon nanotubes are unique. Just think of one at a time. They have been shown – here you see one draped across a few electrodes. They have been shown to be true molecular wires, to conduct electricity like copper – in fact, even better – and have already been assembled into the first molecular transistor ever built; with just a single molecule. (p. 9)

What struck me was the casual manner in which Smalley refers to seeing a single molecule. The idea that a single molecule could be a transistor is itself difficult to grasp, but, that withstanding, the ease with which he speaks of seeing the molecule is of a piece with how he speaks of manipulating them. It is both natural and, in the context of what we mean by "see", illustrative of the point I have been trying to make. The methods, standards and implications of modifying the language to accommodate the new technology comes slowly but of a piece. The stronger thesis that it is a metaphorical extension of standard usage will have to wait for another time for its defense. But just consider another claim by Smalley. This simple statement, so straight forward, and yet so misleading, makes the point. I know what it means to divert a small stream of water threatening to destroy my driveway by removing a tree limb that has blocked a drainage ditch. I pick it up and toss it into the field. By analogy I think I know what it means to put an atom where you want it to go, but I doubt if it is as simple as picking up a stick. Yet, the language of "putting atoms where you want them to go" makes it sound so familiar. What is really entailed? All we are talking about is manipulating atoms. Atoms, remember, are 1/4 of a nanometer in diameter. A nanometer is 1 billionth of a meter. To unpack Smalley's claim about putting atoms where you want them means understanding a lot about the means we have devised for doing this sort of work, the tools we have built and the assumptions we employ about what we are doing. My guess is that putting molecules where you want them is much like seeing through a microscope, it is now a team activity, involving sophisticated instruments and subsidiary techniques, a lot of theory, many theories, a lot of skill, and a lot of luck.

That seeing in the context of using SEMs and very large telescopes has become a team activity is not in itself something negative. The point here is that it is a different sort of thing than seeing a tree. It is important to note this difference because it helps us understanding how science changes. In particular, it is not just that what we mean by "see" has changed, it is that the introduction of these instruments changes how we do science. This is not the obvious point that science is increasingly a team activity, it is that we have a new way of understanding scientific change. The moral of the story is that the older theories of scientific change proposed by Kuhn, Lakatos, and Laudan, seen in the light of the impact of new and innovative technologies such as scientific instruments, are deeply flawed. Scientific change is not merely a matter of the logical conditions under which scientific theories can be abandoned or accepted. It is a far more complicated process heavily influenced by the role of innovative instruments and other technologies that not only change the nature of the enterprise, but change the meaning of concepts like scientific observation, evidence, experiment.[3] The impact of the new techniques required for a robust set of nanotechnologies will be important to watch as they will make a difference also in the manner in which we do the science of the very small.

References

Hooke, R. 1665, 1902. *Extracts from Micrographia*. Endinburgh: The Alembic Club.

Dobell, C. 1932, 1958. *Antony van Leeuwenhoek and his "Little animals"*. New York: Russell and Russell.

Pitt, J.C. 2000. *Thinking About Technology*. 146 pp. New York, NY: Seven Bridges Press.

Bud, R. and Warner, D. eds. 1998.*Instruments of Science.* 709 pp. New York, NY: Garland Publishing, Inc.

Rasmusen, N. 1997. *Picture Control*: *The Electron Microscope And the Transformation of Biology in America, 1940–1960.* 338 pp. Stanford, CA: Stanford University Press.

(1999) *Nanotechnology: The State of Nano-science and its Prospects for the Next Decade*. Washington, DC: U.S. Government Printing Office.

[3] For an elaboration of this theme see my *Thinking About Technology*.

Chapter 18
When Is an Image Not an Image?

The challenge is to tell the truth. In the world of nano this is not as easy as it sounds. Take, for example, the question of images claimed to represent what some nano configuration or another looks like. It is alleged Scanning Tunneling Electron Microscopes (STEMs henceforth) produce such images. Let's rehearse what happens: According to Rasmussen and Hawkes:

> . . . an electron beam that is small compared with the imaged area passes over the specimen in a regular pattern, and a picture of the specimen surface is reconstructed on a video tube. . . interaction of the beam with the specimen produces varying intensities of backscattered and secondarily released electrons for each position in the scan, and these are registered by a detector placed appropriately near the specimen. . . All electron microscopes depend on the capacity of magnetic and electric fields to alter the path of electron beams according to the laws of optics (1998, p. 383).

Using an STEM is one of the ways it is said that we can see what is going on at the nano level. However, I am suspicious. Or, to put it in a less antagonistic way, to accept this claim will, I believe, force us to expand or change our understanding of what it is to see something, and in this case in particular, to understand what constitutes an image. There is nothing wrong with this. The meaning of words do change over time – they often expand, as the meaning of "men" in "All men are created equal' has expanded to include African Americans, other minorities, and women. However, we often do not pay attention to the fact that while we continue to use a word whose meaning we think we understand, in this instance "see" and "image", we also sometimes extend the meaning of that word by applying it to novel situations where they only apply at best metaphorically, as I argue below. Eventually what is at first a metaphorical extension of the meaning of a term may become an accepted part of the meaning of the term, but we should be sensitive to the fact that the meanings of words change over time. This claim is part of a more general thesis I am developing: to explain what we are doing when we employ novel instrumentation, we often employ words whose meanings we already understand

---

Originally appeared in a joint issue of *Hyle*, *the International Journal of Chemistry,* Vol. 11, No.1 and *Techné: Research in Philosophy and Technology,* 2005, Vol. 8, No 2. Reprinted by permission of the editors.

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_18,

in an effort to characterize the sort of thing we think we are now doing with this new instrument, despite the fact that seeing through a microscope is not the same as opening one's eyes and seeing a tree in front of me, if we are to adhere to a strict sense of "seeing". I argue elsewhere that in extending the meaning of words metaphorically we also change the meanings of the family of concepts with which they are associated, such as evidence and explanation.[1]

If we take Rasmussen and Hawkes seriously, what the electron microscope does is to produce an image. But, I suggest, this is unintuitive for the reasons given below. Furthermore, to claim that an image is produced, suggesting by that the image is a genuine and realistic representation of what is really there, has serious ethical and social consequences. I want to talk about images first, and then I will turn to some of disturbing consequences of thinking about "seeing" by way of a STEM.

Imagine if you will, a very accurate tennis ball machine. It is a device that shoots tennis balls at you so you can practice returning them without having a serving partner. Lets assume you take this machine and aim it at a wall built from rough hewed stone. Your job is to construct an accurate representation of the surface of the wall simply by observing the directions of the balls as they bounce off the wall. Well, clearly you need some help to do this. You need to know a lot about the physics of objects colliding and how irregular surfaces change the vectors, etc. You also need to know a lot about translating what you see happening to the balls after they collide with the wall onto paper in a way that captures not the picture of the ball shooting off in this direction and then that, but the texture of the surface of the wall. It is not as if you are directly drawing what you see when you look at the wall. You are interpreting the action of the balls as indicating something about the surface and then you are putting that guess down on paper. That, with some minor modifications, is what the alleged image produced by an STEM is supposed to have accomplished. But instead of a person doing the drawing, a computer program does it. And, we are asked to consider the result an image of the surface. Take your hand, if you will, and run it over your shirt. Now draw what you felt. It is not easy is it? That is why I am asking this question, "when is an image not an image?"

Let us begin by trying to figure out what an image is. This is not an easy task, for we tend to use a substantial vocabulary of what we often take to be more or less synonymous terms when talking about what STEMs produce. Thus, there has been a lot of loose talk about images, representations, etc. Terms like these have been casually interchanged, mangled and generally semantically violated. I will not claim that I offer much of an improvement – but I at least want to alert us to the problem of image talk. In cases like this, my preferred method is to work our way toward a common sense understanding of what ought to count, in this case, as an image.

My intuitions tell me an image is a representation – where a representation is the result of an attempt to capture the salient features of an object, scene, state of

[1] This thesis is being developed in a book length manuscript under construction entitled tentatively, *Seeing Near and Far, A Heraclitian Philosophy of Science and Technology*.

affairs, or idea, etc. Fortunately or unfortunately, what constitutes a salient feature is a function of the person or persons constructing the image. As a first pass, consider the following items as images:

Sculptures
Photographs
Portraits
Still lives
Landscapes
Various kinds of drawings
Motion pictures—both animated and “realisitic”
Visualizations inspired by poetry
Visualizations inspired by music
Performed Plays
Performed Operas
Performed Ballet and interpretive dance

If we accept the fact that these are images, then a Picasso such as the Gernica counts as an image, but it would seem that a Jackson Pollack does not only insofar as it is unclear what a Pollock is supposed to represent.[2] This entails declaring that to be an image is to be representational. But it says nothing about what makes something representational. That said, nevertheless, it is not shocking to note that not all paintings are images, where a painting is nothing more conceptually complicated than paint deliberately applied to a surface. But, if it is true that not all paintings are images, especially when they are not representational, have we not found a way into our topical question, when is an image not an image? It looks like we could reasonably say that an image is not an image when it is not representational. On the other hand, doesn’t that just beg the question? After all, it isn’t at all clear that for an image to be an image it must be an image of something. When you think about it, on the one hand, it seems arbitrary to demand that images be representational, but, on the other hand, to do so seems to beg the question. For example, consider the following as candidates for being added to the list above.

Diagrams
Flow charts
Data tables

The interesting feature of these sorts of things is that while they are not representational, they do convey information in visual form. For, on the surface at least, it seems as if these forms of images have a different semantics than written language. The important point however, is that they do seem to have a semantics,

[2] If turning to art is seen as somehow cheating, it is important to remember that the creation of images began in art.

for they do manage to convey information. The unresolved problem that remains for us is how to determine if the image is an accurate representation. So, if we accept this approach, then one answer to our question is that an image is not an image when we do not know if it is representational but conveys information none-the-less.[3] With your permission, let's accept that for the time being as a first pass.

However, that just moves us back one step, for now we can re-ask the question that our quick look at electron microscopes motivated: when is an alleged representation a representation? The point here is epistemological.

I think it not too radical to suggest that seeing is a complex activity in which after learning to see that as a tree or as a car, we forget that we had to learn that. In our mature state we see the world around us and assume we see it for what it is. That is why philosophical questions like "but are you seeing what is really there?" seem so silly. But, on reflection, we also understand that seeing is an interpretive process and that we bring to our seeings a load of background information and experience. Elsewhere I have argued that to call it a seeing by way of images generated by an electron microscope is a metaphorical extension of our common sense notion of seeing (Chapter 17). But, I have now come to realize that there is a lot involved in appealing to metaphor here. If we unpack it, as I would like to start to do here, we can see that to understand through metaphor is to do a number of things at once. First, we use metaphor to access what is new and different because in a metaphor we take what we know and apply it to the unknown and say that the unknown is like the known in these various ways. It makes the new seem familiar and approachable, usually. Sometimes, as in the example of the tennis gun above, it makes the unknown or the new seems even stranger than we first thought. Second, when using metaphor to make the new and unknown approachable, we are also asked to accept that certain things that we do not really understand are reliable. Metaphors tell you this is like that in certain limited ways, and by the way, just accept that everything else is working just fine, however that happens. In the case of the electron microscope, when asked to accept what it produces as a representative image, we are also asked to accept the fact that the assumptions built into the manner in which that image is constructed are correct and reliable. To use the language of science studies, we black-box the process and merely look at the result. But to call the image created by the electron microscope an image is to ask us to accept in some fundamental way that the science is sound and the technology (programming?) reliable and the people manipulating it reliably are honest.

But, I suggest, this ought to be a lot to ask. What is interesting is that it appears that it is not. It is a measure of the success of the scientific establishment that we, the general public, tend to accept claims based on the use of increasingly

[3] Yes, "information" is not defined. But, I suggest, we have to start somewhere. If we succeed in making progress by proceeding in the manner suggested we can always return and fine-tune the argument by going deeper into concepts like "information". Call this approach "conceptual boot-strapping".

complicated instruments working in the realm of the frontiers of science with increasing readiness. That is, the more complicated the science and the more simplified the public explanations, the more readily we tend to accept those fantasies. That is why it is important to know what really happens in an electron microscope before buying into the claims with which it is associated. Before I explore what that ominous sounding remark is supposed to suggest, let me give you just one example of the kind of phenomenon to which I am referring. I think we are all in awe of the images sent to us by the Hubble Space Telescope. The ones of the horse head and crab nebulae are just breathtaking – and the colors are truly inspiring – just one catch – the colors are computer generated. When I tell my students that, the looks on their faces resemble the one when they learned that there is no Santa Claus. What got me going in this direction was a presentation at the Conference, "Discovering the Nanoscale" at Darmstadt in October 2003 that revealed that the picture of the nano-scale IBM was not just constructed through the assistance of computers, but it too was computer enhanced – with the colors added, for example. This, it turns out is a pervasive problem; even the choice to use grey scale is a decision to create the image in a certain way. So when we say of an image that it must convey information, should we not also be asking (1) whether there is a claim that reality is being representing, and (2) is the image presented of something real or imagined? Perhaps, then, should we not be asking this slightly different question: "When is an image not an imagining?"

The issue here is both epistemological and ethical. The epistemological issue concerns, for lack of a better term, noise. We are familiar with the problem of filtering out noise when searching for an identifiable signal. The problem is multifaceted: what to filter out and on what criteria, what to amplify, to what degree, etc. The problem with color-enhancement and sharpening up of nano-images is that we don't yet know what is important and what is not. Further, the problem may become intractable since we do not have a god's eye view from which to determine if we have it right. In a certain sense then the problem here is an in principle lack of access, or to put it differently, a case of very strong underdetermination. But is this really a problem? We have in-principle-lack-of-access to many astronomical events, like the big bang, and we still claim to know a lot about the early universe. We have images from the Hubble of far distant galaxies that we can never get close to in person, and yet we can still understand a lot of what is going on there – or so we think.

My worry is that, unlike the "images" from the Hubble, we have relatively little experience in enhancing the images produced by STEMs. We have ways of checking up on the Hubble images. For example, we can experiment with filters and use smaller telescopes here on earth to check out their effect when we look at mountains or trees. However, although we have lots of experiences with so-called images from STEMs – we do not have such successes in fixing them up. This is, in a curious way, a new version of the what-are-we-going-to-do-when-we-stain-a-specimen-that-we-are-going-to-examine-under-a-standard-microscope problem (see Chapter 17). Computer enhancement of images is fun, especially with all the nifty colors we can use. But is it producing an honest replication of the object/surface in question? Clearly not, and that raises the ethical issues.[4]

The ethical issues arise in two forms: strong and relatively minor. The relatively minor issues have to do with the relationships between science and the public. For example, we are misleading the public when we fail to disclose fully what we are doing when we computer enhance our electron microscope constructed images. The strong ethical issues center around the fact that these images raise false expectations. Among them, that we know more than we do. The presentation of these beautiful pictures suggests in a very strong way that this is indeed what it is like out there, in there. But more importantly, they mislead in crucial ways. The beautiful computer simulations we see of nano interactions are not only beautiful simulations, they are also almost heart-stopping in their ability to feed the hubris we sometimes exhibit when employing the newest technological toys, computer and advanced programming techniques, among them. Please do not get the wrong impression – I am not suggesting that we should not employ the latest technologies in science. What I am talking about is the illusion we create not just in the general public but sometimes in the practicing scientific community. The illusion is that we know more than we really do. Never underestimate the ability of human beings for self-delusion. These computer generated and enhanced pictures suggest that the world is at rock bottom a simple place. It can be pictured as individuals atoms resting on stable fields that we can manipulate at will, twirl them, enlarge and narrow them, put them to music, make them dance, when in fact nothing of the kind is the case. The world at the nano and quantum mechanical level is a buzzing, shifting, constantly in motion in non-linear and non-classical causal fashion.

This is all heading in one direction. It is not just misleading to suggest that the world is simple at the bottom. It is epistemically suspect. It employs a crucial but faulty assumption. It is the assumption that the world is better understood if we simplify our presentations of it. I humbly suggest that this is wrong-headed. It may in fact be helpful to extract some feature of the world, color it pretty non-natural colors and play with it. But it is more important to put that heuristically altered item back into the buzz and try to understand it in that environment, its "natural" environment. Most importantly it is crucial that we explain to the public and our colleagues the purpose of the heuristic move and what it reveals about what is really going on at the bottom.

So what is wrong with simplification? It suggests that we know more than we do and, crucially, that we can do more than we can. The scientific community has

[4]The "Clearly not..." might be considered contentious, but with a little expansion, I believe it will be obvious. Consider, for example, that the surface on which nano scale objects exist is at the interface between the quantum domain and the atomic. We have no idea how to visually represent what happens in the quantum domain, so we cannot say we are accurately representing the surface on which the atomic structures we are picturing sit. If we cannot claim to be accurately depicting the surface, then how can be sure of the space in which nano structures function, and if that is uncertain, so must be our representation of the nano structures themselves.

done a good job of convincing the public that it has god-like properties – but this situation presents a double-edged sword; the public feeds on gods that fail. Be honest about the mess and you will repeat positive rewards. Further, it is not the simplicity of the universe that makes it the object of our enquiry, it is the complications, the unanswered questions, the mess of it all. The more we look, the more complicated we find it to be. If you cuddle the public and give them simplicity and then in the crunch, when, for instance, in the hospital, you say, well it is more complicated than that, then you will have failed miserably. I love the pictures, but they are not representations. They are heuristic imaginings, extended metaphors, if you will, and they should be recognized as such and treated that way. How will that affect the way in which the work of science is perceived? My guess is that it will enhance it. Doing science is hard work. The public should know that and when they do the successes of science will be all the more appreciated. Telling the truth is also hard.

To conclude, let me summarize. The question is "in what sense is a STEM computer generated picture of nano structures an accurate representation of what is there?" Following some discussion of how "seeing" using a STEM involved a metaphorical extension of the concept of "seeing," it was argued that to be a representation the image must convey information. The problem is in understanding what information is being conveyed, since we cannot directly access the domain that we are purporting to represent. The problem is not that we do not know how to interpret what is presented to us as an image, but, rather, that we have loaded the creation of the representation ahead of time without being able to know if our guess that this is what the STEM and its fellow traveler computer programs are producing is an accurate picture of what is really there. The reason why there is so much discussion of when an image is an image is that this really is a question of whether or not the image that is produced is an accurate portrayal of something that is really there or a mere fabrication.

Consider one last attempt to convey a sense of the magnitude of the problem. If we do a random sample of some domain and then plot the results in three dimensions, assuming that is sample is truly random and that there is no natural clumping of the data, which curve is the correct one? We can draw an infinite number of curves through those data. Without an independently certified decision procedure for selecting the correct curve we are simply left with the data. The problem is further complicated by the fact that there are ethical dimensions. (1) To say that this is what is taking place at the nano-level, is to lie, since we don't, in fact, know that to be the case. (2) To present these standard, nicely colored, enhanced, and simplified pictures as genuine representations of what is going on at the nano-level is to claim falsely that nature is in fact simple and clean and neatly colored at that level. But, nature is not neat and tidy at that level. To suggest otherwise is to mislead by way of making it appear that there are simple answers to very complex problems. That approach gets us into trouble at the political level and it should get us into equally big trouble in our epistemology.

Acknowledgement My thanks to Thomas Staley for his assistance. All mistakes are my own.

References

Rasmussen, N. and Hawkes, P. 1998. “Microscope, Electron”. In Bud R. and Warner D.J. eds., *Instruments of Science*. New York, NY: Garland Publishing.
Pitt, J.C. 2005. “The Epistemology of the Very Small”. In Nordmann A., Schummer J., and Baird D. eds., *Discovering the Nanoscale*. Amsterdam: IOS Press.

Chapter 19
Small Talk: Nanotechnology and Metaphor

19.1 Introduction

The general topic I am addressing concerns the epistemological role of the use of metaphor in the philosophy of science. More specifically, I am concerned with the role metaphor plays in scientific and technological change. In the case in point, nanotechnology, I will explore the role of metaphor in changing our conception of the confirmation of the plausibility of theoretical notions. The basic idea is that metaphors either offer or suggest images that are meant to persuade one to change one's belief.[1] Thus the confirmatory role is variable.

For a while now, I have been arguing against the tradition of perennial philosophy – the claim that there are universal, timeless questions to be solved and that there are in fact universal and timeless answers for these questions. In place of the perennial philosophy with its search for absolutes, I offer a Heraclitean philosophy of science that embraces change and seeks to understand the mechanisms behind changes. The goal is to produce a story that is itself a coherent account of what happened and why – helping us to understand why things are as they are at the present. This approach is rooted in history, but it rejects narrow or narrowly confined case studies in place of tracing and exploring scientific and philosophical problematics.[2]

A problematic is a source of worry and fascination that extends over a long period of time. Problematics are mostly linear clusters of problems. They come into being in a variety of ways – sometimes accidental, sometimes evolving out of other problems. The content may change as various individuals wrestle with it and as new and innovative technologies emerge, but there is a clear causal/social chain of events

Originally appeared in *Spontaneous Generations,* Vol. 2. Reprinted by permission of the editor.

[1] By "metaphor" I mean the use of a term with which we are familiar to characterize features of an object or situation with which we are not. Thus we talk about electrons orbiting the nucleus of an atom, having appropriated the term "orbit" from astronomy.

[2] I first developed the notion of problematics in my 2007. The concept was elaborated in my 2001. (The first paper was originally presented at a conference in 1995 honoring Majorie Grene on the occasion of her 85th birthday, but the volume did not appear until 2007.)

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4_19,

that can be traced. For example, when Galileo turned the telescope toward the heavens in 1609 and saw moons orbiting Jupiter, those observations eventually threw the entire notion of an earth-centred universe into doubt. Galileo's observations allowed him to make the negative argument against the Aristotelian view, but he could not provide a replacement physics to make a positive argument for heliocentrism.[3] His persistence in the face of controversy is remarkable, as are the social, political, and mathematical developments that form the context for his story, but what is at issue here is what Galileo offers in place of a physics: metaphor.[4]

To understand why Galileo was therefore willing to risk his career and life in promoting a new cosmology and scientific method, one has to understand the historical context. On this view then, the phrase "the Scientific Revolution" (contrary to certain efforts to deconstruct the phrase) identifies a problematic that takes over 100 and 50 years to come to some sort of resolution. There was a lot going on during that time period and to understand that requires uncovering a lot of dots and then finding a way to connect them. Sometimes connecting the dots requires understanding where they came from, or, in other words, the influence of this or that on this person or that person. The history gets messy, which history is, and the story is constantly being revised as we learn more and more, which is as it should be. But the more we discover about why something happened and what its impact was or is, the more we should be amazed by the complexity of the dance and the fact that we can actually trace out a sense of progress. I would venture to guess that understanding complexity will give us a greater sense of intellectual security than basking in the light of false simplicity. But all of this concerns Galileo and on the surface may not appear to an assessment of the epistemic value of metaphor per se. So, let's move on with the understanding that we will return to these issues later.

19.2 An Example of an Early Use of Metaphor to Facilitate Theory Change

Eventually I want to look at the role of metaphor in extending knowledge, specifically the role of metaphor in the development of a language to talk about what is going on or what might be possible at the nanoscale. To place this in a context, I want to look at two examples that rely heavily on multiple metaphors to ascertain what, if any, light they ultimately shed on the vexing question of whether self-replicating nano-robots are possible. This is an important question because the possibility of self-replicating nanobots has also raised the spectre of run-away technology and all sorts of dire consequences.

[3]See my 1992 for a fuller discussion of the structure of the Dialogue and Galileo's efforts to circumvent the problems he encountered defending Copernicus' view. While it is true that Galileo did not have a full physics, he did attempt to use the concepts of circular and relative motion to make his point. They simply weren't up to the job.

[4]My thanks to Isaac Record for this formulation. I could not have said it better.

In Chapter 18 I argued that we use metphor to help us understand new theories and hypotheses about the world, when what is new is sufficiently novel and different from that with which we are familiar that, on the surface, it may appear impossible to gain access or understanding. Consider again, for example, Galileo's defense of the heliocentric theory. Two things to keep in mind here:

Neither Copernicus nor Galileo had a physics to explain how the earth could move – they were offering, at best, a different description of the structure of the solar system. There were many reasons in support of the new theory, but they were as much aesthetic as anything else. We will return to this issue below.

Replacing the geocentric view with the heliocentric one was for many also impossible to comprehend. Putting the physics aside, the centrality of the earth was absolutely crucial to the understanding educated folks in Western Europe at that time had of who they were. To see this you need to couple Aristotle's physical argument for the earth' central location with the theology deriving from Genesis, Book 1. From Aristotle we get both centrality and the sphericity of earth – earth is the heaviest of the elements – all four elements seek their natural place – earth, after the chaos coalesces at the centre since the heaviest element falls downward – and if earth falls downward from all possible directions, a sphere is the natural result. From Genesis we are told that God created Man in His image and gave him dominion over the earth. Put the two together and you find Man, in God's image, at the centre of God's universe with dominion over the earth – hence man is the centre of God's attention.

Move the earth out of the centre and it all falls apart. There is no evidence that God sees Man in any special light – the earth is just one of many planets – nothing special there, etc. In the immortal words of Anna Russell (1991), "it all falls down." Aristotle's explanations of why the earth is where it is and has the shape it does are now inadequate.

So, what does Galileo do? He doesn't have a physics to explain how the earth can move. Somehow, however, he has to reassure people that they have not been rendered irrelevant. To overcome both problems he does two things:

He introduces a language that all the contestants use and can understand, giving them the illusion of sharing a common ground. That language is geometry. Most of Aristotle's physical arguments employ a form of conceptual argument: if this means this then this follows. Surprisingly, in his *Dialogue on the Two Chief Sy tems* Galileo does not try to refute Aristotle's account. He is much more subtle. Rather, he shows that Aristotle's arguments could be better put in the language of geometry.

Consider his opening volley. He starts by giving us Aristotle's familiar proof that the world is perfect in an attempt to lay the Peripatetics' grounds for refuting Copernicus's claim that the earth moves.

> For, he (Aristotle) tells us, it is not a mere line, nor a bare surface, but a body ha ing length, breath, and depth. Since there are only these three dimensions, the world, having these, has them all, and having the Whole, is perfect. (p. 9)

Galileo then continues:

> To be sure, I much wish that Aristotle had proved to me by rigorous deductions that simple length constitutes what we call a line, which by the addition of breadth becomes a surface; that by further adding altitude or depth to this there results a body, and that after these three dimensions there is no passing farther – so that be these three alone, completeness, or, so to speak, wholeness is concluded. (pp. 9–10)

Here Galileo is simply appealing to the principles of Euclidean geometry and shortly he actually draws a diagram to illustrate his point. Throughout the rest of *Dialogue* he resorts to this technique – making the diagram, employing Euclidean geometry, the proof of one point after another.

The second move is simply ingenious. With no celestial physics at his command, he nevertheless offers as his crowning argument a pseudo geometric explanation that the earth moves. It is both a proof by elimination and a conditional proof. He argues that only by assuming the earth moves can we explain why there are tides – a most contentious problem at the time. As the earth turns on its axis the seas slosh the same way water in a barge does as the barge accelerates and slows down. That is the punch line – but the argument is loaded with geometric diagrams and lots of handwaving.

Galileo has compounded two techniques here – first, over the course of the book he has essentially legitimized appeals to geometry as a lingua franca. Second, he then turns to the metaphor of the barge carrying water from the mainland to Venice and the way the water sloshes back and forth to make his case for the tides, hence the earth moves, QED. The idea here is to extend something with which Galileo's contemporaries are familiar, the sloshing of water in a barge, to the unfamiliar, the earth in motion.

Does the proof really work? No. But the metaphor of the barge is highly compelling. It was readily understood by the educated lay person of the time. Using metaphors or analogies to make a point was fairly common at the time. Coupled with Galileo's pandering to his audience, convincing them that they *can* understand these sometimes complex geometric proofs – he both offers an explanation and reduces their fears. If we try to understand the epistemic punch of the metaphor of the barge, we have several options. Either Galileo offers a proof that compels us to change our minds on rational grounds or he is doing something else. Clearly it is something else, for, as already noted, the so-called proof fails. I am willing to argue, at this point, that, at best, the combination of the geometric machinations and the barge metaphor can be considered as an attempt at persuasion. On the other hand, Galileo probably believed he was doing more – this theory of the tides having been dear to him since 1595 (when he probably stole it from his friend Scarpi). But what he thought he was accomplishing and what he did accomplish are two different things.

19.3 NANO – Do Mixed Metaphors and a Lot of Mathematics Constitute a Proof?

A lot of scary stuff has been written about what could go on at the nano scale and the implications for us'ins here up top. One of the most controversial ideas put forth is

Eric Drexlers' idea of self-replicating nano-robots – nanobots, which can in itself be considered a metaphor. Fear runs rampant especially among technophobes – since Drexler's nanobots are to be self-programming. What if these things gain some sort of autonomy and take over the world?! In a two page article in September 2001s *Scientific American* (a fairly safe place to offer this argument since its readership is fairly pro-science and technology), Nobel laureate Richard Smalley tried to put those fears to rest. He is not offering a defense of nanotechnology. Rather, he attempts to show that fears based on Drexler's scenario are unwarranted since Drexler's scenario can't occur according to Smalley. While his strategy does not completely map onto Galileo's, there are enough similarities to give us reason to take a close look at his argument and to see if he manages to achieve crowd control. Let me be clear from the start, I am on his side on this one – but I also want to make sure the case has really been made. What I will conclude is that Smalley and Galileo both employ what can be described as a political strategy designed more for persuasion that getting at the truth. Smalley's account runs roughly like this:

1. *love is like chemistry* (metaphor #1) – put two people together and there is a product that results. This is a case of intrinsic affinities.
2. chemical reactions, however are actually a lot more complicated than love, i.e., it is more complicated than just putting two atoms together and seeing what emerges. In the space, roughly one nanometer, within which the appropriate atoms are supposed to react, there are, in fact, many atoms, 12 to 15 – engaged in a *three dimensional waltz* (2nd metaphor).
3. one nanobot would not be useful – generating even a tiny amount of a product would take a solitary nanobot millions of years – lots of math now follows which I have to take on faith:
4. "Making a mole of something – say 30 g, or about one ounce – would require at least 6×10 to the 23rd bonds, one for each atom. At the frenzied rate of 10 to the 9th per second it would take this nanobot 6×10 to the 14th seconds that is, 10 to the 13th minutes – which is 6.9×10 to the ninth days, or 19 millions." Pretty slow, eh? Pretty much like Galileo's proof for the tides, nay?
5. But, if the nanobot could replicate itself and then if the two could replicate themselves we could have an *army* (3rd metaphor) of nanobots at our command and then they could work together and increase the rate of production and maybe the world of plenty *would* be possible.
6. But they could then also *mutate* (4th metaphor) and get out of control reducing the world to *an undifferentiated mass of grey goo* (5th metaphor) or simply eliminate the need for humans.
7. BUT – this is not possible for two reasons, given the already mentioned small space in which atomic reactions occur: a. *fat fingers* (metaphor 6) and b. *sticky fingers* (metaphor 7).
8. Fat fingers: because the arms of a nanobot must itself be made of atoms, there is an irreducible size problem. "There just isn't enough room in the nanometer-size reaction region to accommodate all the fingers of the manipulators necessary to have complete control of the chemistry."

9. Sticky fingers: "The atoms of the manipulator arms will adhere to the atom that is being moved. So it will be impossible to release this miniscule building block in precisely the right spot."
10. In conclusion, Smalley returns to his love theme and also to his waltz theme neatly tying his first two metaphors together and then relating them to the nanobot theme: "Like the dance of love, chemistry is a waltz with its own step-slide-step in three-quarter time. Wishing that a waltz were a merengue – or that we could set down each atom in just the right place – doesn't make it so."

Well, does Smalley's argument work? We have several metaphors at work – seven at least – and they don't obviously compliment each other – love, waltzes, nanobots, army, fat fingers and sticky fingers. The fingers can work together – love and waltzes are a stretch, and "army" stands out like a sore thumb. Let's spell out two of the metaphors in a bit more detail, love and sticky and fat fingers.

The love metaphor has to do with the way in which atoms are attracted to one another. According to Smalley it is not as simple as putting two people together. Put two people together and love develops. Put two atoms together and it isn't clear that anything will happen. Chemistry is about atomic affinities, some atoms bond more readily than others. But isn't that also the case with people? Just putting two people together doesn't guarantee love. On the other hand, Smalley's second point is that you just don't find two atoms in a space, you find several. Further, these several atoms are engaged in a waltz. He doesn't explain this, but what he seems to be suggesting is that the situation at the nanolevel is fairly unstable – the atoms are in constant motion, therefore no nanobot can be constructed. However, this is not an argument, it is a rhetorical sleight of hand suggested by the waltz metaphor.

If we turn to the metaphor of fat and sticky fingers we find more of the same. I *think* I know what he means by both. I try to play the piano. Often I just can't get my fingers between two black keys and I mutter "fat fingers". To me, "fat fingers" means there isn't enough room for me to get my fingers in there. For Smalley, "fat fingers" means something similar,[5] essentially he is saying that whatever implement we use to manipulate the atoms at the nano level, it will be too large to fit the area and therefore we will not be able to control the atoms in the way we think we can. My problem with this argument is that it is not clear why the area is so confined. Smalley has not produced the rationale for thinking that we necessarily have the fat fingers problem.

When it comes to sticky fingers we face a more difficult issue. The "sticky fingers" metaphor is supposed to convey the following: atoms have a tendency to stick to one another; therefore when we use an implement to pick up an atom and deposit it elsewhere the atom will stick to the implement. It will not let go and, hence, we cannot really put it down, get another atom and put it next to it, continuing to build something. There are two problems here. First, this seems to go against

[5]This is the beauty of using metaphors; they rely on what you think you understand it to mean based on your own experience, which may not be exactly what the author intended.

the chemistry-isn't-love theme of the first Smalley metaphor. Second, we seem to have empirical evidence that shows we can in fact pick up an atom and deposit it where we want it. This evidence comes in the form of a picture two IBM researchers published showing the name IBM spelled out atom by atom.[6]

Then we have the razzle-dazzle quick math computation which I see as an effort both to prove his point about the amount of time it would take for a nanobot to accomplish anything, and as a way of establishing the scientific nature of the discussion – since the whole point of talking about chemistry in metaphorical terms is to make non-chemists think they understand what is going on – we also clearly need to establish our scientific credentials. The best way to do this is to appeal to the language of science: mathematics – even if we can't refute it, we know it is math and it therefore must be good. The problem is that as with many metaphors, it becomes very difficult to break the argument apart – each metaphor supports, as it were, a different premise and in each case this can, perhaps, be seen as the argument for the premise. It is not always a very good argument, as we have seen, but it can be persuasive.

So, does the parallel hold up between Galileo and Smalley? Well, they both start with the familiar – Aristotle's method of proving things on the one hand and love, on the other hand. Then they both move to mathematics. Finally, they turn to new metaphors – barges and fat fingers. Galileo's geometrical argument for the tides fails – Smalley's math is probably correct, but on the key points he reverts completely to metaphors. It is not clear that Smalley's metaphors are strong enough by themselves to carry the day – consider the sticky fingers claim – if he is correct that we won't be able to place atoms where we want them, then how do we a count for the IBM image? Likewise for the fat fingers claim. But it is unclear as to whether Smalley is appealing to genuine technical difficulties or engaging in a rhetorical strategy.

In closing we find ourselves with more problems than we can solve. Can metaphors be falsified by empirical data? Can metaphors function successfully as arguments? Galileo's metaphor of the barge fails, not because of empirical evidence, but because he just doesn't have a physics to support extending the metaphor to the earth. In Smalley's case it is not clear. While the IBM image seems to suggest that at least the sticky fingers metaphor fails, we have the problem of asking if the construction of a set of letters is in any way similar to the construction of nanobots and their hoped for ability to construct various products. What seems clear to me is that to make headway on Smalley's "refutation" of Drexler, we need to understand the epistemic value of metaphor, what the criteria are for evaluating their rhetorical power, and how these criteria change over time. When we spell all this out we will find that our older notions of what constitutes evidence and confirmation must change in light of that new understanding. Philosophers for years have resisted

---

[6]In 1990, D.M. Eigler and E.K. Schweizer published a paper in Nature that included a photo of an image spelling out IBM, proving that atoms could be deliberately manipulated to form specific configurations using a scanning tunneling microscope.

crediting any genuine epistemic impact to rhetorical devices – with Popper rejecting a logic of discovery in which metaphors play a crucial role and the positivists concentrating solely on the logic of allegedly universal and unchanging concepts. This resistance is blatant despite the fact that philosophers have used metaphors as far back as Heraclites (stepping in the same river twice) and Plato (the cave). But if there is an epistemology of metaphor, those positions must be re-examined. This is not beating a dead horse. Well, yes, at least these two horses, Popper and Positivism, are dead. It is the perennial philosophy I am after. What STS has forced Philosophy of Science to recognize is that there is a changing plurality of epistemic concepts appropriate to the analysis of science and this enlarges the philosophical field of analysis. In this sense the problematic of devising an appropriate normative epistemological account of science must change to accommodate rhetorical devices scientists as well as philosophers use and the contexts in which they use them. Thus, sometimes the audience is other scientists – as in Smalley's Scientific American article, and sometimes the educated public, as in Drexler's book – which yields the interesting question of whether an undermining metaphor needs to be directed to the same audience to be effective. That is, does Smalley's metaphors of fat and sticky fingers appeal to or even make sense to the folks who reacted to Drexler's nanobots? This, interestingly enough, is a philosophical question that needs empirical work.

In closing I want to end with a simple question: is the philosophical problematic regarding scientific epistemology changing by accommodating the social, and if so, what does this mean for the perennial philosophy? My guess is that it is and it strikes a mortal blow to the misguided search for eternal answers.

References

Eigler, D.M. and Schweizer, E.K. 1990. "Positioning Single Atoms with a Scanning Tunneling Microscope." *Nature*, 344: 525–26.

Pitt, J.C. 2001. "The Dilemma of Case Studies". *Perspectives on Science: Historical, Philosophical, Social*, 9(4). 373–82.

Pitt, J.C. 1992. *Galileo, Human Knowledge, and the Book of Nature*. Dardrecht: Kluwer.

Pitt, J.C. 2007. "Seeing Nature: Origins of Scientific Observation". In Burian R.M. and Gayon J. eds., *Conceptions de la Science: Hier, Aujourd'Hui et Demain*, pp. 272–89. Brusselis: Ousia.

Russell, A. 1991. "The Ring." On *The Anna Russell Album*.

Index

J.C. Pitt, *Doing Philosophy of Technology*, Philosophy of Engineering and Technology 3, DOI 10.1007/978-94-007-0820-4,